

Oracle® Grid Infrastructure

Installation Guide

11g Release 2 (11.2) for Linux

E10812-03

August 2009

Oracle Grid Infrastructure Installation Guide, 11g Release 2 (11.2) for Linux

E10812-03

Copyright © 2007, 2009, Oracle and/or its affiliates. All rights reserved.

Primary Author: Douglas Williams

Contributing Authors: David Brower, Jonathan Creighton, Barb Lundhild, Saar Maoz, Markus Michalewicz, Hanlin Qian, Janet Stern

Contributors: Aiman Al-Khammash, Troy Anthony, Ram Avudaiappan, Mario Barrientos, Mark Bauer, Yuki Feng, Mark Fuller, Barb Glover, Fred Glover, Paul Harter, Yingwei Hu, Masakazu Ito, Mayumi Hayasaka, Prakash Jashnani, Sameer Joshi, Sana Karam, Roland Knapp, Erich Kreisler, Jai Krishnani, Shashi Lanka, John Leys, Karen Li, Rich Long, Scott Lu, Dave Marshall, Greg Mengel, Manami Miyayama, Anil Nair, Van Okamura, Stefan Pommerenk, Srinivas Poovala, Justin Ren, Sudhe Sampath, Shachi Sanklecha, Vishal Saxena, Janelle Simmons, Akshay Shah, Ara Shakian, Vivian Schupmann, James Spiller, Malai Stalin, Dipak Saggi, Bob Thome, S. Matt Taylor, Preethi Vallam, Rick Wessman, Jim Womack, Warren Xu, Sivakumar Yarlagadda, Min Yu, Yunzhou Zheng, Ricky Zhu

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

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software 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 which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on 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. 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.

Contents

Preface	ix
Intended Audience.....	ix
Documentation Accessibility	ix
Related Documents	x
Conventions	xi
What's New in Oracle Grid Infrastructure Installation and Configuration?	xiii
Desupported Options	xiii
New Features for Release 2 (11.2)	xiii
New Features for Release 1 (11.1)	xvi
1 Typical Installation for Oracle Grid Infrastructure for a Cluster	
1.1 Typical and Advanced Installation	1-1
1.2 Preinstallation Steps Completed Using Typical Installation	1-1
1.3 Preinstallation Steps Requiring Manual Tasks.....	1-1
1.3.1 Verify System Requirements.....	1-2
1.3.2 Check Network Requirements.....	1-2
1.3.3 Check Operating System Packages	1-4
1.3.4 Create Groups and Users.....	1-4
1.3.5 Check Storage.....	1-4
1.3.6 Prepare Storage for Automatic Storage Management.....	1-5
1.3.7 Install Oracle Grid Infrastructure Software	1-6
2 Advanced Installation Oracle Grid Infrastructure for a Cluster Preinstallation Tasks	
2.1 Installing the Linux Operating System	2-1
2.1.1 Completing a Minimal Linux Installation.....	2-2
2.1.2 Completing a Default Linux Installation.....	2-2
2.1.3 About the Oracle Validated Configuration RPM.....	2-3
2.1.4 Installing the Oracle Validated Configuration RPM	2-4
2.2 Reviewing Upgrade Best Practices.....	2-4
2.3 Installation Fixup Scripts	2-5
2.4 Logging In to a Remote System Using X Terminal.....	2-6
2.5 Creating Groups, Users and Paths for Oracle Grid Infrastructure.....	2-7
2.5.1 Determining If the Oracle Inventory and Oracle Inventory Group Exists.....	2-7

2.5.2	Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist	2-8
2.5.3	Creating the Oracle Grid Infrastructure User.....	2-8
2.5.4	Creating the Oracle Base Directory Path.....	2-11
2.5.5	Creating Job Role Separation Operating System Privileges Groups and Users.....	2-12
2.5.6	Example of Creating Standard Groups, Users, and Paths.....	2-19
2.5.7	Example of Creating Role-allocated Groups, Users, and Paths.....	2-20
2.6	Checking the Hardware Requirements	2-21
2.7	Checking the Network Requirements.....	2-23
2.7.1	Network Hardware Requirements.....	2-23
2.7.2	IP Address Requirements.....	2-24
2.7.3	DNS Configuration for Domain Delegation to Grid Naming Service	2-26
2.7.4	Grid Naming Service Configuration Example	2-26
2.7.5	Manual IP Address Configuration Example	2-27
2.7.6	Network Interface Configuration Options.....	2-28
2.7.7	Enabling the Name Service Cache Daemon	2-29
2.8	Identifying Software Requirements	2-29
2.8.1	Software Requirements List for x86 (32-bit) Linux Platforms	2-30
2.8.2	Software Requirements List for x86-64 Linux Platforms	2-33
2.9	Checking the Software Requirements.....	2-39
2.10	Network Time Protocol Setting	2-40
2.11	Installing the cvuqdisk Package for Linux.....	2-41
2.12	Enabling Intelligent Platform Management Interface (IPMI).....	2-42
2.12.1	Requirements for Enabling IPMI.....	2-42
2.12.2	Configuring the IPMI Management Network.....	2-42
2.12.3	Configuring the IPMI Driver	2-42
2.13	Automatic SSH Configuration During Installation	2-47
2.14	Configuring Grid Infrastructure Software Owner User Environments	2-47
2.14.1	Environment Requirements for Oracle Grid Infrastructure Software Owner.....	2-48
2.14.2	Procedure for Configuring Oracle Software Owner Environments	2-48
2.14.3	Setting Resource Limits for the Oracle Software Installation Users.....	2-50
2.14.4	Setting Display and X11 Forwarding Configuration.....	2-51
2.14.5	Preventing Installation Errors Caused by stty Commands.....	2-51
2.15	Requirements for Creating an Oracle Grid Infrastructure Home Directory.....	2-52

3 Configuring Storage for Grid Infrastructure for a Cluster and Oracle Real Application Clusters (Oracle RAC)

3.1	Reviewing Oracle Grid Infrastructure Storage Options.....	3-1
3.1.1	Overview of Oracle Clusterware and Oracle RAC Storage Options	3-1
3.1.2	General Storage Considerations for Oracle Grid Infrastructure and Oracle RAC.....	3-2
3.1.3	Supported Storage Options	3-3
3.1.4	After You Have Selected Disk Storage Options	3-4
3.2	Shared File System Storage Configuration.....	3-4
3.2.1	Requirements for Using a Shared File System	3-5
3.2.2	Deciding to Use a Cluster File System for Oracle Clusterware Files	3-7
3.2.3	Deciding to Use Direct NFS for Data Files.....	3-7
3.2.4	Deciding to Use NFS for Data Files.....	3-9
3.2.5	Configuring Storage NFS Mount and Buffer Size Parameters.....	3-9

3.2.6	Checking NFS Mount and Buffer Size Parameters for Oracle Clusterware	3-10
3.2.7	Checking NFS Mount and Buffer Size Parameters for Oracle RAC.....	3-10
3.2.8	Enabling Direct NFS Client Oracle Disk Manager Control of NFS.....	3-12
3.2.9	Creating Directories for Oracle Clusterware Files on Shared File Systems	3-13
3.2.10	Creating Directories for Oracle Database Files on Shared File Systems.....	3-14
3.2.11	Disabling Direct NFS Client Oracle Disk Management Control of NFS	3-15
3.3	Automatic Storage Management Storage Configuration.....	3-16
3.3.1	Configuring Storage for Automatic Storage Management.....	3-16
3.3.2	Using Diskgroups with Oracle Database Files on ASM	3-28
3.3.3	Configuring Oracle Automatic Storage Management Cluster File System (ACFS)	3-29
3.3.4	Migrating Existing Oracle ASM Instances	3-30
3.3.5	Converting Standalone Oracle ASM Installations to Clustered Installations.....	3-31
3.4	Desupport of Block and Raw Devices.....	3-31

4 Installing Oracle Grid Infrastructure for a Cluster

4.1	Preparing to Install Oracle Grid Infrastructure with OUI	4-1
4.2	Installing Grid Infrastructure.....	4-6
4.2.1	Running OUI to Install Grid Infrastructure.....	4-6
4.2.2	Installing Grid Infrastructure Using a Cluster Configuration File.....	4-7
4.3	Installing Grid Infrastructure Using a Software-Only Installation	4-7
4.3.1	Installing the Software Binaries	4-7
4.3.2	Configuring the Software Binaries	4-8
4.4	Confirming Oracle Clusterware Function.....	4-10
4.5	Confirming Oracle ASM Function for Oracle Clusterware Files.....	4-10

5 Oracle Grid Infrastructure Postinstallation Procedures

5.1	Required Postinstallation Tasks.....	5-1
5.1.1	Download and Install Patch Updates	5-1
5.2	Recommended Postinstallation Tasks	5-2
5.2.1	Back Up the root.sh Script	5-2
5.2.2	Install Cluster Health Management	5-3
5.2.3	Tune Semaphore Parameters	5-4
5.2.4	Create a Fast Recovery Area Disk Group.....	5-4
5.3	Using Older Oracle Database Versions with Grid Infrastructure	5-5
5.3.1	General Restrictions for Using Older Oracle Database Versions	5-5
5.3.2	Pinning Cluster Nodes for Oracle Database Release 10.x or 11.x.....	5-6
5.3.3	Enabling The Global Services Daemon (GSD) for Oracle Database Release 9.2	5-7
5.3.4	Using the Correct LSNRCTL Commands	5-7
5.4	Modifying Oracle Clusterware Binaries After Installation.....	5-7

6 How to Modify or Deinstall Oracle Grid Infrastructure

6.1	Deciding When to Deinstall Oracle Clusterware	6-1
6.2	Migrating From Oracle Restart to Oracle Clusterware	6-1
6.3	Adding Standalone Grid Infrastructure Servers to a Cluster.....	6-2
6.4	Relinking Oracle Grid Infrastructure for a Cluster Binaries.....	6-3
6.5	Removing Oracle Clusterware and ASM.....	6-3

6.5.1	About the Deinstallation Tool.....	6-3
6.5.2	Example of Running the Deinstall Command for Oracle Clusterware and ASM	6-5
6.5.3	Example of a Deinstallation Parameter File for Grid Infrastructure for a Cluster.....	6-5

A Troubleshooting the Oracle Grid Infrastructure Installation Process

A.1	General Installation Issues.....	A-1
A.2	Interpreting CVU "Unknown" Output Messages Using Verbose Mode	A-4
A.3	Interpreting CVU Messages About Oracle Grid Infrastructure Setup.....	A-5
A.4	About the Oracle Clusterware Alert Log	A-7
A.5	Missing Operating System Packages On Linux	A-7
A.6	Performing Cluster Diagnostics During Oracle Grid Infrastructure Installations.....	A-8
A.7	Interconnect Configuration Issues.....	A-8

B Installing and Configuring Oracle Database Using Response Files

n	How Response Files Work.....	B-1
n	Reasons for Using Silent Mode or Response File Mode	B-2
n	General Procedure for Using Response Files	B-2
4.	Creating the oraInst.loc File.....	B-3
4.	Preparing a Response File.....	B-3
n	Editing a Response File Template	B-3
4.	Recording a Response File.....	B-5
7.	Running the Installer Using a Response File	B-6
5.	Running Net Configuration Assistant Using a Response File	B-7
n	Running Database Configuration Assistants Using Response Files	B-7
n	About the Database Configuration Assistant in Response File Mode.....	B-8
n	Running Database Configuration Assistant in Response File or Silent Mode	B-8
n	Postinstallation Configuration Using a Response File	B-9
n	About the Postinstallation Configuration File.....	B-9
n	Running Postinstallation Configuration Using a Response File.....	B-10

C Configuring Large Memory Optimization

C.1	Overview of HugePages	C-1
C.1.1	What HugePages Provides.....	C-1
C.2	Restrictions for HugePage Configurations	C-2

D Oracle Grid Infrastructure for a Cluster Installation Concepts

n	Understanding Preinstallation Configuration.....	D-1
n	Understanding Oracle Groups and Users.....	D-1
n	Understanding the Oracle Base Directory Path	D-3
n	Understanding Network Addresses	D-3
n	Understanding Network Time Requirements.....	D-5
n	Understanding Storage Configuration	D-6
n	Understanding Automatic Storage Management Cluster File System (ACFS).....	D-6
n	About Migrating Existing Oracle ASM Instances.....	D-6
n	About Converting Standalone Oracle ASM Installations to Clustered Installations	D-7
n	Understanding Out-of-Place Upgrade.....	D-7

E How to Complete Installation Prerequisite Tasks Manually

n	Configuring SSH Manually on All Cluster Nodes.....	E-1
n	Checking Existing SSH Configuration on the System.....	E-1
n	Configuring SSH on Cluster Nodes.....	E-2
3.	Enabling SSH User Equivalency on Cluster Nodes.....	E-4
3.	Configuring Kernel Parameters.....	E-5
n	Minimum Parameter Settings for Installation.....	E-5
n	Additional Parameter and Kernel Settings for SUSE Linux.....	E-6
4.	Checking OCFS2 Version Manually.....	E-7

F How to Upgrade to Oracle Grid Infrastructure 11g Release 2

F.1	Back Up the Oracle Software Before Upgrades.....	F-1
F.2	Restrictions for Clusterware and ASM Upgrades to Grid Infrastructure.....	F-1
F.3	Verify System Readiness for Upgrades.....	F-2
F.4	Upgrading an Existing Oracle Clusterware Installation.....	F-3
F.4.1	Preparing to Upgrade an Existing Oracle Clusterware Installation.....	F-3
F.5	Performing Rolling Upgrades From an Earlier Release.....	F-3
F.5.1	Verify System Readiness for Upgrades.....	F-4
F.5.2	Performing a Rolling Upgrade of Oracle Clusterware.....	F-4
F.5.3	Performing a Rolling Upgrade of Automatic Storage Management.....	F-5
F.6	Updating DB Control and Grid Control Target Parameters.....	F-6
F.7	Downgrading Oracle Clusterware After an Upgrade.....	F-7

Index

List of Tables

2-1	Swap Space Required as a Multiple of RAM.....	2-21
2-2	Grid Naming Service Example Network.....	2-27
2-3	Manual Network Configuration Example	2-28
2-4	Linux x86 (32-bit) Operating System Kernel Requirements	2-30
2-5	Linux x86 (32-bit) Grid Infrastructure and Oracle RAC Package Requirements	2-31
2-6	Linux x86 Oracle Database Features Package Requirements.....	2-32
2-7	Linux x86-64 Operating System Kernel Requirements	2-34
2-8	Linux x86-64 Oracle Grid Infrastructure and Oracle RAC Package Requirements	2-35
2-9	Linux x86-64 Oracle Database Features Package Requirements.....	2-37
3-1	Supported Storage Options for Oracle Clusterware and Oracle RAC.....	3-4
3-2	Oracle Clusterware Shared File System Volume Size Requirements.....	3-6
3-3	Oracle RAC Shared File System Volume Size Requirements.....	3-6
3-4	NFS Mount Options for Oracle RAC	3-11
3-5	Total Oracle Clusterware Storage Space Required by Redundancy Type	3-18
3-6	Total Oracle Database Storage Space Required by Redundancy Type.....	3-18
3-7	ORACLEASM Script Options	3-26
B-1	Response Files for Oracle Database.....	B-4
B-2	Response files for Oracle Grid Infrastructure.....	B-4

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 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 Automatic Storage Management in a grid infrastructure for a cluster installation.

For customers with specialized system roles who intend to install Oracle Real Application Clusters (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 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. For additional information, refer to books in the Related Documents list.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an

otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at <http://www.fcc.gov/cgb/consumerfacts/trs.html>, and a list of phone numbers is available at <http://www.fcc.gov/cgb/dro/trsphonebk.html>.

Related Documents

For more information, refer to the following Oracle resources:

Oracle Clusterware and Oracle Real Application Clusters Documentation

This installation guide reviews 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 Database or Oracle RAC, then complete preinstallation tasks as described in this installation guide, complete grid infrastructure installation, and review those installation guides for additional information. You can install either Oracle databases for a standalone server on a grid infrastructure installation, or install an Oracle RAC database. If you want to install an Oracle Restart deployment of grid infrastructure, then refer to *Oracle Database Installation Guide for Linux*

Most Oracle error message documentation is only available in HTML format. If you only have access to the Oracle Documentation media, then browse the error messages by range. When you find a range, use your browser's "find in page" 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. However, error messages for Oracle Clusterware and Oracle RAC tools are included in *Oracle Clusterware Administration and Deployment Guide*, or *Oracle Real Application Clusters Administration and Deployment Guide*.

Installation Guides

- *Oracle Diagnostics Pack Installation Guide*
- *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, 11g Release 2 (11.2) for UNIX Systems*
- *Oracle Database Platform Guide for Microsoft Windows*

Oracle Clusterware and Oracle Automatic Storage Management Administrative Guides

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

Oracle Real Application Clusters Administrative Guides

- *Oracle Real Application Clusters Administration and Deployment Guide*
- *Oracle Database 2 Day + Real Application Clusters Guide*
- *Oracle Database 2 Day DBA*
- *Getting Started with the Oracle Diagnostics Pack*

Generic Documentation

- *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 Web site:

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

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

<http://otn.oracle.com/membership/>

If you already have a username and password for OTN, then you can go directly to the documentation section of the OTN Web site at the following Web site:

<http://otn.oracle.com/documentation/>

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 "find in page" 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.

If you already have a username and password for OTN, then you can go directly to the documentation section of the OTN Web Site:

<http://otn.oracle.com/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.

Convention	Meaning
<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.

What's New in Oracle Grid Infrastructure Installation and Configuration?

This section describes new features as they pertain to the installation and configuration of Oracle grid infrastructure (Oracle Clusterware and Automatic Storage Management), and Oracle Real Application Clusters (Oracle RAC). The topics in this section are:

- [Desupported Options](#)
- [New Features for Release 2 \(11.2\)](#)
- [New Features for Release 1 \(11.1\)](#)

Desupported Options

The following is a list of options desupported with this release:

Block and Raw Devices Not Supported with OUI

With this release, OUI no longer supports installation of Oracle Clusterware files on block or raw devices. Install Oracle Clusterware files either on Automatic Storage Management diskgroups, or in a supported shared file system.

New Features for Release 2 (11.2)

The following is a list of new features for installation of Oracle Clusterware and Oracle ASM 11g release 2 (11.2):

Automatic Storage Management and Oracle Clusterware Installation

With Oracle grid infrastructure 11g release 2 (11.2), Automatic Storage Management (ASM) and Oracle Clusterware are installed into a single home directory, which is referred to as the Grid Infrastructure home. Configuration assistants start after the installer interview process that configure ASM and Oracle Clusterware.

The installation of the combined products is called Oracle grid infrastructure. However, Oracle Clusterware and Automatic Storage Manager remain separate products.

Automatic Storage Management and Oracle Clusterware Files

With this release, Oracle Cluster Registry (OCR) and voting disks can be placed on Automatic Storage Management (ASM).

This feature enables ASM to provide a unified storage solution, storing all the data for the clusterware and the database, without the need for third-party volume managers or cluster filesystems.

For new installations, OCR and voting disk files can be placed either on ASM, or on a cluster file system or NFS system. Installing Oracle Clusterware files on raw or block devices is no longer supported, unless an existing system is being upgraded.

Automatic Storage Management Cluster File System (ACFS)

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a new multi-platform, scalable file system and storage management design that extends Oracle Automatic Storage Management (ASM) technology to support all application data. Oracle ACFS provides dynamic file system resizing, and improved performance using the distribution, balancing and striping technology across all available disks, and provides storage reliability through ASM's mirroring and parity protection.

ASM Job Role Separation Option with SYSASM

The SYSASM privilege that was introduced in Oracle ASM 11g release 1 (11.1) is now fully separated from the SYSDBA privilege. If you choose to use this optional feature, and designate different operating system groups as the OSASM and the OSDBA groups, then the SYSASM administrative privilege is available only to members of the OSASM group. The SYSASM privilege also can be granted using password authentication on the Oracle ASM instance.

You can designate OPERATOR privileges (a subset of the SYSASM privileges, including starting and stopping ASM) to members of the OSOPER for ASM group.

Providing system privileges for the storage tier using the SYSASM privilege instead of the SYSDBA privilege provides a clearer division of responsibility between ASM administration and database administration, and helps to prevent different databases using the same storage from accidentally overwriting each other's files.

See Also: *Oracle Database Storage Administrator's Guide*

Cluster Time Synchronization Service

Cluster node times should be synchronized. With this release, Oracle Clusterware provides Cluster Time Synchronization Service (CTSS), which ensures that there is a synchronization service in the cluster. If Network Time Protocol (NTP) is not found during cluster configuration, then CTSS is configured to ensure time synchronization.

Enterprise Manager Database Control Provisioning

Enterprise Manager Database Control 11g provides the capability to automatically provision Oracle grid infrastructure and Oracle RAC installations on new nodes, and then extend the existing Oracle grid infrastructure and Oracle RAC database to these provisioned nodes. This provisioning procedure requires a successful Oracle RAC installation before you can use this feature.

See Also: *Oracle Real Application Clusters Administration and Deployment Guide* for information about this feature

Fixup Scripts and Grid Infrastructure Checks

With Oracle Clusterware 11g release 2 (11.2), the installer (OUI) 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 generating the fixup script by clicking the **Fix & Check Again** button.

The fixup script is generated during installation. You are prompted to run the script as `root` in a separate terminal session. When you run the script, it raises kernel values to required minimums, if necessary, and completes other operating system configuration tasks.

You also can have Cluster Verification Utility (CVU) generate fixup scripts before installation.

Grid Plug and Play

In the past, adding or removing servers in a cluster required extensive manual preparation. With this release, you can continue to configure server nodes manually, or use Grid Plug and Play to configure them dynamically as nodes are added or removed from the cluster.

Grid Plug and Play reduces the costs of installing, configuring, and managing server nodes by starting a grid naming service within the cluster to allow each node to perform the following tasks dynamically:

- Negotiating appropriate network identities for itself
- Acquiring additional information it needs to operate from a configuration profile
- Configuring or reconfiguring itself using profile data, making hostnames and addresses resolvable on the network

Because servers perform these tasks dynamically, the number of steps required to add or nodes is minimized.

Intelligent Platform Management Interface (IPMI) Integration

Intelligent Platform Management Interface (IPMI) is an industry standard management protocol that is included with many servers today. IPMI operates independently of the operating system, and can operate even if the system is not powered on. Servers with IPMI contain a baseboard management controller (BMC) which is used to communicate to the server.

If IPMI is configured, then Oracle Cluster uses IPMI when node fencing is required and the server is not responding.

Oracle Clusterware Out-of-place Upgrade

With this release, you can install a new version of Oracle Clusterware into a separate home from an existing Oracle Clusterware installation. This feature reduces the downtime required to upgrade a node in the cluster. When performing an out-of-place upgrade, the old and new version of the software are present on the nodes at the same time, each in a different home location, but only one version of the software is active.

Oracle Clusterware Administration with Oracle Enterprise Manager

With this release, you can use Enterprise Manager Cluster Home page to perform full administrative and monitoring support for both standalone database and Oracle RAC environments, using High Availability Application and Oracle Cluster Resource Management.

When Oracle Enterprise Manager is installed with Oracle Clusterware, it can provide a set of users that have the Oracle Clusterware Administrator role in Enterprise Manager, and provide full administrative and monitoring support for High Availability application and Oracle Clusterware resource management. After you have

completed installation and have Enterprise Manager deployed, you can provision additional nodes added to the cluster using Enterprise Manager.

SCAN for Simplified Client Access

With this release, the single client access name (SCAN) is the hostname to provide for all clients connecting to the cluster. The SCAN is a domain name registered to at least one and up to three IP addresses, either in the domain name service (DNS) or the Grid Naming Service (GNS). The SCAN eliminates the need to change clients when nodes are added to or removed from the cluster. Clients using the SCAN can also access the cluster using EZCONNECT.

SRVCTL Command Enhancements for Patching

With this release, you can use `srvctl` to shut down all Oracle software running within an Oracle home, in preparation for patching. Oracle grid infrastructure patching is automated across all nodes, and patches can be applied in a multi-node, multi-patch fashion.

Typical Installation Option

To streamline cluster installations, especially for those customers who are new to clustering, Oracle introduces the Typical Installation path. Typical installation defaults as many options as possible to those recommended as best practices.

Voting Disk Backup Procedure Change

In prior releases, backing up the voting disks using a `dd` command was a required postinstallation task. With Oracle Clusterware release 11.2 and later, backing up and restoring a voting disk using the `dd` command is not supported.

Backing up voting disks manually is no longer required, as voting disks are backed up automatically in the OCR as part of any configuration change and voting disk data is automatically restored to any added voting disks.

See Also: *Oracle Clusterware Administration and Deployment Guide*

New Features for Release 1 (11.1)

The following is a list of new features for release 1 (11.1)

Changes in Installation Documentation

With Oracle Database 11g release 1, Oracle Clusterware can be installed or configured as an independent product, and additional documentation is provided on storage administration. For installation planning, note the following documentation:

Oracle Database 2 Day + Real Application Clusters Guide

This book provides an overview and examples of the procedures to install and configure a two-node Oracle Clusterware and Oracle RAC environment.

Oracle Clusterware Installation Guide

This book (the guide that you are reading) provides procedures either to install Oracle Clusterware as a standalone product, or to install Oracle Clusterware with either Oracle Database, or Oracle RAC. It contains system configuration instructions that require system administrator privileges.

Oracle Real Application Clusters Installation Guide

This platform-specific book provides procedures to install Oracle RAC after you have completed successfully an Oracle Clusterware installation. It contains database configuration instructions for database administrators.

Oracle Database Storage Administrator's Guide

This book provides information for database and storage administrators who administer and manage storage, or who configure and administer Automatic Storage Management (ASM).

Oracle Clusterware Administration and Deployment Guide

This is the administrator's reference for Oracle Clusterware. It contains information about administrative tasks, including those that involve changes to operating system configurations and cloning Oracle Clusterware.

Oracle Real Application Clusters Administration and Deployment Guide

This is the administrator's reference for Oracle RAC. It contains information about administrative tasks. These tasks include database cloning, node addition and deletion, Oracle Cluster Registry (OCR) administration, use of SRVCTL and other database administration utilities, and tuning changes to operating system configurations.

Release 1 (11.1) Enhancements and New Features for Installation

The following is a list of enhancements and new features for Oracle Database 11g release 1 (11.1):

New SYSASM Privilege and OSASM operating system group for ASM Administration

This feature introduces a new `SYSASM` privilege that is specifically intended for performing ASM administration tasks. Using the `SYSASM` privilege instead of the `SYSDBA` privilege provides a clearer division of responsibility between ASM administration and database administration.

OSASM is a new operating system group that is used exclusively for ASM. Members of the OSASM group can connect as `SYSASM` using operating system authentication and have full access to ASM.

OPROCD and Hangcheck Timer Provide Fencing for Cluster Nodes

With Oracle Clusterware 11g, the Oracle Clusterware Process Monitor Daemon (`oprocd`) monitors the system state of cluster nodes. You should use both `OPROCD` and the hangcheck timer to provide I/O fencing for the cluster, and to ensure no stray writes occur from an evicted node in an Oracle RAC cluster.

Oracle Validated Configuration RPM

The Oracle Validated Configuration RPM can be installed on servers using Red Hat Enterprise Linux or Oracle Enterprise Linux operating systems. It is available from the Unbreakable Linux Network (ULN). This RPM completes most preinstallation configuration tasks required for installing Oracle Clusterware, including setting kernel parameters and adding required kernel RPMs. It also prepares a server for installation of Oracle RAC.

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

<https://linux.oracle.com>

Typical Installation for Oracle Grid Infrastructure for a Cluster

This chapter describes the difference between a Typical and Advanced installation for Oracle grid infrastructure for a cluster, and describes the steps required to complete a Typical installation.

This chapter contains the following sections:

- [Typical and Advanced Installation](#)
- [Preinstallation Steps Completed Using Typical Installation](#)
- [Preinstallation Steps Requiring Manual Tasks](#)

1.1 Typical and Advanced Installation

You are given two installation options for Oracle grid infrastructure installations:

- **Typical Installation:** The Typical installation option is a simplified installation with a minimal number of manual configuration choices. Oracle recommends that you select this installation type for most cluster implementations.
- **Advanced Installation:** The Advanced Installation option is an advanced procedure that requires a higher degree of system knowledge. It enables you to select particular configuration choices, including additional storage and network choices, use of operating system group authentication for role-based administrative privileges, integration with IPMI, or more granularity in specifying Automatic Storage Management roles.

1.2 Preinstallation Steps Completed Using Typical Installation

With Oracle Clusterware 11g release 2 (11.2), during installation Oracle Universal Installer (OUI) generates Fixup scripts (`runfixup.sh`) that you can run to complete required preinstallation steps.

The fixup script is generated during installation. You are prompted to run the script as `root` in a separate terminal session. When you run the script, it completes the following configuration tasks:

- If necessary sets kernel parameters required for installation and runtime to at least the minimum value.

1.3 Preinstallation Steps Requiring Manual Tasks

Complete the following manual configuration tasks

- [Verify System Requirements](#)
- [Check Network Requirements](#)
- [Check Operating System Packages](#)
- [Create Groups and Users](#)
- [Check Storage](#)
- [Prepare Storage for Automatic Storage Management](#)
- [Install Oracle Grid Infrastructure Software](#)

See Also: [Chapter 2, "Advanced Installation Oracle Grid Infrastructure for a Cluster Preinstallation Tasks"](#) and [Chapter 3, "Configuring Storage for Grid Infrastructure for a Cluster and Oracle Real Application Clusters \(Oracle RAC\)"](#) if you need any information about how to complete these tasks

1.3.1 Verify System Requirements

Enter the following commands to check available memory:

```
grep MemTotal /proc/meminfo
grep SwapTotal /proc/meminfo
```

The minimum required RAM is 1.5 GB for grid infrastructure for a cluster, or 2.5 GB for grid infrastructure for a cluster and Oracle RAC. The minimum required swap space is 1.5 GB. Oracle recommends that you set swap space to 1.5 times the amount of RAM for systems with 2 GB of RAM or less. For systems with 2 GB to 16 GB RAM, use swap space equal to RAM. For systems with more than 16 GB RAM, use 16 GB of RAM for swap space.

```
df -h
```

This command checks the available space on file systems. If you use standard redundancy for Oracle Clusterware files, which is three Oracle Cluster Registry (OCR) locations and three voting disk locations, then you should have at least 2 GB of file space available on shared storage volumes reserved for Oracle grid infrastructure files.

If you plan to install on Oracle ASM, then to ensure high availability of Oracle Clusterware files on Oracle ASM, you need to have at least 2 GB of disk space for Oracle Clusterware files in three separate failure groups, with at least three physical disks. Each disk must have at least 1 GB of capacity to ensure that there is sufficient space to create Oracle Clusterware files.

Ensure you have at least 4.5 GB of space for the grid infrastructure for a cluster home (Grid home) This includes Oracle Clusterware and Automatic Storage Management (Oracle ASM) files and log files.

```
df -h /tmp
```

Ensure that you have at least 1 GB of space in /tmp. If this space is not available, then increase the size, or delete unnecessary files in /tmp.

For more information, review the following section in [Chapter 2:](#)

["Checking the Hardware Requirements"](#)

1.3.2 Check Network Requirements

Ensure that you have the following available:

- [Single Client Access Name \(SCAN\) for the Cluster](#)
- [IP Address Requirements](#)
- [Intended Use of Network Interfaces](#)

1.3.2.1 Single Client Access Name (SCAN) for the Cluster

During Typical installation, you are prompted to confirm the default Single Client Access Name (SCAN), which is used to connect to databases within the cluster irrespective of which nodes they are running on. By default, the name used as the SCAN is also the name of the cluster. The default value for the SCAN is based on the local node name. If you change the SCAN from the default, then the name that you use must be globally unique throughout your enterprise.

In a Typical installation, the SCAN is also the name of the cluster. The SCAN and cluster name must be at least one character long and no more than 15 characters in length, must be alphanumeric, and may contain hyphens (-).

For example:

```
sales
```

If you require a SCAN that is longer than 15 characters, then select an Advanced installation.

Refer to the following section for the SCAN address requirements.

1.3.2.2 IP Address Requirements

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.

1.3.2.2.1 IP Address Requirements for Manual Configuration The public and virtual IP addresses must be static addresses, configured before installation, and the virtual IP addresses for each node must not currently be in use. Oracle Clusterware manages private IP addresses in the private subnet on interfaces you identify as private during the installation interview.

Configure the following addresses:

- A public IP address for each node
- A virtual IP address for each node
- A single client access name (SCAN) configured on the domain name server (DNS) for Round Robin resolution to three addresses (recommended) or at least one address.

The single client access name (SCAN) is a hostname 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: If you manually configure addresses, then Oracle strongly recommends that you use DNS resolution for SCAN VIPs. If you use the hosts file to resolve SCANS, then you must provide a hosts file entry for each SCAN address.

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

1.3.2.3 Intended Use of Network Interfaces

During installation, you are asked to identify the planned use for each network interface that OUI detects on your cluster node. You must identify each interface as a public or private interface, and you must use the same private interfaces for both Oracle Clusterware and Oracle RAC. For interfaces that you plan to have used for other purposes--for example, an interface dedicated to a network file system--you must identify those instances as "do not use" interfaces, so that Oracle Clusterware ignores them.

You can bond separate interfaces to a common interface to provide redundancy, in case of a NIC failure, but Oracle recommends that you do not create separate interfaces for Oracle Clusterware and Oracle RAC. If you use more than one NIC for the private interconnect, then Oracle recommends that you use NIC bonding. Note that multiple private interfaces provide load balancing but not failover, unless bonded.

1.3.3 Check Operating System Packages

Refer to the tables listed in [Section 2.8, "Identifying Software Requirements"](#) for the list of required packages for your operating system, or use a system configuration script such as the Oracle Validated RPM, as described in [Section 2.1, "Installing the Linux Operating System"](#).

1.3.4 Create Groups and Users

Enter the following commands to create default groups and users:

One system privileges group for all operating system-authenticated administration privileges, including Oracle RAC (if installed):

```
# groupadd -g 1000 oinstall
# groupadd -g 1200 dba
# useradd -u 1100 -g oinstall -G dba oracle
# mkdir -p /u01/app/11.2.0/grid
# mkdir -p /u01/app/oracle
# chown -R oracle:oinstall /u01
# chmod -R 775 /u01/
```

1.3.5 Check Storage

You must have space available on Automatic Storage Management for Oracle Clusterware files (voting disks and Oracle Cluster Registries), and for Oracle Database files, if you install standalone or Oracle Real Application Clusters Databases. Creating Oracle Clusterware files on block or raw devices is no longer supported for new installations.

Note: When using Oracle Automatic Storage Management (Oracle ASM) for either the Oracle Clusterware files or Oracle Database files, Oracle creates one Oracle ASM instance on each node in the cluster, regardless of the number of databases.

1.3.6 Prepare Storage for Automatic Storage Management

Install the Linux ASMLIB RPMs to simplify storage administration. ASMLIB provides persistent paths and permissions for storage devices used with ASM, eliminating the need for updating `udev` or `devlabel` files with storage device paths and permissions.

If you cannot install the ASMLIB RPMs, or choose to configure your storage devices manually, then review the relevant sections in Chapter 3.

See Also: [Chapter 3, "Configuring Storage for Grid Infrastructure for a Cluster and Oracle Real Application Clusters \(Oracle RAC\)"](#) if you require detailed storage configuration information, or require more information about ASMLIB

ASMLib 2.0 is delivered as a set of three Linux packages:

- `oracleasm-lib-2.0` - the ASM libraries
- `oracleasm-support-2.0` - utilities needed to administer ASMLib
- `oracleasm` - a kernel module for the ASM library

Each Linux distribution has its own set of ASMLib 2.0 packages, and within each distribution, each kernel version has a corresponding `oracleasm` package.

1.3.6.1 Installing ASMLIB Packages Automatically with ULN

If you are registered with the Oracle Unbreakable Linux Network (ULN), then you can download and install ASMLIB packages for your system automatically. To install ASMLIB from ULN:

1. Log in as `root`
2. Run the following command:

```
# up2date -i oracleasm-support oracleasm-lib oracleasm-'uname -r'
```

This command installs the support tools, the library, and the kernel driver for the Linux kernel version running on your system.

See Also: [Chapter 2.1, "Installing the Linux Operating System"](#) for information about how to register for ULN

1.3.6.2 Installing ASMLIB Packages Manually

If you are not a member of ULN, or are using a Red Hat or SUSE Linux kernel, then complete the following procedures on each node that you intend to make a member of the cluster:

- [Determine the Correct Oracleasm Package](#)
- [Download and Install the Oracleasm Package](#)
- [Configure ASMLib](#)

1.3.6.2.1 Determine the Correct Oracleasm Package Determine which kernel you are using by logging in as `root` and running the following command:

```
uname -rm
```

For example:

```
# uname -rm
2.6.9-5.ELsmp i686
```

The example shows that this is a 2.6.9-5 kernel for an SMP (multiprocessor) server using Intel i686 CPUs.

1.3.6.2.2 Download and Install the Oracleasm Package After you determine the kernel version for your system, complete the following task:

1. Open a Web browser using the following URL:

```
http://www.oracle.com/technology/tech/linux/asmlib/index.html
```

2. Select the link for your version of Linux.
3. Download the `oracleasm-lib` and `oracleasm-support` packages for your version of Linux.
4. Download the `oracleasm` package corresponding to your kernel version.
5. Log in as `root` and install the Oracle ASM packages.

1.3.6.2.3 Configure ASMLib Log in as `root`, and enter the following command:

```
# oracleasm configure -i
```

Provide information as prompted for your system. The `oracleasm` command by default is in the path `/usr/sbin`. If you enter the command `oracleasm configure` without the `-i` flag, then you are shown the current configuration.

See Also: [Section 3.3.1.4.1, "Installing and Configuring the ASM Library Driver Software"](#) for additional information about configuration

1.3.6.3 Mark ASM Candidate Disk Partitions

For OUI to recognize partitions as Oracle ASM disk candidates, you must log in as `root` and mark the disk partitions that Oracle ASM can use. To mark a disk for use by ASM, enter the following command syntax, where `ASM_DISK_NAME` is the name of the Oracle ASM disk group, and `candidate_disk` is the name of the disk device that you want to assign to that disk group:

```
oracleasm createdisk ASM_DISK_NAME candidate_disk
```

For example:

```
# oracleasm createdisk data1 /dev/sdf
```

1.3.7 Install Oracle Grid Infrastructure Software

1. Start OUI from the root level of the installation media. For example:

```
./runInstaller
```

2. Select **Install and Configure Grid Infrastructure for a Cluster**, then select **Typical Installation**. In the installation screens that follow, enter the configuration information as prompted.

If you receive an installation verification error that cannot be fixed using a fixup script, then review [Chapter 2, "Advanced Installation Oracle Grid Infrastructure for a Cluster Preinstallation Tasks"](#) to find the section for configuring cluster nodes. After completing the fix, continue with the installation until it is complete.

See Also: [Chapter 4 "Preparing to Install Oracle Grid Infrastructure with OUI"](#)

Advanced Installation Oracle Grid Infrastructure for a Cluster Preinstallation Tasks

This chapter describes the system configuration tasks that you must complete before you start Oracle Universal Installer (OUI) to install Oracle grid infrastructure.

This chapter contains the following topics:

- [Installing the Linux Operating System](#)
- [Reviewing Upgrade Best Practices](#)
- [Installation Fixup Scripts](#)
- [Logging In to a Remote System Using X Terminal](#)
- [Creating Groups, Users and Paths for Oracle Grid Infrastructure](#)
- [Checking the Hardware Requirements](#)
- [Checking the Network Requirements](#)
- [Identifying Software Requirements](#)
- [Checking the Software Requirements](#)
- [Network Time Protocol Setting](#)
- [Installing the cvuqdisk Package for Linux](#)
- [Enabling Intelligent Platform Management Interface \(IPMI\)](#)
- [Automatic SSH Configuration During Installation](#)
- [Configuring Grid Infrastructure Software Owner User Environments](#)
- [Requirements for Creating an Oracle Grid Infrastructure Home Directory](#)

2.1 Installing the Linux Operating System

Note: Ensure you have minimum hardware configuration completed before installing the operating system.

This section provides information about installing a supported Linux distribution. It contains the following topics:

- [Completing a Minimal Linux Installation](#)

- [Completing a Default Linux Installation](#)
- [About the Oracle Validated Configuration RPM](#)
- [Installing the Oracle Validated Configuration RPM](#)

2.1.1 Completing a Minimal Linux Installation

To complete a minimal Linux installation, select a minimal install option (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 installation. However, when you install the Oracle Validated RPM for your platform, the RPM downloads the minimum number of packages required to run Oracle Clusterware and Oracle Database.

Unbreakable Linux Network (ULN) customers can obtain the Oracle Validated RPM by using up2date. If you are not a ULN customer, and your operating system is Red Hat or Oracle Enterprise Linux, then you can obtain the Oracle Validated RPM at the following URLs:

Enterprise Linux 4:

<http://oss.oracle.com/el4/oracle-validated/>

Enterprise Linux 5:

<http://oss.oracle.com/el5/oracle-validated/>

Note: If you are not a member of ULN or RHN (Red Hat support network) and you are an Oracle support customer, then you can download instructions to configure a script that replicates Oracle Validated RPM package downloads at the following URL:

<https://metalink.oracle.com>

Search for "minimal Linux".

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.

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

2.1.2 Completing a Default Linux Installation

If you do not install the Oracle Validated 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 checks

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://metalink.oracle.com>

Search for "Default RPM."

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 for your system configuration.

2.1.3 About the Oracle Validated Configuration RPM

If your Linux distribution is Oracle Enterprise Linux, or Red Hat Enterprise Linux, and you are an Unbreakable Linux customer, then you can complete most preinstallation configuration tasks by using the Oracle Validated Configurations Setup RPM, available from the Unbreakable Linux Network (ULN).

When it is installed, the Oracle Validated Configuration RPM does the following:

- Automatically installs any additional packages needed for installing Oracle grid infrastructure and Oracle Database
- Creates an `oracle` user, and creates the `oraInventory` (`oinstall`) and `OSDBA` (`dba`) groups for that user
- Sets and verifies `sysctl.conf` settings, system startup parameters, user limits, and driver parameters to values based on recommendations from the Oracle Validated Configurations program

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

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

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

<https://linux.oracle.com>

If you are using Oracle Enterprise Linux 4.7 and higher, or Oracle Enterprise Linux 5.2 and higher, then the Oracle Validated RPM is included on the install media.

Note: The Oracle Validated RPM 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.

In addition, users and groups are created using the next available ID numbers. If server group and user IDs are not identical on the cluster nodes where you run the Oracle Validated RPM, then it is possible that these ids will be different on the nodes, and the different group and user IDs will cause installation to fail.

To avoid this issue, run the command `id user` on all nodes, where `user` is the software owner user account. If group or user IDs are not identical on all the nodes, then change them as necessary to ensure that they are identical.

2.1.4 Installing the Oracle Validated Configuration RPM

Use the following procedure to subscribe to Oracle Unbreakable Linux channels, and to add the Oracle Software for Enterprise Linux channel that distributes the Oracle Validated Configurations Setup RPM:

1. Complete a default Oracle Enterprise Linux workstation installation, or a default Red Hat Enterprise Linux installation.
2. Register your server with Unbreakable Linux Network (ULN). By default, you are registered for the Enterprise Linux Latest channel for your operating system and hardware.
3. Log in to ULN at the following URL:

<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. From the Available Channels list, select the Oracle Software for Enterprise Linux channel that is appropriate for your installation of Linux (for example: "Oracle Software for Enterprise Linux 4 (x86_64).")
6. Click **Subscribe**.
7. From a terminal session, as `root`, enter the following command:

```
# up2date --nox --show-channels
```

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

```
el4_i386_latest
el4_i386_oracle
```

8. Open a terminal session as `root`, and install the Oracle Validated Configurations Setup RPM with `up2date`, using the following command:
9. Repeat steps 1 through 8 on all other servers in your cluster.

Note: Check the Oracle Validated Configuration RPM log file to review system configuration changes:

```
/etc/sysconfig/oracle-validated/results/orakernel.log
```

2.2 Reviewing Upgrade Best Practices

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

If you have an existing Oracle installation, then document version numbers, patches, and other configuration information, and review upgrade procedures for your existing installation. Review Oracle upgrade documentation before proceeding with installation, to decide how you want to proceed.

You can upgrade Oracle ASM 11g release 1 (11.1) 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.

If you have an existing Automatic Storage Management (ASM) installation, then review Oracle upgrade documentation. The location of the Oracle ASM home changes in this release, and you may want to consider other configuration changes to simplify or customize storage administration.

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.

Note: Using mixed operating system versions is only supported for the duration of an upgrade, over the period of a few hours. Oracle does not support operating a cluster with mixed operating systems for an extended period. Oracle does not support running Oracle grid infrastructure and Oracle Real Application Clusters on heterogeneous platforms (servers with different chip architectures) in the same cluster.

To find the most recent software updates, and to find best practices recommendations about preupgrade, postupgrade, compatibility, and interoperability, refer to "Oracle Upgrade Companion." "Oracle Upgrade Companion" is available through Note 785351.1 on My Oracle Support:

<https://metalink.oracle.com>

2.3 Installation Fixup Scripts

With Oracle Clusterware 11g release 2, 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 fixup scripts (`runfixup.sh`). You can run the fixup script after you click the **Fix and Check Again** Button.

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

The Fixup script does the following:

- If necessary sets kernel parameters to values required for successful installation, including:
 - Shared memory parameters.
 - Open file descriptor and UDP send/receive parameters.
- Sets permissions on the Oracle Inventory (central inventory) directory.

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
```

2.4 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 a local terminal session, for example, an X terminal (`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 hostname. 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 hostname. 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, refer to your X server documentation for more information about completing this procedure. Depending on the X server software that you are using, you may need to complete the tasks in a different order.

1. Start the X server software.
2. Configure the security settings of the X server 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.

2.5 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](#)
- [Creating the Oracle Base Directory Path](#)
- [Creating Job Role Separation Operating System Privileges Groups and Users](#)
- [Example of Creating Standard Groups, Users, and Paths](#)
- [Example of Creating Role-allocated Groups, Users, and Paths](#)

Note: During a grid infrastructure installation, both Oracle Clusterware and Automatic Storage Management are installed. You no longer can have separate Oracle Clusterware installation owners and Automatic Storage Management installation owners.

2.5.1 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 group privilege).

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:1000:crs,oracle
```

2.5.2 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 1000 oinstall
```

The preceding command creates the group `oinstall`, with the group ID number 1000. Members of the `OINSTALL` group are granted privileges to write to the Oracle central inventory (`oraInventory`).

By default, if an `oraInventory` group does not exist, then the installer lists the primary group of the installation owner for the grid infrastructure for a cluster 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 grid infrastructure for a cluster installation owner has the same name and group ID.

2.5.3 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 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.

2.5.3.1 Understanding Restrictions for Oracle Software Installation Owners

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.

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 (ASM, in the grid infrastructure home, or the database in the Oracle home), using command syntax such as the following example, where `grid` is the Oracle grid infrastructure home:

```
ORACLE_HOME=/u01/app/11.2.0/grid; export ORACLE_HOME
```

If you try to administer an instance using `sqlplus`, `lsnrctl`, or `asmcmd` commands while `$ORACLE_HOME` is set to a different binary path, then you will encounter errors. When starting `srvctl` from a database home, `$ORACLE_HOME` should be set. or `srvctl` fails. But if you are using `srvctl` in the grid infrastructure home, then `$ORACLE_HOME` is ignored, and the oracle home path does not affect `srvctl` commands. You always have to change `$ORACLE_HOME` to the instance that you want to administer.

To create separate Oracle software owners to create separate users 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 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. 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 group. You cannot 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.

Caution: For 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.

2.5.3.2 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=501(oracle) gid=501(oinstall) groups=502(dba),503(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.

To use the existing user, ensure that the user's primary group is the Oracle Inventory group (`oinstall`). If this user account will be used for Oracle Database installations, and you plan to have a different user account as the owner of the Oracle Clusterware and Oracle ASM binaries, then ensure that the Oracle account is also a member of the group you plan to designate as the OSDBA for ASM group (the group whose members are permitted to write to Oracle ASM storage).

2.5.3.3 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 ID and group IDs are the same on each cluster member node. The following procedures uses `grid` as the name of the Oracle software owner, and `dba` as the OSASM group. To create separate system privilege groups to separate administration privileges, complete group creation before you create the user [Section 2.5.5, "Creating Job Role Separation Operating System Privileges Groups and Users,"](#) on page 2-12.

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 1100 -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, 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, and you must have 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 privileges, 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` privileges 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.

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

For example:

```
# id oracle
uid=501(oracle) gid=501(oracle) groups=501(oracle)
# /usr/sbin/usermod -u 1001 -g 1000 -G 1000,1001 oracle
# id oracle
uid=1001(oracle) gid=1000(oinstall) groups=1000(oinstall),1001(oracle)
```

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.

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.

2.5.4 Creating the Oracle Base Directory Path

The Oracle base directory for the grid installation owner is the location where diagnostic and administrative logs, and other logs associated with Oracle ASM and Oracle Clusterware are stored.

If you have created a path for the Oracle Clusterware home that is compliant with Oracle Optimal Flexible Architecture (OFA) guidelines for Oracle software paths then you do not need to create an Oracle base directory. When OUI finds an OFA-compliant path, it creates the Oracle base directory in that path.

For OUI to recognize the path as an Oracle software path, it must be in the form `u0[1-9]/app`, and it must be writable by any member of the `oraInventory` (`oinstall`) group. The Optimal Flexible Architecture path for the Oracle base is `/u01/app/user`, where `user` is the name of the Oracle software installation owner.

Oracle recommends that you create an Oracle base path manually, particularly if you have separate grid infrastructure for a cluster and Oracle Database software owners, so that you can separate log files.

For example:

```
# mkdir -p /u01/app/grid
# chown -R grid:oinstall /u01/app/grid
# chmod -R 775 /u01/app/grid
# mkdir -p /u01/app/oracle
# chown -R oracle:oinstall /u01/app/oracle
```

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 mmmaps.

2.5.5 Creating Job Role Separation Operating System Privileges Groups and Users

A Job Role Separation privileges configuration of Oracle ASM is a configuration with groups and users that divide administrative access privileges to the Oracle ASM installation from other administrative privileges users and groups associated with other Oracle installations. Administrative privileges access is granted by membership in separate operating system groups, and installation privileges are granted by using different installation owners for each Oracle installation.

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

If you prefer, you can allocate operating system user privileges so that you can use 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, Automatic Storage Management, and all Oracle Databases on the servers, and all privileges as installation owners. This group must also be the Oracle Inventory group.

Oracle recommends that you use at least two groups: A system privileges group whose members are granted administrative system privileges, and an installation owner group (the `oraInventory` group) to provide separate installation privileges the `OINSTALL` 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`.

- [Overview of Creating Job Role Separation Groups and Users](#)
- [Creating Database Groups and Users with Job Role Separation](#)

Note: To use a directory service, such as Network Information Services (NIS), refer to your operating system documentation for further information.

2.5.5.1 Overview of Creating Job Role Separation Groups and Users

This section provides an overview of how to create users and groups to use Job Role Separation. Log in as `root` to create these groups and users.

- [Users for Oracle Installations with Job Role Separation](#)
- [Database Groups for Job Role Separation Installations](#)
- [ASM Groups for Job Role Separation Installations](#)

2.5.5.1.1 Users for Oracle Installations with Job Role Separation Oracle recommends that you create the following operating system groups and users for all installations where you create separate software installation owners:

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 it) the OSOPER group 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 Clusterware Administration and Deployment Guide* and *Oracle Database Administrator's Guide* for more information about the OSDBA, OSASM and OSOPER groups and the SYSDBA, SYSASM and SYSOPER privileges

2.5.5.1.2 Database Groups for Job Role Separation Installations The following operating system groups and user are required if you are installing Oracle Database:

- The OSDBA group (typically, `dba`)

You must create this group the first time you install Oracle Database software on the system. This group identifies operating system user accounts that have database administrative privileges (the SYSDBA privilege). If you do not create separate OSDBA, OSOPER and OSASM groups for the Oracle ASM instance, then operating system user accounts that have the SYSOPER and SYSASM privileges must be members of this group. The name used for this group in Oracle code examples is `dba`. If you do not designate a separate group as the OSASM group, then the OSDBA group you define is also by default the OSASM group.

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

Members of the OSDBA group formerly were granted SYSASM privileges on Oracle ASM instances, including mounting and dismounting disk groups. This privileges grant is removed with 11g release 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 privilege is retained.

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

This is an optional group. Create this group if you want a separate group of operating system users to have a limited set of database administrative privileges

(the `SYSOPER` privilege). By default, members of the `OSDBA` group also have all privileges granted by the `SYSOPER` privilege.

To use the `OSOPER` group to create a database administrator group with fewer privileges than the default `dba` group, then you must choose the Advanced installation type to install the software or start OUI as a user that is not a member of the `dba` group. In this case, OUI prompts you to specify the name of this group. The usual name chosen for this group is `oper`.

2.5.5.1.3 ASM Groups for Job Role Separation Installations `SYSASM` is a new system privilege that enables the separation of the Oracle ASM storage administration privilege from `SYSDBA`. With Oracle Automatic Storage Management 11g release 2 (11.2), members of the database `OSDBA` group are not granted `SYSASM` privileges, unless the operating system group designated as the `OSASM` group is the same group designated as the `OSDBA` group.

Select separate operating system groups as the operating system authentication groups for privileges on Oracle ASM. Before you start OUI, create the following groups and users for Oracle ASM

- The Oracle Automatic Storage Management Group (typically `asmadmin`)

This is a required group. Create this group as a separate group if you want to have separate administration privilege groups for Oracle ASM and Oracle Database administrators. 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`.

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 separate `OSASM` group, and use a separate user from the database users to own the grid infrastructure installation (Oracle Clusterware and Oracle ASM). Oracle ASM can support multiple databases.

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.

- The ASM Database Administrator group (`OSDBA` 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 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.

- Members of the ASM Operator Group (`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 ASM Operator group to create an 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`.

If you want to have an OSOPER for ASM group, then the grid infrastructure for a cluster software owner must be a member of this group.

2.5.5.2 Creating Database Groups and Users with Job Role Separation

The following sections describe how to create the required operating system user and groups:

- [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](#)
- [When to Create 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](#)
- [Creating Identical Database Users and Groups on Other Cluster Nodes](#)

2.5.5.2.1 Creating the OSDBA Group to Prepare for Database Installations If you intend to install Oracle Database to use with the 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 as follows. Use the group name dba unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 1200 dba
```

2.5.5.2.2 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 you require a new OSOPER group, then create it as follows. Use the group name oper unless a group with that name already exists.

```
# /usr/sbin/groupadd -g 1201 oper
```

2.5.5.2.3 Creating the OSASM Group If the OSASM group does not exist or if you require a new OSASM group, then create it as follows. Use the group name `asmadmin` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 1000 asmadmin
```

2.5.5.2.4 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 you require a new OSOPER for ASM group, then create it as follows. In the following, use the group name `asmoper` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 1301 asmoper
```

2.5.5.2.5 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 as follows. Use the group name `asmdba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 1300 asmdba
```

2.5.5.2.6 When to Create 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`.

2.5.5.2.7 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=501(oracle) gid=501(oinstall) groups=502(dba),503(oper)
```

Determine whether you want to use the existing user, or create another user. To use the existing user, then 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. Refer to one of the following sections for more information:

- To modify an existing user, refer to the ["Modifying an Existing Oracle Software Owner User"](#) section on page 2-17.
- To create a user, refer to the following section.

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.

2.5.5.2.8 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 as follows. Use the user name `oracle` unless a user with that name already exists.

1. To create an `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -u 1101 -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
```

2.5.5.2.9 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 enter a command similar to the following to modify it. Specify the primary group using the `-g` option and any required secondary group using the `-G` option:

```
# /usr/sbin/usermod -g oinstall -G dba,asmdba oracle
```

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

2.5.5.2.10 Creating Identical Database Users and Groups on 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.

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.

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=502(oracle) gid=501(oinstall) groups=502(dba),503(oper),506(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 next cluster node as `root`.
2. Enter commands similar to the following to create the `oinstall`, `asmadmin`, and `asmdba` groups, and if required, the `asmoper`, `dba`, and `oper` groups. Use the `-g` option to specify the correct GID for each group.

```
# /usr/sbin/groupadd -g 1000 oinstall
# /usr/sbin/groupadd -g 1100 asmadmin
# /usr/sbin/groupadd -g 1200 dba
# /usr/sbin/groupadd -g 1201 oper
# /usr/sbin/groupadd -g 1300 asmdba
# /usr/sbin/groupadd -g 1301 asmoper
```

Note: If the 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 this 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 grid infrastructure (`grid`) user, enter a command similar to the following (in this example, to create the `oracle` user):

```
# /usr/sbin/useradd -u 1100 -g oinstall -G asmdba,dba oracle
```

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, which must be the Oracle Inventory group, for example `oinstall`
- The `-G` option specifies the secondary groups, which can include the OSASM, OSDBA, OSDBA for ASM, and OSOPER or OSOPER for ASM groups. For example:

- A grid installation owner: OSASM (`asmadmin`), whose members are granted the SYSASM privilege
- An Oracle Database installation owner without SYSASM privileges access: OSDBA (`dba`), OSDBA for ASM (`asmdba`), OSOPER for ASM (`asmoper`)

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 [Configuring Grid Infrastructure Software Owner User Environments](#) on page 2-47.

2.5.6 Example of Creating Standard Groups, Users, and Paths

The following is an example of how to create the Oracle Inventory group (`oinstall`), and a single group (`dba`) as the OSDBA, OSASM and OSDBA for ASM groups. In addition, it shows how to create the grid infrastructure software owner (`grid`), and one Oracle Database owner (`oracle`) with correct group memberships. This example also shows how to configure an Oracle base path compliant with OFA structure with correct permissions:

```
# groupadd -g 1000 oinstall
# groupadd -g 1200 dba
# useradd -u 1100 -g oinstall -G dba grid
# useradd -u 1101 -g oinstall -G dba oracle
# mkdir -p /u01/app/11.2.0/grid
# chown -R grid:oinstall /u01
# mkdir /u01/app/oracle
# 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.
- A single system privileges group that is used as the OSASM, OSDBA, OSDBA for ASM, and OSOPER for ASM group (`dba`), whose members are granted the SYSASM and SYSDBA privilege 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 installation for a cluster owner (`grid`), with the oraInventory group as its primary group, and with the OSASM group as the secondary group.
- An Oracle Database owner (`oracle`) with the oraInventory group as its primary group, and the OSDBA group as its secondary group.

- /u01/app owned by `grid:oinstall` with 775 permissions. This ownership and permissions enables OUI to create the Oracle Inventory directory, in the path `/u01/app/oraInventory`.
- /u01 owned by `root`.
- /u01/app/11.2.0/grid owned by `grid:oinstall` with 775 permissions. These permissions are required for installation, and are changed during the installation process.
- /u01/app/oracle owned by `oracle:oinstall` with 775 permissions.

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

The following is an example of how to create role-allocated groups and users that is compliant with an Optimal Flexible Architecture (OFA) deployment:

```
# groupadd -g 1000 oinstall
# groupadd -g 1100 asmadmin
# groupadd -g 1200 dba1
# groupadd -g 1250 dba2
# groupadd -g 1300 asmdba
# groupadd -g 1301 asmoper
# useradd -u 1100 -g oinstall -G asmadmin,asmdba,asmoper grid
# useradd -u 1101 -g oinstall -G dba1,asmdba oracle1
# useradd -u 1102 -g oinstall -G dba2,asmdba oracle2
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/11.2.0/grid
# chown -R grid:oinstall /u01
# mkdir -p /u01/app/oracle1
# chown oracle1:oinstall /u01/app/oracle1
# mkdir -p /u01/app/oracle2
# chown oracle2:oinstall /u01/app/oracle2
# chmod -R 775 /u01
```

After running these commands, you have the following groups and users:

- An Oracle central inventory group, or `oraInventory` group (`oinstall`), whose members that have the central inventory group as their primary group are granted permissions to write to the `oraInventory` directory.
- A separate OSASM group (`asmadmin`), whose members are granted the SYSASM privilege to administer Oracle Clusterware and Oracle ASM.
- A separate OSDBA for ASM group (`asmdba`), whose members include `grid`, `oracle1` and `oracle2`, and who are granted access to Oracle ASM.
- A separate OSOPER for ASM group (`asmoper`), whose members include `grid`, and who are granted limited Oracle ASM administrator privileges, including the permissions to start and stop the Oracle ASM instance.
- An Oracle grid installation for a cluster owner (`grid`), with the `oraInventory` group as its primary group, and with the OSASM (`asmadmin`), OSDBA for ASM (`asmdba`) and OSOPER for ASM groups as secondary groups.
- Two separate OSDBA groups for two different databases (`dba1` and `dba2`) to establish separate SYSDBA privileges for each database.
- Two Oracle Database software owners (`oracle1` and `oracle2`), to divide ownership of the Oracle database binaries, with the `OraInventory` group as their primary group, and the OSDBA group for their database (`dba1` or `dba2`) and the OSDBA for ASM group (`asmdba`) as their secondary groups.

- An OFA-compliant mount point `/u01` owned by `grid:oinstall` before installation.
- An Oracle base `/u01/app/oracle1` owned by `oracle1:oinstall` with 775 permissions.
- An Oracle base `/u01/app/oracle2` owned by `oracle2:oinstall` with 775 permissions.
- A Grid home `/u01/app/11.2.0/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`).
- 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. The grid installation owner Oracle base directory is the location where Oracle ASM diagnostic and administrative log files are placed. During installation, OUI creates the Oracle Inventory directory, in the path `/u01/app/oraInventory`. This path remains owned by `grid:oinstall`, to enable other Oracle software owners to write to the central inventory.

2.6 Checking the Hardware Requirements

- Select servers with the same chip 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 run level 3 or run level 5.
- 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.

Each system must meet the following minimum hardware requirements:

- At least 1.5 GB of physical RAM for grid infrastructure for a cluster installations without Oracle RAC; at least 2.5 GB of physical RAM if you plan to install Oracle RAC after installing grid infrastructure for a cluster.
- At least 1024 x 768 display resolution, so that Oracle Universal Installer (OUI) displays correctly
- Swap space equivalent to the multiple of the available RAM, as indicated in the following table:

Table 2–1 Swap Space Required as a Multiple of RAM

Available RAM	Swap Space Required
Between 1 GB and 2 GB	1.5 times the size of RAM
Between 2 GB and 16 GB	Equal to the size of RAM
More than 16 GB	16 GB

Note: On Linux, the Hugepages feature allocates non-swappable memory with 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.

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.

- 1 GB of space in the `/tmp` directory
- 4.5 GB of space for the grid infrastructure for a cluster home (Grid home) This includes Oracle Clusterware and Automatic Storage Management (ASM) files and log files.
- On Linux x86, if you intend to install Oracle Database, then allocate 4 GB of space for Oracle base.
- For Linux x86_64 platforms, if you intend to install Oracle Database, then allocate 4.6 GB of disk space for the Oracle base.

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

See Also: *Oracle Database Storage Administrator's Guide*

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

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, refer to 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 -k /tmp
```

This command displays disk space in 1 kilobyte blocks. On most systems, you can use the `df` command with the `-h` flag (`df -h`) to display output in "human-readable" format, such as "24G" and "10M." If there is less than 1 GB of disk space available in the `/tmp` directory (less than 1048576 1-k blocks), 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.

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

```
# free
```

Note that available RAM and swap space change, depending on user activity on the server.

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

```
# uname -m
```

Note: This command displays the processor type. If you intend to install on a 64-bit architecture, then the output should be "x86_64."

2.7 Checking the Network Requirements

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:

- [Network Hardware Requirements](#)
- [IP Address Requirements](#)
- [DNS Configuration for Domain Delegation to Grid Naming Service](#)
- [Manual IP Address Configuration Example](#)
- [Grid Naming Service Configuration Example](#)
- [Network Interface Configuration Options](#)
- [Enabling the Name Service Cache Daemon](#)

Note: For the most up-to-date information about supported network protocols and hardware for Oracle RAC installations, refer to the Certify pages on the My Oracle Support Web site at the following URL:

<https://metalink.oracle.com>

2.7.1 Network 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).

To use multiple NICs for the public network or for the private network, Oracle recommends that you use NIC bonding. Use separate bonding for the public and private networks, because during installation each interface is defined as a public or private interface.

- The public interface names associated with the network adapters for each network must be the same on all nodes, and the private interface names associated with the network adaptors should be the same on all nodes.

For example: With a two-node cluster, you cannot configure network adapters on `node1` with `eth0` as the public interface, but on `node2` have `eth1` as the public interface. Public interface names must be the same, so you must configure `eth0` as public on *both* nodes. You should configure the private interfaces on the same network adapters as well. If `eth1` is the private interface for `node1`, then `eth1` should be the private interface for `node2`.

- For the public network, each network adapter must support TCP/IP.
- For the private network, the interconnect 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 interconnect protocol for Oracle RAC, and TCP is the interconnect protocol for 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.

- 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`.
- During installation, you are asked to identify the planned use for each network interface that OUI detects on your cluster node. You must identify each interface as a public or private interface, and you must use the same private interfaces for both Oracle Clusterware and Oracle RAC.

You can bond separate interfaces to a common interface to provide redundancy, in case of a NIC failure, but Oracle recommends that you do not create separate interfaces for Oracle Clusterware and Oracle RAC. If you use more than one NIC for the private interconnect, then Oracle recommends that you use NIC bonding. Note that multiple private interfaces provide load balancing but not failover, unless bonded.

IP addresses on the subnet you identify as private are assigned as private IP addresses for cluster member nodes. You do not need to configure these addresses manually in a hosts directory.

2.7.2 IP Address Requirements

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:

- Oracle Grid Naming Service (GNS) using a static public node address and dynamically allocated IPs for the Oracle Clusterware provided VIP addresses, DHCP server assigned, and resolved using a multicast domain name server configured as part of Oracle Clusterware within the cluster. If you plan to use GNS, then you must have the following:
 - A DHCP service running on the public network for the cluster

- Enough addresses on the DHCP to provide 1 IP address for each node's virtual IP, and 3 IP addresses for the cluster used by the Single Client Access Name (SCAN) for the cluster
- Static addresses network administrators assign on a network domain name server (DNS) for each node. Selecting this option requires that you request network administration updates when you modify the cluster.

2.7.2.1 IP Address Requirements with Grid Naming Service

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. You define this address in the DNS domain before installation. The 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, before installation the DNS administrator must establish DNS Lookup to direct DNS resolution of a subdomain to the cluster. If you enable GNS, then you must have a DHCP service on the public network that allows the cluster to dynamically allocate the virtual IP addresses as required by the cluster.

2.7.2.2 IP Address Requirements for Manual Configuration

If you do not enable GNS, then the public and virtual IP addresses for each node must be static IP addresses, configured before installation for each node, but not currently in use. Public and virtual IP addresses must be on the same subnet.

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 addresses configured:

- A public IP address for each node
- A virtual IP address for each node
- A single client access name (SCAN) configured on the domain name server (DNS) for Round Robin resolution to three addresses (recommended) or at least one address.

The single client access name (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.

The SCAN addresses must be on the same subnet as virtual IP addresses and public IP addresses. For high availability and scalability, Oracle recommends that you configure the SCAN to use Round Robin resolution to three addresses. The name for the SCAN cannot begin with a numeral. For installation to succeed, the SCAN must resolve to at least one address.

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 you will only be able to resolve to one IP address and you will have only one SCAN address.

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

2.7.3 DNS Configuration for Domain Delegation to Grid Naming Service

If you plan to use GNS, then before grid infrastructure installation, you must configure your domain name server (DNS) to send to GNS name resolution requests for the subdomain GNS serves, which are the cluster member nodes.

You must configure the DNS to send GNS name resolution requests using delegation. Configure delegation using the following procedure:

1. In the DNS, create an entry for the GNS virtual IP address. For example:

```
gns-server.clustertype.com: 192.0.2.1
```

The address you provide must be routable.

2. In the DNS, create an entry similar to the following for the delegated domain, where `clusterdomain.example.com` is the subdomain you want to delegate:

```
clusterdomain.example.com: NS gns-server.clustertype.com
```

When using GNS, you must configure the `resolve.conf` on the nodes in the cluster to contain name server entries that are resolvable to corporate DNS servers. The total timeout period configured—a combination of options attempts (retries) and options timeout (exponential backoff)—should be less than 30 seconds. For example, where `xxx.xxx.xxx.42` and `xxx.xxx.xxx.15` are valid name server addresses in your network, provide an entry similar to the following in

`/etc/resolv.conf`:

```
options attempts: 2
options timeout: 1

search clusterdomain.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 Oracle 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
```

2.7.4 Grid Naming Service Configuration Example

If you use GNS, then you need to specify a static IP address for the GNS VIP address, and delegate a subdomain to be delegated to that 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 `example.com`, the subdomain is `grid.example.com`, 192.0.2 in the IP addresses represent the cluster public IP address network, and 192.168.0 represents the private IP address subnet:

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

¹ Node hostnames may resolve to multiple addresses, including any private IP addresses or VIP addresses currently running on that host.

2.7.5 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 2–3 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 hostnames 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.

2.7.6 Network Interface Configuration Options

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.

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.

2.7.7 Enabling the Name Service Cache Daemon

To prevent public network failures with Oracle RAC databases using 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`:

On Red Hat and Oracle Enterprise Linux 4 and 5:

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

On SUSE SLES10 and11:

```
# 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
```

2.8 Identifying Software Requirements

Depending on the products that you intend to install, verify that the following operating system software is installed on the system. To check these requirements refer to the section "[Checking the Software Requirements](#)" on page 2-39, following this section.

OUI performs checks 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.

Oracle recommends that you install your Linux operating system with the default software packages (RPMs), unless you specifically intend to perform a minimal installation, and follow the directions for performing such an installation to ensure that you have all required packages for Oracle software.

Oracle recommends that you do not customize RPMs during a default operating system installation. A default installation includes most required packages, and will help you to limit manual checks of package dependencies

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.

The following is the list of supported Linux versions and requirements at the time of release:

- [Software Requirements List for x86 \(32-bit\) Linux Platforms](#)
- [Software Requirements List for x86-64 Linux Platforms](#)

2.8.1 Software Requirements List for x86 (32-bit) Linux Platforms

For installations of Oracle grid infrastructure, ensure that you have the kernel versions and packages listed in [Table 2-4](#) and [Table 2-5](#).

If you intend to install Oracle Database or Oracle RAC in addition to Oracle grid infrastructure, then check [Table 2-6](#) to determine if you must install additional packages for the features you plan to use.

Note: For Asianux Server, Oracle Enterprise Linux, and Red Hat Enterprise Linux, system requirements are identical by kernel version. Specifically:

Asianux 2, Oracle Enterprise Linux 4, and Red Hat Enterprise Linux 4 requirements are the same.

Asianux Server 3, Oracle Enterprise Linux 5, and Red Hat Enterprise Linux 5 requirements are the same.

Table 2-4 Linux x86 (32-bit) Operating System Kernel Requirements

Linux Distribution	Requirements
Asianux Distributions	<ul style="list-style-type: none"> ■ Asianux 2, kernel 2.6.9 or later ■ Asianux Server 3, kernel 2.6.18 or later
Enterprise Linux Distributions	<ul style="list-style-type: none"> ■ Enterprise Linux 4 Update 7, kernel 2.6.9 or later ■ Enterprise Linux 5 Update 2, kernel 2.6.18 or later
Red Hat Enterprise Linux Distributions	<ul style="list-style-type: none"> ■ Red Hat Enterprise Linux 4 Update 7, kernel 2.6.9 or later ■ Red Hat Enterprise Linux 5 Update 2, kernel 2.6.18 or later
SUSE Enterprise Linux Distributions	<ul style="list-style-type: none"> ■ SUSE 10, kernel 2.6.16.21 or later ■ SUSE 11, kernel 2.6.27.19 or later

Table 2–5 Linux x86 (32-bit) Grid Infrastructure and Oracle RAC Package Requirements

Item	Requirements
Asianux 2, Enterprise Linux 4, and Red Hat Enterprise Linux 4	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.15.92.0.2 compat-libstdc++-33.2.3 elfutils-libelf-0.97 elfutils-libelf-devel-0.97 gcc-3.4.6 gcc-c++-3.4.6 glibc-2.3.4-2.41 glibc-common-2.3.4 glibc-devel-2.3.4 glibc-headers-2.3.4 libaio-devel-0.3.105 libaio-0.3.105 libgcc-3.4.6 libstdc++-3.4.6 libstdc++-devel-3.4.6 make-3.80 pdksh-5.2.14 sysstat-5.0.5 unixODBC-2.2.11 unixODBC-devel-2.2.11 </pre>
Asianux Server 3, Enterprise 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 elfutils-libelf-0.125 elfutils-libelf-devel-0.125 elfutils-libelf-devel-static-0.125 gcc-4.1.2 gcc-c++-4.1.2 glibc-2.5-24 glibc-common-2.5 glibc-devel-2.5 glibc-headers-2.5 kernel-headers-2.6.18 ksh-20060214 libaio-0.3.106 libaio-devel-0.3.106 libgcc-4.1.2 libgomp-4.1.2 libstdc++-4.1.2 libstdc++-devel-4.1.2 make-3.81 sysstat-7.0.2 unixODBC-2.2.11 unixODBC-devel-2.2.11 </pre>

Table 2–5 (Cont.) Linux x86 (32-bit) Grid Infrastructure and Oracle RAC Package

Item	Requirements
SUSE 10 Packages	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.16.91.0.5 compat-libstdc++-5.0.7 gcc-4.1.2 gcc-c++-4.1.2 glibc-2.5-24 glibc-devel-2.4 ksh-93r-12.9 libaio-0.3.104 libaio-devel-0.3.104 libelf-0.8.5 libgcc-4.1.2 libstdc++-4.1.2 libstdc++-devel-4.1.2 make-3.80 sysstat-8.0.4 </pre>
SUSE 11 packages	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.19 gcc-4.3 gcc-c++-4.3 glibc-2.9 glibc-devel-2.9 ksh-93t libstdc++33-3.3.3 libstdc++43-4.3.3_20081022 libstdc++43-devel-4.3.3_20081022 libaio-0.3.104 libaio-devel-0.3.104 libgcc43-4.3.3_20081022 libstdc++-devel-4.3 make-3.81 sysstat-8.1.5 </pre>

Table 2–6 Linux x86 Oracle Database Features Package Requirements

Item	Requirement
LDAP package	<p>If you did not perform a default Linux installation, you intend to use LDAP, and you want to use the scripts <code>odisrvreg</code>, <code>oidca</code>, or <code>schemasync</code>, then install the Korn shell RPM for your Linux distribution.</p>
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>Intel C/C++ Compiler 10.1 or later and the version of GNU C and C++ compilers listed previously for the distribution are supported for use with these products.</p> <p>Note: Intel C/C++ Compiler v10.1 can be used only with gcc 3.4.5, gcc 4.0 or gcc 4.1 standard template libraries to build OCCI applications.</p> <p>Oracle XML Developer's Kit is supported with the same compilers as OCCI.</p>

Table 2–6 (Cont.) Linux x86 Oracle Database Features Package Requirements

Item	Requirement
Oracle ODBC Drivers	<p>If you intend to use Open Database Connectivity (ODBC), then you should install the most recent ODBC Driver Manager for Linux.</p> <p>You should obtain the most current ODBC driver from your operating system vendor. You can read about ODBC at the following URL:</p> <p>http://www.unixodbc.org</p> <p>The ODBC RPMs are only needed if you plan on using ODBC. If you do not plan to use ODBC, then you do not need to install the ODBC RPMs for Oracle Clusterware, Oracle ASM, or Oracle RAC.</p>
Oracle JDBC/OCI Drivers	<p>You can use the following optional JDK version with the Oracle JDBC/OCI drivers; however, it is not required for the installation:</p> <ul style="list-style-type: none"> ■ Sun JDK 1.5.0-06 (JDK 5.0) with the JNDI extension ■ IBM Java 5.0 32-bit (SR1) or later <p>Note: By default, IBM Java 5.0 (32-bit) is installed with this release.</p>
Oracle Real Application Clusters	<p>For a cluster file system, use one of the following options:</p> <p>OCFS2</p> <ul style="list-style-type: none"> ■ Version 1.2.9 or later <p>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.</p> <p>To check your OCFS version, enter the following command:</p> <pre>modinfo ocfs2</pre> <p>To check the version of OCFS2 tools and OCFS2 console, enter the following command:</p> <pre>rpm -qa grep ocfs2</pre> <p>Ensure that <code>ocfs2console</code> and <code>ocfs2-tools</code> are at least version 1.2.9</p> <p>For information about OCFS2, refer to the following Web site:</p> <p>http://oss.oracle.com/projects/ocfs2/</p> <p>For OCFS2 certification status, refer to the Certify page on My Oracle Support.</p>

2.8.2 Software Requirements List for x86-64 Linux Platforms

For installations only of Oracle grid infrastructure, ensure that you have the kernel versions and packages listed in [Table 2–7](#) and [Table 2–8](#).

If you intend to install Oracle Database or Oracle RAC in addition to Oracle grid infrastructure, then check [Table 2–9](#) to determine if you must install additional packages for the features you plan to use.

Note: For Asianux Server, Oracle Enterprise Linux, and Red Hat Enterprise Linux, system requirements are identical by kernel version. Specifically:

Asianux 2, Oracle Enterprise Linux 4, and Red Hat Enterprise Linux 4 requirements are the same.

Asianux Server 3, Oracle Enterprise Linux 5, and Red Hat Enterprise Linux 5 requirements are the same.

Table 2–7 Linux x86-64 Operating System Kernel Requirements

Linux Distribution	Requirements
Asianux Distributions	<ul style="list-style-type: none">■ Asianux 2, kernel 2.6.9 or later■ Asianux 3, kernel 2.6.18 or later
Enterprise Linux Distributions	<ul style="list-style-type: none">■ Enterprise Linux 4 Update 7, kernel 2.6.9 or later■ Enterprise Linux 5 Update 2, kernel 2.6.18 or later
Red Hat Enterprise Linux Distributions	<ul style="list-style-type: none">■ Red Hat Enterprise Linux 4 Update 7, kernel 2.6.9 or later■ Red Hat Enterprise Linux 5 Update 2, kernel 2.6.18 or later
SUSE Enterprise Linux Distributions	<ul style="list-style-type: none">■ SUSE 10, kernel 2.6.16.21 or later■ SUSE 11, kernel 2.6.27.19 or later

Table 2–8 Linux x86-64 Oracle Grid Infrastructure and Oracle RAC Package Requirements

Item	Requirements
Asianux 2, Enterprise Linux 4, and Red Hat Enterprise Linux 4	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.15.92.0.2 compat-libstdc++-33-3.2.3 compat-libstdc++-33-3.2.3 (32 bit) elfutils-libelf-0.97 elfutils-libelf-devel-0.97 expat-1.95.7 gcc-3.4.6 gcc-c++-3.4.6 glibc-2.3.4-2.41 glibc-2.3.4-2.41 (32 bit) glibc-common-2.3.4 glibc-devel-2.3.4 glibc-headers-2.3.4 libaio-0.3.105 libaio-0.3.105 (32 bit) libaio-devel-0.3.105 libaio-devel-0.3.105 (32 bit) libgcc-3.4.6 libgcc-3.4.6 (32-bit) libstdc++-3.4.6 libstdc++-3.4.6 (32 bit) libstdc++-devel 3.4.6 make-3.80 pdksh-5.2.14 sysstat-5.0.5 unixODBC-2.2.11 unixODBC-2.2.11 (32 bit) unixODBC-devel-2.2.11 unixODBC-devel-2.2.11 (32 bit) </pre>

Table 2–8 (Cont.) Linux x86-64 Oracle Grid Infrastructure and Oracle RAC Package Requirements

Item	Requirements
Asianux Server 3, Enterprise 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) elfutils-libelf-0.125 elfutils-libelf-devel-0.125 gcc-4.1.2 gcc-c++-4.1.2 glibc-2.5-24 glibc-2.5-24 (32 bit) glibc-common-2.5 glibc-devel-2.5 glibc-devel-2.5 (32 bit) glibc-headers-2.5 ksh-20060214 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 make-3.81 sysstat-7.0.2 unixODBC-2.2.11 unixODBC-2.2.11 (32 bit) unixODBC-devel-2.2.11 unixODBC-devel-2.2.11 (32 bit) </pre>
SUSE 10	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.16.91.0.5 compat-libstdc++-5.0.7 gcc-4.1.0 gcc-c++-4.1.2 glibc-2.5-24 glibc-devel-2.4 glibc-devel-32bit-2.4 ksh-93r-12.9 libaio-0.3.104 libaio-32bit-0.3.104 libaio-devel-0.3.104 libaio-devel-32bit-0.3.104 libelf-0.8.5 libgcc-4.1.2 libstdc++-4.1.2 make-3.80 sysstat-8.0.4 </pre>

Table 2–8 (Cont.) Linux x86-64 Oracle Grid Infrastructure and Oracle RAC Package Requirements

Item	Requirements
SUSE 11	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.19 gcc-4.3 gcc-32bit-4.3 gcc-c++-4.3 glibc-2.9 glibc-32bit-2.9 glibc-devel-2.9 glibc-devel-32bit-2.9 ksh-93t libaio-0.3.104 libaio-32bit-0.3.104 libaio-devel-0.3.104 libaio-devel-32bit-0.3.104 libstdc++33-3.3.3 libstdc++33-32bit-3.3.3 libstdc++43-4.3.3_20081022 libstdc++43-32bit-4.3.3_20081022 libstdc++43-devel-4.3.3_20081022 libstdc++43-devel-32bit-4.3.3_20081022 libgcc43-4.3.3_20081022 libstdc++-devel-4.3 make-3.81 sysstat-8.1.5 </pre>

Table 2–9 Linux x86-64 Oracle Database Features Package Requirements

Item	Requirement
LDAP package	If you did not perform a default Linux installation, you intend to use LDAP, and you want to use the scripts <code>odisrvreg</code> , <code>oidca</code> , or <code>schemasync</code> , then install the Korn shell RPM for your Linux distribution.
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>Intel C/C++ Compiler 10.1 or later and the version of GNU C and C++ compilers listed previously for the distribution are supported for use with these products.</p> <p>Note: Intel C/C++ Compiler v10.1 can be used only with gcc 3.4.5, gcc 4.0 or gcc 4.1 standard template libraries to build OCCI applications.</p> <p>Oracle XML Developer's Kit is supported with the same compilers as OCCI.</p>

Table 2–9 (Cont.) Linux x86-64 Oracle Database Features Package Requirements

Item	Requirement
Oracle ODBC Drivers	<p>If you intend to use Open Database Connectivity (ODBC), then you should install the most recent ODBC Driver Manager for Linux.</p> <p>You should obtain the most current ODBC driver from your operating system vendor. You can read about ODBC at the following URL:</p> <p>http://www.unixodbc.org</p> <p>The ODBC RPMs are only needed if you plan on using ODBC. If you do not plan to use ODBC, then you do not need to install the ODBC RPMs for Oracle Clusterware, Oracle ASM, or Oracle RAC.</p> <p>To use ODBC, you must also install the following additional 32-bit ODBC RPMs, depending on your operating system:</p> <p>Asianux 2, Enterprise Linux 4, and Red Hat Enterprise Linux 4:</p> <p>unixODBC-2.2.11 (32 bit) or later unixODBC-devel-2.2.11 (32 bit) or later</p> <p>Asianux Server 3, Enterprise Linux 5, and Red Hat Enterprise Linux 5:</p> <p>unixODBC-2.2.11 (32 bit) or later unixODBC-devel-2.2.11 (32 bit) or later</p>
Oracle JDBC/OCI Drivers	<p>You can use the following optional JDK version with the Oracle JDBC/OCI drivers; however, it is not required for the installation:</p> <ul style="list-style-type: none"> ■ Sun JDK 1.5.0-06 (JDK 5.0) with the JNDI extension ■ IBM Java 5.0 32-bit (SR1) or later <p>Note: By default, IBM Java 5.0 (32-bit) is installed with this release.</p>
Oracle Real Application Clusters	<p>For a cluster file system, use the following option:</p> <p>OCFS2</p> <ul style="list-style-type: none"> ■ Version 1.2.9 or later <p>To check your OCFS2 version, enter the following command:</p> <pre>modinfo ocfs2</pre> <p>To check the version of OCFS2 tools and OCFS2 console, enter the following command:</p> <pre>rpm -qa grep ocfs2</pre> <p>Ensure that <code>ocfs2console</code> and <code>ocfs2-tools</code> are at least version 1.2.7</p> <p>For information about OCFS2, refer to the following Web site:</p> <p>http://oss.oracle.com/projects/ocfs2/</p> <p>For OCFS2 certification status, refer to the Certify page on My Oracle Support.</p> <p>http://oss.oracle.com/projects/ocfs2/</p> <p>For OCFS2 certification status, refer to the Certify page on My Oracle Support.</p>

2.9 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 following command:

```
# cat /proc/version
```

Note: Only the distributions and versions listed in the previous table are supported. Do not install the software on other versions of Linux.

2. 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 a Red Hat Enterprise Linux 4.0 system:

```
2.6.9-55.0.0.0.2.ELsmp
```

In this example, the output shows the kernel version (2.6.9) and errata level (55.0.0.0.2.ELsmp) on the system.

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

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

```
# rpm -q package_name
```

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

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

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

```
# rpm -q binutils compat-libstdc++ 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 Web site.

4. To determine if OCFS2 is installed, enter the following command:

```
# /sbin/modinfo ocfs2
```

To install the Oracle Database files on an OCFS2 file system and the packages are not installed, then download them from the following Web site. Follow the instructions listed with the kit to install the packages and configure the file system:

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

2.10 Network Time Protocol Setting

Oracle Clusterware 11g release 2 (11.2) and later requires time synchronization across all nodes within a cluster when Oracle RAC is deployed. Oracle provide two options for time synchronization: an operating system configured network time protocol (NTP), or Oracle Cluster Time Synchronization Service (ctssd). Oracle Cluster Time Synchronization Service is designed for organizations whose Oracle RAC databases are unable to access NTP services.

If you want to use Cluster Time Synchronization Service to provide synchronization service in the cluster, then deconfigure and deinstall the Network Time Protocol (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 Enterprise Linux, and Asianux systems, run the following commands as the root user

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

or, `mv /etc/ntp.conf to /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 initialization file 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 Enterprise 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 ntp 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 ntp restart
```

2.11 Installing the cvuqdisk Package for Linux

Install the operating system package `cvuqdisk`. 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`, or `i386`).

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

Note: If you prefer, you can choose to disable Cluster Verification Utility shared disk checks by adding the following line to the file `grid_home/cv/admin/cvuconfig`:

```
CV_RAW_CHECK_ENABLED=FALSE
```

1. Locate the `cvuqdisk` RPM package, which is in the directory `rpm` on the installation media. If you have already installed Oracle grid infrastructure, then it is located in the directory `grid_home/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-1.0.7-1.rpm
```

2.12 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. With Oracle 11g release 2, 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 a node-termination protocol, such as IPMI. 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

2.12.1 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 will not function properly.

2.12.2 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) you still must configure separate addresses for the IPMI interfaces. As the IPMI adapter is not seen directly by the host, the IPMI adapter is not visible to GNS as an address on the host.

2.12.3 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 Enterprise Linux, Red Hat Enterprise Linux, and SUSE Enterprise Linux distributions supported with this release.

2.12.3.1 Configuring the Open IPMI Driver

On Linux systems, the OpenIPMI driver is the supported driver for Oracle Clusterware deployments with IPMI.

You can install the driver dynamically by manually loading the required modules. To install the driver manually:

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
```

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.

2.12.3.2 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.

2.12.3.2.1 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 2-43 and ensure that the IPMI driver is running.

3. Enable IPMI over LAN using the following procedure

- a. Determine the channel number for the channel used for IPMI over LAN. Beginning with channel 1, run the following command until you find the channel that displays LAN attributes (for example, the IP address):

```
# ipmitool lan print 1

. . .
IP Address Source      : 0x01
IP Address             : 140.87.155.89
. . .
```

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

```
# ipmitool 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) will be associated only with the BMC, and will not respond to normal pings.

5. Establish an administration account with a username and password, using the following procedure (assuming the channel is 1):

- a. Set BMC to require password authentication for ADMIN access over LAN. For example:

```
# ipmitool lan set 1 auth ADMIN MD5,PASSWORD
```

- b. List the account slots on the BMC, and identify an unused slot (a User ID with an empty user name field). For example:

```
# ipmitool channel getaccess 1

. . .
User ID           : 4
User Name         :
```

```

Fixed Name          : No
Access Available    : call-in / callback
Link Authentication : disabled
IPMI Messaging      : disabled
Privilege Level     : NO ACCESS
. . .

```

- c. Assign the desired administrator user name and password and enable messaging for the identified slot. (Note that for IPMI v1.5 the user name and password can be at most 16 characters). Also, set the privilege level for that slot when accessed over LAN (channel 1) to ADMIN (level 4). For example, where *username* is the administrative user name, and *password* is the password:

```

# ipmitool user set name 4 username
# ipmitool user set password 4 password
# ipmitool user enable 4
# ipmitool channel setaccess 1 4 privilege=4
# ipmitool channel setaccess 1 4 link=on
# ipmitool channel setaccess 1 4 ipmi=on

```

- d. Verify the setup using the command `lan print 1`. The output should appear similar to the following. Note that the items in **bold** text are the settings made in the preceding configuration steps, and comments or alternative options are indicated within brackets []:

```

# ipmitool lan print 1
Set in Progress          : Set Complete
Auth Type Support        : NONE MD2 MD5 PASSWORD
Auth Type Enable         : Callback : MD2 MD5
                        : User      : MD2 MD5
                        : Operator  : MD2 MD5
                        : Admin   : MD5 PASSWORD
                        : OEM      : MD2 MD5
IP Address Source      : DHCP Address [or Static Address]
IP Address             : 192.168.0.55
Subnet Mask           : 255.255.255.0
MAC Address              : 00:14:22:23:fa:f9
SNMP Community String    : public
IP Header                 : TTL=0x40 Flags=0x40 Precedence=...
Default Gateway IP     : 192.168.0.1
Default Gateway MAC      : 00:00:00:00:00:00
.
.
.
# ipmitool channel getaccess 1 4
Maximum User IDs         : 10
Enabled User IDs         : 2

User ID                  : 4
User Name              : username [This is the administration user]
Fixed Name               : No
Access Available         : call-in / callback
Link Authentication    : enabled
IPMI Messaging        : enabled
Privilege Level       : ADMINISTRATOR

```

- 6. Verify that the BMC is accessible and controllable from a remote node in your cluster using the `bmc info` command. For example, if `node2-ipmi` is the network hostname 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.

2.13 Automatic SSH Configuration During Installation

To install Oracle software, Secure Shell (SSH) connectivity must 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: SSH is used by Oracle configuration assistants for configuration operations from local to remote nodes. It is also used by Enterprise Manager.

You can configure SSH from the Oracle Universal Installer (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 Oracle software installation owners, 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. If they are not disabled, then SSH must be configured manually before an installation can be run.

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

2.14 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. Also, create other required Oracle software owners, if needed.

This section contains the following topics:

- [Environment Requirements for Oracle Grid Infrastructure Software Owner](#)
- [Procedure for Configuring Oracle Software Owner Environments](#)
- [Setting Resource Limits for the Oracle Software Installation Users](#)
- [Setting Display and X11 Forwarding Configuration](#)
- [Preventing Installation Errors Caused by stty Commands](#)

2.14.1 Environment Requirements for Oracle Grid Infrastructure Software Owner

You must make the following changes to configure the Oracle grid infrastructure software owner environment:

- 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 the Oracle grid infrastructure installation

2.14.2 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 a new terminal session; for example, start an X terminal (`xterm`).
2. Enter the following command to ensure that X Window applications can display on this system:

```
$ xhost + hostname
```

The *hostname* is the name of the local host.
3. If you are not already logged in to the system where you want to install the software, then log in to that system as the software owner user.
4. If you are not logged in as the user, then switch to the software owner user you are configuring. For example, with the `grid` user:

```
$ su - grid
```
5. To determine the default shell for the user, enter the following command:

```
$ echo $SHELL
```
6. 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
```
7. Enter or edit the following line, specifying a value of `022` for the default file mode creation mask:

```
umask 022
```
8. If the `ORACLE_SID`, `ORACLE_HOME`, or `ORACLE_BASE` environment variables are set in the file, then remove these lines from the file.
9. Save the file, and exit from the text editor.

10. 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
```

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:

```
$ DISPLAY=local_host:0.0 ; export DISPLAY
```

- 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 you determined that 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 `TEMP` 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:

```
$ su - root
# mkdir /mount_point/tmp
# chmod 775 /mount_point/tmp
# exit
```

- c. Enter commands similar to the following to set the `TEMP` and `TMPDIR` environment variables:

- * Bourne, Bash, or Korn shell:

```
$ TEMP=/mount_point/tmp
$ TMPDIR=/mount_point/tmp
$ export TEMP TMPDIR
```

- * C shell:

```
% setenv TEMP /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.

2.14.3 Setting Resource Limits for the Oracle Software Installation Users

To improve the performance of the software on Linux systems, you must increase the following resource limits for the Oracle software owner users (`grid`, `oracle`):

Process Resource Shell Limit	Item in <code>limits.conf</code>	Hard Limit
Maximum number of open file descriptors	<code>nofile</code>	65536
Maximum number of processes available to a single user	<code>nproc</code>	16384
Maximum size of the stack segment of the process	<code>stack</code>	10240

In addition, to increase the shell limits:

1. On each node, add the following lines to the `/etc/security/limits.conf` file (the following example shows the software account owners `oracle` and `grid`):

```
grid          soft  nproc  2047
grid          hard  nproc  16384
grid          soft  nofile 1024
grid          hard  nofile 65536
oracle        soft  nproc  2047
oracle        hard  nproc  16384
oracle        soft  nofile 1024
oracle        hard  nofile 65536
```

2. On each node, add or edit the following line in the `/etc/pam.d/login` file, if it does not already exist:

```
session    required    pam_limits.so
```

3. Depending on your shell environment, make the following changes to the default shell startup file, to change `ulimit` setting for all Oracle installation owners (note that these examples show the users `oracle` and `grid`):

- For the Bourne, Bash, or Korn shell, add lines similar to the following to the `/etc/profile` file (or the file on SUSE systems) `/etc/profile.local`

```
if [ $USER = "oracle" ] || [ $USER = "grid" ]; then
    if [ $SHELL = "/bin/ksh" ]; then
        ulimit -p 16384
        ulimit -n 65536
    else
        ulimit -u 16384 -n 65536
    fi
fi
umask 022
fi
```

- For the C shell (`cs`h or `tc`sh), on Red Hat, OEL, or Asianux, add the following lines to the `/etc/csh.login` file. On SUSE systems add the lines to `/etc/csh.login.local`:

```
if ( $USER == "oracle" || $USER == "grid" ) then
    limit maxproc 16384
    limit descriptors 65536
endif
```

4. Repeat this procedure on all other nodes in the cluster, and for all Oracle software owners that you intend to use to install Oracle software.

2.14.4 Setting 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 the `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 hostname is `node1`, then enter the following command:

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

To ensure that X11 forwarding will 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. Make sure that the `ForwardX11` attribute is set to `no`. For example:

```
Host *
    ForwardX11 no
```

2.14.5 Preventing Installation Errors Caused by `stty` 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 `stty` commands.

To avoid this problem, you must modify these files in each Oracle installation owner user home directory to suppress all output on `STDERR`, 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.

2.15 Requirements for Creating an Oracle Grid Infrastructure Home Directory

During installation, you are prompted to provide a path to a home directory to store Oracle grid infrastructure software. Ensure that the directory path you provide meets the following requirements:

- 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 should be created either as a subdirectory in a path where all files can be owned by `root`, or in a unique path.
- 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.

For installations with Oracle grid infrastructure only, Oracle recommends that you create a 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 `u0[1-9]/app`.

When OUI finds an OFA-compliant path, it creates the Oracle grid infrastructure and Oracle Inventory (`oraInventory`) directories for you.

To create an Oracle grid infrastructure path manually, ensure that it is in a separate path, not under an existing Oracle base path. For example:

```
# mkdir -p /u01/app/11.2.0/grid
# chown grid:oinstall /u01/app/11.2.0/grid
# chmod -R 775 /u01/app/11.2.0/grid
```

With this path, if the installation owner is named `grid`, then by default OUI creates the following path for the `grid` home:

```
/u01/app/11.2.0/grid
```

Create an Oracle base path for database installations, owned by the Oracle Database installation owner account. The OFA path for an Oracle base is `/u01/app/user`, where `user` is the name of the Oracle software installation owner account. For example, use the following commands to create an Oracle base for the database installation owner account `oracle`:

```
# mkdir -p /u01/app/oracle
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

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 grid installation owner Oracle base or the Oracle Database installation owner 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.

Homes for Oracle grid infrastructure for a standalone server (Oracle Restart) can be under Oracle base. Refer to *Oracle Database Installation Guide* for your platform for more information about Oracle Restart.

Configuring Storage for Grid Infrastructure for a Cluster and Oracle Real Application Clusters (Oracle RAC)

This chapter describes the storage configuration tasks that you must complete before you start the installer to install Oracle Clusterware and Automatic Storage Management (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](#)
- [Shared File System Storage Configuration](#)
- [Automatic Storage Management Storage Configuration](#)
- [Desupport of Block and Raw Devices](#)

3.1 Reviewing Oracle Grid Infrastructure Storage Options

This section describes supported options for storing Oracle grid infrastructure for a cluster storage options. It contains the following sections:

- [Overview of Oracle Clusterware and Oracle RAC Storage Options](#)
- [General Storage Considerations for Oracle Grid Infrastructure and Oracle RAC](#)
- [Supported Storage Options](#)
- [After You Have Selected Disk Storage Options](#)

See Also: The Oracle Certify site for a list of supported vendors for Network Attached Storage options:

<http://www.oracle.com/technology/support/metalink/>

Refer also to the Certify site on My Oracle Support for the most current information about certified storage options:

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

3.1.1 Overview of Oracle Clusterware and Oracle RAC Storage Options

There are two ways of storing Oracle Clusterware files:

- **Automatic Storage Management (Oracle ASM):** You can install Oracle Clusterware files (OCR and voting disks) in Oracle ASM diskgroups.

Oracle ASM is the required database storage option for Typical installations, and for Standard Edition Oracle RAC installations. It is an integrated, high-performance database file system and disk manager for Oracle Clusterware and Oracle Database files. It performs striping and mirroring of database files automatically.

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

Only one Oracle ASM instance is permitted for each node regardless of the number of database instances on the node.

Note: For Oracle Automatic Storage Management (Oracle ASM) 11g release 2 (11.2) for Linux, Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) are only supported in the following environments:

- Red Hat and Oracle Enterprise Linux 5, 32-bit
- Red Hat and Oracle Enterprise Linux 5, 64-bit

For OVM environments, Red Hat and Oracle Enterprise Linux 5 Update 4 or later is required.

You cannot put Oracle Clusterware binaries and files on Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

You cannot put Oracle Database files on Oracle ACFS.

You can put Oracle Database binaries on Oracle ACFS.

If you plan to install an Oracle RAC home on a shared OCFS2 location, then you must upgrade OCFS2 to at least version 1.4.1, which supports shared writable mmmaps.

ACFS provides a general purpose file system for other files.

- **A supported shared file system:** Supported file systems include the following:
 - **Network File System (NFS):** Note that if you intend to use NFS for your data files, then you should create partitions large enough for the database files when you create partitions for Oracle grid infrastructure. NFS mounts differ for software binaries, Oracle Clusterware files, and database files.

Note: Placing Oracle grid infrastructure for a cluster binaries on a cluster file system is not supported.

You can no longer use OUI to install Oracle Clusterware or Oracle Database files on block or raw devices.

See Also: My Oracle Support for supported file systems and NFS or NAS filers

3.1.2 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 databases (Oracle RAC). To enable automated backups during the installation, you must also choose the storage option to use for recovery files (the Fast Recovery Area). You do not have to use the same storage option for each file type.

3.1.2.1 General Storage Considerations for Oracle Clusterware

Oracle Clusterware voting disks are used to monitor cluster node status, and Oracle Cluster Registry (OCR) files contain configuration information about the cluster. You can place voting disks and OCR files either in an ASM diskgroup, or on a cluster file system or shared network file system. Storage must be shared; any node that does not have access to an absolute majority of voting disks (more than half) will be restarted.

3.1.2.2 General Storage Considerations for Oracle RAC

Use the following guidelines when choosing the storage options to use for each file type:

- 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.
- Oracle recommends that you choose Oracle ASM as the storage option for database and recovery files.
- For Standard Edition Oracle RAC installations, Oracle ASM is the only supported storage option for database or recovery files.
- 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 11g release 2 (11.2) installed as part of an Oracle grid infrastructure for a cluster installation.
 - Any existing Oracle ASM instance on any node in the cluster is shut down.
- Raw or block devices are supported only when upgrading an existing installation using the partitions already configured. On new installations, using raw or block device partitions is not supported by Automatic Storage Management Configuration Assistant (ASMCA) or Oracle Universal Installer (OUI), but is supported by the software if you perform manual configuration.

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 disk areas to provide voting disk redundancy.

3.1.3 Supported Storage Options

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

Note: For information about OCFS2, refer to the following Web site:

<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, refer to the Certify page on My Oracle Support.

Table 3–1 Supported Storage Options for Oracle Clusterware and Oracle RAC

Storage Option	OCR and Voting Disks	Oracle Clusterware binaries	Oracle RAC binaries	Oracle Database Files	Oracle Recovery Files
Automatic Storage Management	Yes	No	No	Yes	Yes
Automatic Storage Management Cluster File System (ACFS)	No	No	Yes	No	No
NFS file system on a certified NAS filer	Yes	Yes	Yes	Yes	Yes
Note: Direct NFS does not support Oracle Clusterware files.					
Shared disk partitions (block devices or raw devices)	Not supported by OUI or ASMCA, but supported by the software. They can be added or removed after installation.	No	No	Not supported by OUI or ASMCA, but supported by the software. They can be added or removed after installation.	No

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 11g release 2 (11.2) to store Oracle Clusterware files. You cannot use prior Oracle ASM releases to do this.
- If you do not have a storage option that provides external file redundancy, then you must configure at least three voting disk locations and at least three Oracle Cluster Registry locations to provide redundancy.

3.1.4 After You Have Selected Disk Storage Options

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

- **To use a file system**, refer to [Shared File System Storage Configuration](#) on page 3-4.
- **To use Automatic Storage Management**, refer to ["Using Diskgroups with Oracle Database Files on ASM"](#) on page 3-28

3.2 Shared File System Storage Configuration

The installer does not suggest a default location for the Oracle Cluster Registry (OCR) or the Oracle Clusterware voting disk. If you choose to create these files on a file

system, then review the following sections to complete storage requirements for Oracle Clusterware files:

- [Requirements for Using a Shared File System](#)
- [Deciding to Use a Cluster File System for Oracle Clusterware Files](#)
- [Deciding to Use Direct NFS for Data Files](#)
- [Deciding to Use NFS for Data Files](#)
- [Configuring Storage NFS Mount and Buffer Size Parameters](#)
- [Checking NFS Mount and Buffer Size Parameters for Oracle Clusterware](#)
- [Checking NFS Mount and Buffer Size Parameters for Oracle RAC](#)
- [Enabling Direct NFS Client Oracle Disk Manager Control of NFS](#)
- [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](#)

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.

3.2.1 Requirements for Using a Shared File System

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 certified NAS device. Log in to My Oracle Support at the following URL, and click the Certify tab to find a list of certified NAS devices.
<https://metalink.oracle.com/>
- If you choose to place your Oracle Cluster Registry (OCR) files on a shared file system, then Oracle recommends that one of the following is true:
 - 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 11g release 2 (11.2) 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 disk partitions, then you can continue to use those partition sizes.

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

Use [Table 3–2](#) and [Table 3–3](#) to determine the minimum size for shared file systems:

Table 3–2 Oracle Clusterware Shared File System Volume Size Requirements

File Types Stored	Number of Volumes	Volume Size
Voting disks with external redundancy	3	At least 280 MB for each voting disk volume.
Oracle Cluster Registry (OCR) with external redundancy	1	At least 280 MB for each OCR volume
Oracle Clusterware files (OCR and voting disks) with redundancy provided by Oracle software.	1	At least 280 MB for each OCR volume At least 280 MB for each voting disk volume

Table 3–3 Oracle RAC Shared File System Volume Size 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 3–2](#) and [Table 3–3](#), 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 disk on each volume). You should have a minimum of three physical disks, each at least 500 MB, to ensure that voting disks 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 5.5 GB available total for all volumes.

Note: If you create partitions on shared partitions with `fdisk` by specifying a device size, such as `+300M`, 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.

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

For new installations, Oracle recommends that you use Automatic Storage Management (Oracle ASM) to store voting disk and OCR files. For Linux x86 (32-bit) and x86-64 (64-bit) platforms, Oracle provides a cluster file system, OCFS2. However, Oracle does not recommend using OCFS2 for Oracle Clusterware files.

3.2.3 Deciding to Use Direct NFS for Data Files

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

- [About Direct NFS Storage](#)
- [Using the Oranfstab File with Direct NFS](#)
- [Mounting NFS Storage Devices with Direct NFS](#)

3.2.3.1 About Direct NFS Storage

With Oracle Database 11g release 2 (11.2), instead of using the operating system kernel NFS client, you can configure Oracle Database to access NFS V3 servers directly using an Oracle internal Direct NFS client.

To enable Oracle Database to use Direct NFS, the NFS file systems must be mounted and available over regular NFS mounts before you start installation. Direct NFS manages settings after installation. You should still set the kernel mount options as a backup, but for normal operation, Direct NFS will manage NFS mounts.

Refer to your vendor documentation to complete NFS configuration and mounting.

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 it for Direct NFS to operate. To disable reserved port checking, consult your NFS file server documentation.

Note: Use NFS servers certified for Oracle RAC. Refer to the following URL for certification information:

<https://metalink.oracle.com>

3.2.3.2 Using the Oranfstab File with Direct NFS

If you use Direct NFS, then you can choose to use a new file specific for Oracle data file management, `oranfstab`, to specify additional options specific for Oracle Database to Direct NFS. 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 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 3.2.5, "Configuring Storage 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.

Caution: Direct NFS will not serve an NFS server with write size values (`wtmp`) less than 32768.

3.2.3.3 Mounting NFS Storage Devices with Direct NFS

Direct NFS 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 searches for mount entries in the following order:

1. `$_ORACLE_HOME/dbs/oranfstab`
2. `/etc/oranfstab`
3. `/etc/mtab`

Direct NFS uses the first matching entry found.

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

If Oracle Database uses Direct NFS 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 logs an informational message, and does not operate.

If Oracle Database cannot open an NFS server using Direct NFS, 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 3-10. Additionally, an informational message is logged into the Oracle alert and trace files indicating that Direct NFS could not be established.

The Oracle files resident on the NFS server that are served by the Direct NFS Client are also accessible through the operating system kernel NFS client. The usual considerations for maintaining integrity of the Oracle files apply in this situation.

3.2.3.4 Specifying Network Paths with the Oranfstab File

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

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

- `gv$dnfs_servers`: Shows a table of servers accessed using Direct NFS.
- `gv$dnfs_files`: Shows a table of files currently open using Direct NFS.

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

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

3.2.4 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.

3.2.5 Configuring Storage 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 so that they allow `root` on the clients mounting to the storage to be considered `root` instead of being mapped to an anonymous user, and allow `root` on the client server to create files on the NFS filesystem that are owned by `root`.

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`, `node2`, 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, as using this syntax allows 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 refer to 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
```

3.2.6 Checking NFS Mount and Buffer Size Parameters for Oracle Clusterware

On the 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 are:

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

If you have Oracle grid infrastructure binaries on an NFS mount, then you must include the `suid` option.

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,rsize=32768,wsize=32768,tcp,vers=3,timeo=600,actimeo=0
```

Note that mount point options are different for Oracle software binaries, Oracle Clusterware files (OCR and voting disks), 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 -yes \  
rw,bg,hard,nointr,rsize=32768,wsize=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://metalink.oracle.com>

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

3.2.7 Checking NFS Mount and Buffer Size Parameters for Oracle RAC

If you use kernel-managed NFS mounts, 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.

In general, most vendors recommend that you use the NFS mount options listed in [Table 3–4](#).

Table 3–4 NFS Mount Options for Oracle RAC

Option	Requirement	Description
hard	Mandatory	Generate a hard mount of the NFS file system. If the connection to the server fails or is temporarily lost, then connection attempts are made until the NAS device responds
bg	Optional	Try to connect in the background if connection fails.
rw	Mandatory	Read and write access
tcp	Optional	Use the TCP protocol rather than UDP. TCP is more reliable than UDP.
vers=3	Optional	Use NFS version 3. Oracle recommends that you use NFS version 3 where available, unless the performance of version 2 is higher.
suid	Optional	Allow clients to run software binaries with SUID enabled. SUID is required for all NFS mounts that contain Oracle software.
rsize	Mandatory	The number of bytes used when reading from the NAS device. This value should be set to the maximum database block size supported by this platform. A value of 8192 is often recommended for NFS version 2 and 32768 is often recommended for NFS version 3.
wsize	Mandatory	The number of bytes used when writing to the NAS device. This value should be set to the maximum database block size supported by this platform. A value of 8192 is often recommended for NFS version 2 and 32768 is often recommended for NFS version 3.
nointr (or intr)	Optional	Do not allow (or allow) keyboard interrupts to stop a process that is hung while waiting for a response on a hard-mounted file system. Note: Different vendors have different recommendations about this option. Contact your vendor for advice.
actime=0	Optional	Disable attribute caching. Note: You must specify this option for NFS file systems where you want to install the software binaries. If you do not use this option, then the installer will not install the software in the directory that you specify.
actimeo	Optional	Using actimeo sets all of acregmin, acregmax, acdirmin, and acdirmax to the same value. There is no default value.
timeo	Optional	Timeout setting. Better overall performance may be achieved by increasing the timeout when mounting on a busy network, to a slow server, or through several routers or gateways. Oracle recommends that you set the timeout value to the maximum timeout for TCP, which is 600 seconds.

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,nfsvers=3,timeo=600,rsize=32768,wsize=32768,actimeo=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. Refer to 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://metalink.oracle.com>

3.2.8 Enabling Direct NFS Client Oracle Disk Manager Control of NFS

Complete the following procedure to enable Direct NFS:

1. Create an `oranfstab` file with the following attributes for each NFS server to be accessed using Direct NFS:
 - **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).
 - **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.

The examples that follow show three possible NFS server entries in `oranfstab`. A single `oranfstab` can have multiple NFS server entries.

Example 3-1 Using Local and Path NFS Server Entries

The following example uses both local and path. Since they are in different subnets, we do not have to specify `dontroute`.

```
server: MyDataServer1
local: 192.0.2.0
path: 192.0.2.1
local: 192.0.100.0
path: 192.0.100.1
export: /vol/oradata1 mount: /mnt/oradata1
```

Example 3-2 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

```

Example 3–3 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

```

2. Oracle Database uses an ODM library, `libnfsodm11.so`, to enable Direct NFS. To replace the standard ODM library, `$ORACLE_HOME/lib/libodm11.so`, with the ODM NFS library, `libnfsodm11.so`, complete the following steps on all nodes unless the Oracle home directory is shared:
 - a. Change directory to `$ORACLE_HOME/lib`.
 - b. Enter the following commands:

```

cp libodm11.so libodm11.so_stub
ln -s libnfsodm11.so libodm11.so

```

3.2.9 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 disk, 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 Oracle Cluster Registry (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.

3.2.10 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/fast_recovery_area
# chown oracle:oinstall /mount_point/fast_recovery_area
# chmod 775 /mount_point/fast_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.

3.2.11 Disabling Direct NFS Client Oracle Disk Management Control of NFS

Use one of the following methods to disable the Direct NFS client:

- Remove the `oranfstab` file.
- Restore the stub `libodm11.so` file by reversing the process you completed in step 2b, "[Enabling Direct NFS Client Oracle Disk Manager Control of NFS](#)"
- On all nodes, remove the specific NFS server or export paths in the `oranfstab` file.

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

3.3 Automatic Storage Management Storage Configuration

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

- [Configuring Storage for Automatic Storage Management](#)
- [Using Diskgroups with Oracle Database Files on ASM](#)
- [Configuring Oracle Automatic Storage Management Cluster File System \(ACFS\)](#)
- [Migrating Existing Oracle ASM Instances](#)
- [Converting Standalone Oracle ASM Installations to Clustered Installations](#)

3.3.1 Configuring Storage for Automatic Storage Management

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

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

3.3.1.1 Identifying Storage Requirements for Automatic Storage Management

To identify the storage requirements for using Automatic Storage Management, 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 Automatic Storage Management for Oracle Clusterware files (OCR and voting disks), 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 Automatic Storage Management for the other.

If you choose to enable automated backups and you do not have a shared file system available, then you must choose Automatic Storage Management for recovery file storage.

If you enable automated backups during the installation, then you can select Automatic Storage Management as the storage mechanism for recovery files by specifying an Automatic Storage Management disk group for the Fast Recovery Area. Depending on how you choose to create a database during the installation, you have the following options:

- If you select an installation method that runs ASMCA in interactive mode (for example, by choosing the Advanced database configuration option) then you can decide whether you want to use the same Automatic Storage Management

disk group for database files and recovery files, or use different failure groups for each file type.

- If you select an installation method that runs DBCA in noninteractive mode, then you must use the same Automatic Storage Management disk group for database files and recovery files.
2. Choose the Automatic Storage Management redundancy level to use for the Automatic Storage Management disk group.

The redundancy level that you choose for the Automatic Storage Management disk group determines how Automatic Storage Management mirrors files in the disk group and determines the number of disks and amount of free disk space that you require, 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 Automatic Storage Management 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, Automatic Storage Management 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, Normal redundancy disk groups provide 3 voting disk files, 1 OCR and 2 copies (one primary and one secondary mirror). 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, Automatic Storage Management 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, High redundancy disk groups provide 5 voting disk files, 1 OCR and 3 copies (one primary and two secondary mirrors). 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.

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 3–5](#) and [Table 3–6](#) 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 disks in a separate disk group:

Table 3–5 Total Oracle Clusterware Storage Space Required by Redundancy Type

Redundancy Level	Minimum Number of Disks	Oracle Cluster Registry (OCR) Files	Voting Disk Files	Both File Types
External	1	280 MB	280 MB	560 MB
Normal	3	560 MB	840 MB	1.4 GB ¹
High	5	840 MB	1.4 GB	2.3 GB

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

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

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

Table 3–6 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
High	3	4.5 GB	9 GB	13.5 GB

4. For Oracle Clusterware installations, you must also add additional disk space for the Automatic Storage Management metadata. You can use the following formula to calculate the additional disk space requirements (in MB):

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

Where:

- redundancy = Number of mirrors: external = 1, normal = 2, high = 3.
- ausize = Metadata AU size in megabytes.
- 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 X 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, you need to have at least 2 GB of disk space for Oracle Clusterware files in three separate failure groups, with at least three physical disks. Each disk must have at least 1 GB of capacity to ensure that there is sufficient space to create Oracle Clusterware files.

5. For Oracle RAC installations, you must also add additional disk space for the Automatic Storage Management metadata. You can use the following formula to calculate the additional disk space requirements (in MB):

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

Where:

- *ausize* = Metadata AU size in megabytes.
- *clients* = Number of database instances for each node.
- *disks* = Number of disks in disk group.
- *nodes* = Number of nodes in cluster.
- *redundancy* = Number of mirrors: external = 1, normal = 2, high = 3.

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 disk space:

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

If an Automatic Storage Management instance is already running on the system, then you can use an existing disk group to meet these storage requirements. If necessary, you can add disks to an existing disk group during the installation.

6. Optionally, identify failure groups for the Automatic Storage Management disk group devices.

Note: Complete this step only if you intend to use an installation method that runs Database Configuration Assistant in interactive mode; for example, if you intend to choose the Custom installation type or the Advanced database configuration option, then complete this step. Other installation types do not enable you to specify failure groups.

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 `asmctl`, 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 disks, 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 Automatic Storage Management 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. Automatic Storage Management expects each disk group device to be on a separate physical disk.
 - Although you can specify a logical volume as a device in an Automatic Storage Management disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing Automatic Storage Management from optimizing I/O across the physical devices. They are not supported with Oracle RAC.

3.3.1.2 Creating Files on a NAS Device for Use with Automatic Storage Management

If you have a certified NAS storage device, then you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an Automatic Storage Management 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://metalink.oracle.com>

For more information about editing the mount file for the operating system, refer to the man pages. For more information about recommended mount options, refer to the section "[Checking NFS Mount and Buffer Size Parameters for Oracle RAC](#)" on page 3-10.

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
```

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 Automatic Storage Management disk discovery string to specify a regular expression that matches the file names you created. For example:

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

Note: During installation, disk paths mounted on Oracle ASM and registered on ASMLIB with the string `ORCL:*` are listed as default database storage candidate disks.

3.3.1.3 Using an Existing Automatic Storage Management Disk Group

To store either database or recovery files in an existing Automatic Storage Management disk group, then you have the following choices, depending on the installation method that you select:

- If you select an installation method that runs Database Configuration Assistant in interactive mode (for example, by choosing the Advanced database configuration option), 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 Automatic Storage Management instance that manages the existing disk group can be running in a different Oracle home directory.

To determine if an existing Automatic Storage Management disk group exists, or to determine if there is sufficient disk space in a disk group, you can use the ASM command line tool (`asmcmd`), Oracle Enterprise Manager Grid Control or Database Control. Alternatively, you can use the following procedure:

1. View the contents of the `oratab` file to determine if an Automatic Storage Management instance is configured on the system:

```
$ more /etc/oratab
```

If an Automatic Storage Management instance is configured on the system, then the `oratab` file should contain a line similar to the following:

```
+ASM2:oracle_home_path
```

In this example, `+ASM2` is the system identifier (SID) of the Automatic Storage Management instance, with the node number appended, and `oracle_home_path` is the Oracle home directory where it is installed. By convention, the SID for an Automatic Storage Management instance begins with a plus sign.

2. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the Automatic Storage Management instance.
3. Connect to the Automatic Storage Management instance and start the instance if necessary:

```
$ $ORACLE_HOME/bin/asmcmd  
ASMCMD> startup
```

4. 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> lsdb
```

or:

```
$ORACLE_HOME/bin/asmcmd -p lsdg
```

5. From the output, identify a disk group with the appropriate redundancy level and note the free space that it contains.
6. 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.

3.3.1.4 Configuring Disks for Automatic Storage Management with ASMLIB

The Automatic Storage Management library driver (ASMLIB) simplifies the configuration and management of the disk devices by eliminating the need to rebind disk devices used with Oracle ASM each time the system is restarted.

Without ASMLIB Linux 2.6 kernel and later, block device paths do not maintain permissions and path persistence unless you create a `permissions` or `rules` file on each cluster member node; block device paths that were `/dev/sda` can appear as `/dev/sdb` after a system restart. Adding new disks requires you to modify the `udev` file to provide permissions and path persistence for the new disk.

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.

If you intend to use Automatic Storage Management on block devices for database storage for Linux, then Oracle recommends that you install the ASMLIB driver and associated utilities, and use them to configure the disks for ASM.

To use the Automatic Storage Management library driver (ASMLIB) to configure Automatic Storage Management devices, complete the following tasks.

- [Installing and Configuring the ASM Library Driver Software](#)
- [Configuring Disk Devices to Use ASM Library Driver on x86 Systems](#)
- [Administering the ASM Library Driver and Disks](#)

Note: To create a database during the installation using the 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:*`.

3.3.1.4.1 Installing and Configuring the ASM Library Driver Software If you are a member of the Unbreakable Linux Network, then you can install the ASMLIB rpms by subscribing to the Oracle Software for Enterprise Linux channel, and using `up2date` to retrieve the most current package for your system and kernel. For additional information, refer to the following URL:

<http://www.oracle.com/technology/tech/linux/asmlib/uln.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 OTN Web site:

<http://www.oracle.com/technology/tech/linux/asmlib/index.html>

Note: You must install `oracleasm-support` package version 2.0.1 or later to use ASMLIB on Red Hat Enterprise Linux Advanced Server, or SUSE Linux Enterprise Server.

You must install the following packages, 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
oracleasm-lib-version.arch.rpm
```

3. Switch user to the `root` user:

```
$ su -
```

4. Enter a command similar to the following to install the packages:

```
# rpm -Uvh oracleasm-support-version.arch.rpm \
    oracleasm-kernel-version.arch.rpm \
    oracleasm-lib-version.arch.rpm
```

For example, if you are using the Red Hat Enterprise Linux AS 4 enterprise kernel on an AMD64 system, then enter a command similar to the following:

```
# rpm -Uvh oracleasm-support-2.0.1.x86_64.rpm \
    oracleasm-lib-2.0.1.x86_64.rpm \
    oracleasm-2.6.9-11.EL-2.0.1.x86_64.rpm
```

5. Enter the following command to run the `oracleasm` initialization script with the `configure` option:

```
# /usr/sbin/oracleasm configure -i
```

Note: The `oracleasm` command in `/usr/sbin` is the command you should use. The `/etc/init.d` path is not deprecated, but the `oracleasm` binary in that path is now used typically for internal commands.

6. Enter the following information in response to the prompts that the script displays:

Prompt	Suggested Response
Default user to own the driver interface:	<p>Standard groups and users configuration: Specify the Oracle software owner user (for example, <code>oracle</code>)</p> <p>Job role separation groups and users configuration: Specify the Grid Infrastructure software owner (for example, <code>grid</code>)</p>
Default group to own the driver interface:	<p>Standard groups and users configuration: Specify the OSDBA group for the database (for example, <code>dba</code>).</p> <p>Job role separation groups and users configuration: Specify the OSASM group for storage administration (for example, <code>asmadmin</code>).</p>

Prompt	Suggested Response
Start Oracle Automatic Storage Management Library driver on boot (y/n):	Enter <i>y</i> to start the Oracle Automatic Storage Management library driver when the system starts.
Fix permissions of Oracle ASM disks on boot? (y/n)	Enter <i>y</i> to fix permissions of 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 Automatic Storage Management library to communicate with the Automatic Storage Management 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.

3.3.1.4.2 Configuring Disk Devices to Use ASM Library Driver on x86 Systems To configure the disk devices to use in an Automatic Storage Management disk group, follow these steps:

1. If you intend to use IDE, SCSI, or RAID devices in the Automatic Storage Management 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. 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:

Disk Type	Device Name Format	Description
IDE disk	<code>/dev/hdxn</code>	In this example, <i>x</i> is a letter that identifies the IDE disk and <i>n</i> is the partition number. For example, <code>/dev/hda</code> is the first disk on the first IDE bus.
SCSI disk	<code>/dev/sdxn</code>	In this example, <i>x</i> is a letter that identifies the SCSI disk and <i>n</i> is the partition number. For example, <code>/dev/sda</code> is the first disk on the first SCSI bus.

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 Automatic Storage Management 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 Automatic Storage Management, 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 Automatic Storage Management disks.

3.3.1.4.3 Administering the ASM Library Driver and Disks To administer the Automatic Storage Management library driver and disks, use the `oracleasm` initialization script with different options, as described in [Table 3-7](#).

Table 3-7 ORACLEASM Script Options

Option	Description
configure	Use the <code>configure</code> option to reconfigure the Automatic Storage Management library driver, if necessary: <pre># /usr/sbin/oracleasm configure -i</pre> <p>To see command options, enter <code>oracleasm configure</code> without the <code>-i</code> flag.</p>
enable disable	Use the <code>disable</code> and <code>enable</code> options to change the actions of the Automatic Storage Management library driver when the system starts. The <code>enable</code> option causes the Automatic Storage Management library driver to load when the system starts: <pre># /usr/sbin/oracleasm enable</pre>
start stop restart	Use the <code>start</code> , <code>stop</code> , and <code>restart</code> options to load or unload the Automatic Storage Management library driver without restarting the system: <pre># /usr/sbin/oracleasm restart</pre>
createdisk	Use the <code>createdisk</code> option to mark a disk device for use with the Automatic Storage Management library driver and give it a name: <pre># /usr/sbin/oracleasm createdisk DISKNAME devicename</pre>

Table 3-7 (Cont.) ORACLEASM Script Options

Option	Description
deletedisk	Use the <code>deletedisk</code> option to unmark a named disk device: <pre># /usr/sbin/oracleasm deletedisk DISKNAME</pre> <p>Caution: Do not use this command to unmark disks that are being used by an Automatic Storage Management disk group. You must delete the disk from the Automatic Storage Management disk group before you unmark it.</p>
querydisk	Use the <code>querydisk</code> option to determine if a disk device or disk name is being used by the Automatic Storage Management library driver: <pre># /usr/sbin/oracleasm querydisk {DISKNAME devicename}</pre>
listdisks	Use the <code>listdisks</code> option to list the disk names of marked Automatic Storage Management library driver disks: <pre># /usr/sbin/oracleasm listdisks</pre>
scandisks	Use the <code>scandisks</code> option to enable cluster nodes to identify which shared disks have been marked as Automatic Storage Management library driver disks on another node: <pre># /usr/sbin/oracleasm scandisks</pre>

3.3.1.5 Configuring Disk Devices Manually for Oracle ASM

By default, the 2.6 kernel 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 disks or Oracle Cluster Registry partitions, corrupting them when the server is restarted. For example, a voting disk on a device named `/dev/sdd` owned by the user `grid` may be on a device named `/dev/sdf` owned by `root` after restarting the server.

If you use ASMLIB, then you do not need to ensure permissions and device path persistency in `udev`.

If you do not use ASMLIB, 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`.

Where rules files describe the same devices, on Asianux, Red Hat, and Oracle Enterprise Linux, the last file read is the one that is applied. On SUSE 2.6 kernels, the first file read is the one that is applied.

To configure a permissions file for disk devices, complete the following tasks:

1. 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
```

2. 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 grid infrastructure software. For example, using the installation owner `grid` and using a role-based group configuration, with the OSASM group `asmadmin`:

```
# vi /etc/udev/rules.d/99-oracle-asmdevices.rules

KERNEL=="sd?1", BUS=="scsi", PROGRAM=="/sbin/scsi_id",
RESULT=="14f70656e66696c00000000", OWNER="grid", GROUP="asmadmin", MODE="0660"
KERNEL=="sd?2", BUS=="scsi", PROGRAM=="/sbin/scsi_id",
RESULT=="14f70656e66696c00000000", OWNER="grid", GROUP="asmadmin", MODE="0660"
KERNEL=="sd?3", BUS=="scsi", PROGRAM=="/sbin/scsi_id",
RESULT=="14f70656e66696c00000000", OWNER="grid", GROUP="asmadmin", MODE="0660"
```

3. 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
```

4. Load updated block device partition tables on all member nodes of the cluster, using `/sbin/partprobe devicename`. You must do this as `root`.
5. Enter the command to restart the UDEV service.

On Asianux, OEL5, and RHEL5, the commands are:

```
# /sbin/udevcontrol reload_rules
# /sbin/start_udev
```

On SLES10, the command is:

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

Check to ensure that your system is configured correctly.

3.3.2 Using Diskgroups with Oracle Database Files on ASM

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

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

3.3.2.1 Identifying and Using Existing Oracle Database Diskgroups on ASM

The following section describes how to identify existing diskgroups and determine the free disk space that they contain.

- Optionally, identify failure groups for the Automatic Storage Management disk group devices.

Note: Complete this step only if you intend to use an installation method that runs Automatic Storage Management Configuration Assistant (ASMCA) in interactive mode; for example, if you intend to choose the Advanced database configuration option, then complete this step. Typical installation does not enable you to specify failure groups.

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.

3.3.2.2 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 Automatic Storage Management 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. Automatic Storage Management expects each disk group device to be on a separate physical disk.
- Although you can specify a logical volume as a device in an Automatic Storage Management disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing Automatic Storage Management from optimizing I/O across the physical devices. They are not supported with Oracle RAC.

3.3.3 Configuring Oracle Automatic Storage Management Cluster File System (ACFS)

Oracle ACFS is installed as part of an Oracle grid infrastructure installation (Oracle Clusterware and Automatic Storage Management) for 11g release 2 (11.2).

Note: Oracle ACFS is supported only on Oracle Enterprise Linux 5.0 and Red Hat Enterprise Linux 5.0. All other Linux releases supported with Oracle grid infrastructure for a cluster 11g release 2 (11.2) are not supported for Oracle ACFS.

To configure Automatic Storage Management Cluster File System for an Oracle Database home for an Oracle RAC database:

1. Install Oracle grid infrastructure for a cluster (Oracle Clusterware and Automatic Storage Management)
2. Change directory to the grid infrastructure home. For example:


```
$ cd /u01/app/11.2.0/grid
```
3. Start ASM Configuration Assistant as the grid installation owner. For example:


```
./asmca
```
4. 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.
5. On the ASM Cluster File Systems page, right-click the Data disk, then select **Create ACFS for Database Home**.

6. 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 mountpoint. For example: `/u02/acfsmounts/dbase_01`
Make a note of this mountpoint 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.
7. During Oracle RAC installation, ensure that you or the DBA who installs Oracle RAC selects for the Oracle home the mountpoint you provided in the Database Home Mountpoint field (in the preceding example, `/u02/acfsmounts/dbase_01`).

See Also: *Oracle Database Storage Administrator's Guide* for more information about configuring and managing your storage with Oracle ACFS

3.3.4 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 Automatic Storage Management Configuration Assistant (ASMCA, located in the path `Grid_home/bin`) to upgrade the existing Oracle ASM instance to 11g release 2 (11.2), and subsequently configure failure groups, ASM volumes and Automatic Storage Management Cluster File System (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 ASM home, then after installing the Oracle ASM 11g release 2 (11.2) binaries, you can start ASMCA to upgrade the existing Oracle ASM instance. You can then choose to configure an ACFS deployment by creating ASM volumes and using the upgraded Oracle ASM to create the ACFS.

On an existing Oracle Clusterware or Oracle RAC installation, if the prior version of Oracle ASM instances on all nodes is 11g release 1, then you are provided with the option to perform a rolling upgrade of Oracle ASM instances. If the prior version of Oracle ASM instances on an Oracle RAC installation are from a release prior to 11g release 1, then rolling upgrades cannot be performed. Oracle ASM on all nodes will be upgraded to 11g release 2 (11.2).

3.3.5 Converting Standalone Oracle ASM Installations to Clustered Installations

If you have existing standalone Oracle ASM installations on one or more nodes you select as 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 disks) 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 11g release 2 (11.2) 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, diskgroup names on the cluster-enabled Oracle ASM instances must be different from existing standalone diskgroup names.

3.4 Desupport of Block and Raw Devices

With the release of Oracle Database 11g release 2 (11.2) and Oracle RAC 11g release 2 (11.2), using Database Configuration Assistant or the installer to store Oracle Clusterware or Oracle Database files on block or raw devices is not supported.

If you intend to upgrade an existing Oracle RAC database, or an Oracle RAC database with Oracle ASM instances, then you can use an existing raw or block device partition, and perform a rolling upgrade of your existing installation. Performing a new installation using block or raw devices is not allowed.

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 Automatic Storage Management. 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:

- [Preparing to Install Oracle Grid Infrastructure with OUI](#)
- [Installing Grid Infrastructure](#)
- [Installing Grid Infrastructure Using a Software-Only Installation](#)
- [Confirming Oracle Clusterware Function](#)
- [Confirming Oracle ASM Function for Oracle Clusterware Files](#)

4.1 Preparing to Install Oracle Grid Infrastructure with OUI

Before you install Oracle grid infrastructure with the installer, use the following checklist to ensure that you have all the information you will need during installation, and to ensure that you have completed all tasks that must be done before starting your installation. Check off each task in the following list as you complete it, and write down the information needed, so that you can provide it during installation.

Shut Down Running Oracle Processes

You may need to shut down running Oracle processes:

Installing 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 software.

Installing on a node that already has a standalone Oracle Database 11g release 2 (11.2) installation running on Oracle ASM: Stop the existing Oracle ASM instances. After the software is installed, start the Oracle ASM instances again.

Installing 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.

Note: If you are upgrading an Oracle RAC 9i release 2 (9.2) node, and the TNSLSNR is listening to the same port on which the SCAN listens (default 1521), then the TNSLSNR should be shut down.

If a Global Services Daemon (GSD) from Oracle9i Release 9.2 or earlier is running, then stop it before installing Oracle grid infrastructure by running the following command:

```
$ Oracle_home/bin/gsdctl stop
```

where *Oracle_home* is the Oracle Database home that is running the GSD.

Caution: If you have an existing Oracle9i release 2 (9.2) Oracle Cluster Manager (Oracle CM) installation, then *do not* shut down the Oracle CM service. Shutting down the Oracle CM service prevents the Oracle grid infrastructure 11g release 2 (11.2) software from detecting the Oracle9i release 2 node list, and causes failure of the Oracle grid infrastructure installation.

Note: If you receive a warning to stop all Oracle services after starting OUI, then run the command

```
Oracle_home/bin/localconfig delete
```

where *Oracle_home* is the existing Oracle Clusterware home.

❑ **Prepare for Automatic Storage Management and Oracle Clusterware Upgrade If You Have Existing Installations**

During installation, you can upgrade existing Oracle Clusterware and Oracle ASM installations. You can upgrade an existing standalone Oracle ASM installation if it is on the node where you are installing Oracle grid infrastructure 11g release 2 (11.2).

When all member nodes of the cluster are running Oracle grid infrastructure 11g release 2 (11.2), then 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 11g release 2 (11.2) on all cluster nodes before you install the Oracle Database 11g release 2 (11.2) version of Oracle RAC.

Note: All Oracle grid infrastructure upgrades (upgrades of existing Oracle Clusterware and Oracle ASM installations) are out-of-place upgrades.

❑ **Determine the Oracle Inventory (oraInventory) location**

If you have already installed Oracle software on your system, then OUI detects the existing Oracle Inventory (oraInventory) directory from the `/etc/orainst.loc` file, and uses this location. This 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.

If you are installing Oracle software for the first time on your system, and your system does not have an oraInventory directory, then the installer 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.

Note: The oraInventory directory cannot be placed on a shared file system.

See Also: The preinstallation chapters in [Chapter 2](#) for information about creating the Oracle Inventory, and completing required system configuration

❑ **Obtain root account access**

During installation, you are asked to run configuration scripts as the `root` user. You must run these scripts as `root`, or be prepared to have your system administrator run them for you. Note that these scripts must be run in sequence. If you attempt to run scripts simultaneously, then the installation will fail.

❑ **Decide if you want to install other languages**

During installation, you are asked if you want translation of user interface text into languages other than the default, which is English.

Note: If the language set for the operating system is not supported by the installer, then by default the installer runs in the English language.

See Also: *Oracle Database Globalization Support Guide* for detailed information on character sets and language configuration

❑ **Determine your cluster name, public node names, the SCAN, virtual node names, GNS VIP and planned interface use for each node in the cluster**

During installation, you are prompted to provide the public and virtual hostname, unless you use a third party cluster software. In that case, the public hostname information will be filled in. You are also prompted to identify which interfaces are public, private, or interfaces in use for another purpose, such as a network file system.

If you use Grid Naming Service (GNS), then OUI displays the public and virtual hostname addresses labeled as "AUTO" because they are configured automatically.

Note: 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 hostname 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.

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.

In addition:

- Provide a cluster name with the following characteristics:
 - * It must be globally unique throughout your host domain.
 - * It must be at least one character long and less than 15 characters long.
 - * It must 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.
- 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*.
- 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.

❑ **Identify shared storage for Oracle Clusterware files and prepare storage if necessary**

During installation, you are asked to provide paths for the following Oracle Clusterware files. These files must be shared across all nodes of the cluster, either on Automatic Storage Management, or on a supported file system:

- Voting disks are files that Oracle Clusterware uses to verify cluster node membership and status.

Voting disk files must be owned by the user performing the installation (`oracle` or `grid`), and must have permissions set to 640.

- Oracle Cluster Registry files (OCR) contain cluster and database configuration information for Oracle Clusterware.

Before installation, OCR files must be owned by the user performing the installation (`grid` or `oracle`). That installation user must have `oinstall` as its primary group. During installation, OUI changes ownership of the OCR files to `root`.

If your file system does not have external storage redundancy, then Oracle recommends that you provide two additional locations for the OCR disk, and two additional locations for the voting disks, for a total of six partitions (three for OCR, and three for voting disks). Creating redundant storage locations protects the OCR and voting disk in the event of a failure. To completely protect your cluster, the storage locations given for the copies of the OCR and voting disks should have completely separate paths, controllers, and disks, so that no single point of failure is shared by storage locations.

When you select to store the OCR on Oracle ASM, the default configuration is to create the OCR on one ASM diskgroup. If you create the disk group with normal or high redundancy, then the OCR is protected from physical disk failure.

To protect the OCR from logical disk failure, create another ASM diskgroup after installation and add the OCR to the second diskgroup using the `ocrconfig` command.

See Also: [Chapter 2, "Advanced Installation Oracle Grid Infrastructure for a Cluster Preinstallation Tasks"](#) and *Oracle Database Storage Administrator's Guide* for information about adding disks to diskgroups

❑ **Ensure cron jobs do not run during installation**

If the installer is running when daily cron jobs start, then you may encounter unexplained installation problems if your cron job is performing cleanup, and temporary files are deleted before the installation is finished. Oracle recommends that you complete installation before daily cron jobs are run, or disable daily cron jobs that perform cleanup until after the installation is completed.

❑ **Have IPMI Configuration completed and have IPMI administrator account information**

If you intend to use IPMI, then ensure BMC interfaces are configured, and have an administration account username and password to provide when prompted during installation.

For nonstandard installations, if you must change configuration on one or more nodes after installation (for example, if you have different administrator usernames 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 as described in [Chapter 5](#).

❑ **Ensure that the Oracle home path you select for the grid infrastructure home uses only ASCII characters**

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.

❑ **Unset Oracle environment variables.** If you have set `ORA_CRS_HOME` as an environment variable, then unset it before starting an installation or upgrade. You should never use `ORA_CRS_HOME` as an environment variable.

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`

4.2 Installing 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 Grid Infrastructure](#)
- [Installing Grid Infrastructure Using a Cluster Configuration File](#)

4.2.1 Running OUI to Install Grid Infrastructure

Complete the following steps to install grid infrastructure (Oracle Clusterware and 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. Change to the `/Disk1` directory on the installation media, or where you have downloaded the installation binaries, and start the `runInstaller` command. For example:

```
$ cd /home/grid/oracle_sw/Disk1
$ ./runInstaller
```

2. Select **Typical** or **Advanced** installation.
3. Provide information or run scripts as `root` when prompted by OUI. If you need assistance during installation, click **Help**. Click **Details** to see the log file.

Note: You must run the `root.sh` script on the first node and wait for it to finish. If your cluster has four or more nodes, then `root.sh` can be run concurrently on all nodes but the first and last. As with the first node, the `root.sh` script on the last node must be run separately.

4. After you run `root.sh` on all the nodes, OUI runs Net Configuration Assistant (`netca`) and Cluster Verification Utility. These programs run without user intervention.
5. Automatic Storage Management Configuration Assistant (`asmca`) configures Oracle ASM during the installation.

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.

If you intend to install Oracle Database 11g release 2 (11.2) with Oracle RAC, then refer to *Oracle Real Application Clusters Installation Guide for Linux*.

See Also: *Oracle Real Application Clusters Administration and Deployment Guide* for information about using cloning and node addition procedures, and *Oracle Clusterware Administration and Deployment Guide* for cloning Oracle grid infrastructure

4.2.2 Installing Grid Infrastructure Using a Cluster Configuration File

During installation of 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 information about the cluster name and node names that it requires to configure the cluster.

Oracle suggests 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:

1. On the installation media, navigate to the directory `/response`.
2. Using a text editor, open the response file `crs_install.rsp`.
3. Follow the directions in that section for creating a cluster configuration file.

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 Automatic Storage Management (ASM) during installation, then ASM default discovery finds all disks that ASMLIB marks as ASM disks.

4.3 Installing Grid Infrastructure Using a Software-Only Installation

Note: Oracle recommends that only advanced users should perform the software-only installation, as this installation method provides no validation of the installation, and as this installation option requires manual postinstallation steps to enable the grid infrastructure software.

A software-only installation consists of installing Oracle grid infrastructure for a cluster on one node, configuring it, and then running the installer on each node that you want to have as a cluster member node, and then joining other nodes to the cluster.

To perform a software-only installation:

4.3.1 Installing the Software Binaries

1. Start the `runInstaller` command from the relevant directory on the Oracle Database 11g release 2 (11.2) installation media or download directory. For example:

```
$ cd /home/grid/oracle_sw/Disk1
$ ./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. Follow

the output of the `root.sh` script, and run the command `Grid_home/crs/install/rootcrs.pl` or `Grid_home/crs/install/rootcrs.pl` as required. For example:

```
$ cd /u01/app/grid/11.2.0/crs/install
$ perl rootcrs.pl
```

Use `Grid_home/crs/install/rootcrs.pl` to configure Oracle Grid Infrastructure for a standalone server. Use `Grid_home/crs/install/rootcrs.pl` to configure Oracle Grid Infrastructure for a cluster.

Note: Oracle grid infrastructure can be used for standalone servers and for clusters. However, if you first configure Oracle grid infrastructure for a standalone server, and then decide you want to configure Oracle grid infrastructure for a cluster, then you must re-link the Oracle software before you run `rootcrs.pl` to configure Oracle grid infrastructure for a clusters. The **Install Grid Infrastructure Software Only** installation option does not assume a cluster configuration, and therefore does not automatically link the Oracle RAC option.

5. To relink Oracle Clusterware with the Oracle RAC option enabled, run commands similar to the following (in this example, the Grid home is `/u01/app/grid/11.2.0`):

```
$ cd /u01/app/grid/11.2.0
$ set env ORACLE_HOME pwd
$ cd rdbms/lib
$ make -f ins_rdbms.mk rac_on ioracle
```

6. On each remaining node, verify that the cluster node meets installation requirements using the command `runcluvfy.sh stage -pre crsinst`. Ensure that you have completed all storage and server preinstallation requirements.
7. 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.
8. If required relink the Oracle RAC binaries as described in step 5 on every node where you installed the Oracle grid infrastructure software.

4.3.2 Configuring the Software Binaries

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.

To configure and activate software-only grid infrastructure for a cluster installations, complete the following tasks:

1. Using a text editor, modify the template file `/Grid_home/install/crs/crsconfig_params` to create a parameter file for the installer to use to configure the cluster. For example:

```
ORACLE_OWNER=grid
ORA_DBA_GROUP=oinstall
```



```

ORA_ASM_GROUP=asm
LANGUAGE_ID='AMERICAN_AMERICA.WE8ISO8859P1'
ORACLE_HOME=/u01/crs
ORACLE_BASE=/u01/crsbase
OCR_LOCATIONS=/u02/stor1/ocr,/u03/stor2/ocr
CLUSTER_NAME=example_cluster
HOST_NAME_LIST=node1,node2
NODE_NAME_LIST=node1,node2
VOTING_DISKS=/u02/stor1/vdsk,/u03/stor2/vdsk,/u04/stor3/vdsk
CRS_STORAGE_OPTION=2
CRS_NODEVIPS='node1-vip/255.255.252.0/eth0,node2-vip/255.255.252.0/eth0'
NODELIST=node1,node2
NETWORKS="eth0"/192.0.2.64:public,"eth1"/192.0.2.65:cluster_interconnect
SCAN_NAME=example-scan.domain
SCAN_PORT=1522

```

2. On all nodes, place the `crsconfig_params` file in the path `Grid_home/crs/install/crsconfig_params`, where `Grid_home` is the path to the Oracle grid infrastructure home for a cluster. For example:

```
$ cp crsconfig_params /u01/app/11.2.0/grid/crs/install/crsconfig_params
```

3. After configuring the `crsconfig_params` file, log in as root, and run the script `Grid_home/crs/install/rootcrs.pl`, using the following syntax:

```
Grid_home/perl/lib/perl -IGRID_HOME/perl/lib -IGrid_home/crs/install Grid_
home/crs/install/rootcrs.pl
```

For example:

```
# /u01/app/grid/11.2.0/perl/lib/perl -I/u01/app/grid/11.2.0/perl/lib \
-I/u01/app/grid/11.2.0/crs/install /u01/app/grid/11.2.0/crs/install/rootcrs.pl
```

4. Change directory to `Grid_home/oui/bin`, where `Grid_home` is the path of the Grid Infrastructure home on each cluster member node.
5. Enter the following command syntax, where `Grid_home` is the path of the Grid Infrastructure home on each cluster member node, and `node_list` is a comma-delimited list of nodes on which you want the software enabled:

```
runInstaller -updateNodeList ORACLE_HOME=Grid_home -defaultHomeName
```

For example

```
$ ./runInstaller -updateNodeList ORACLE_HOME=/u01/app/11.2.0/grid
-defaultHomeName
"CLUSTER_NODES={node_list}" CRS=TRUE
```

To enable the Oracle Clusterware installation on the local node only, enter the following command, where `Grid_home` is the Grid home on the local node, and `node_list` is a comma-delimited list of nodes on which you want the software enabled:

```
runInstaller -updateNodeList ORACLE_HOME=Grid_home -defaultHomeName
"CLUSTER_NODES={node_list}" CRS=TRUE
```

For example:

```
$ ./runInstaller -updateNodeList ORACLE_HOME=/u01/app/11.2.0/grid
-defaultHomeName
"CLUSTER_NODES={node_list}" CRS=TRUE
```

4.4 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 crs
```

For example:

```
$ crsctl check crs
```

```
CRS-4638: Oracle High Availability Services is online
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.

4.5 Confirming Oracle ASM Function for Oracle Clusterware Files

If you installed the OCR and voting disk files on Oracle ASM, then use the following command syntax as the Grid Infrastructure installation owner to confirm that your Oracle ASM installation is running:

```
srvctl status asm
```

For example:

```
$ srvctl status asm
ASM instance +ASM1 is running on node node1.
ASM instance +ASM2 is running on node node2.
```

Oracle ASM is running only if it is needed for Oracle Clusterware files. If you have not installed OCR and voting disks 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.

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 Older Oracle Database Versions with Grid Infrastructure](#)
- [Modifying Oracle Clusterware Binaries After Installation](#)

5.1 Required Postinstallation Tasks

You must perform the following tasks after completing your installation:

- [Download and Install Patch Updates](#)

Note: In prior releases, backing up the voting disks using a `dd` command was a required postinstallation task. With Oracle Clusterware release 11.2 and later, backing up and restoring a voting disk using the `dd` command may result in the loss of the voting disk, so this procedure is not supported.

5.1.1 Download and Install Patch Updates

Refer to the My Oracle Support Web site for required patch updates for your installation.

Note: Browsers require an Adobe Flash plug-in, version 9.0.115 or higher to use My Oracle Support. Check your browser for the correct version of Flash plug-in by going to the Adobe Flash checker page, and installing the latest version of Adobe Flash.

If you do not have Flash installed, then download the latest version of the Flash Player from the Adobe Web site:

<http://www.adobe.com/go/getflashplayer>

To download required patch updates:

1. Use a Web browser to view the My Oracle Support Web site:

<https://metalink.oracle.com>

2. Log in to My Oracle Support Web site.

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 & Update page, click **Advanced Search**.
5. On the Advanced Search page, click the search icon next to the Product or Product Family field.
6. In the Search and Select: Product Family field, select Database and Tools in the Search list field, enter RDBMS Server in the text field, and click **Go**.

RDBMS Server appears in the Product or Product Family field. The current release appears in the Release field.
7. Select your platform from the list in the Platform field, and at the bottom of the selection list, click **Go**.
8. Any available patch updates appear under the Results heading.
9. Click the patch number to download the patch.
10. On the Patch Set page, click **View README** and read the page that appears. 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 11g release 2 (11.2) 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. Refer to [Appendix F](#) on page F-1 for information about how to stop database processes in preparation for installing patches.

5.2 Recommended Postinstallation Tasks

Oracle recommends that you complete the following tasks as needed after installing Oracle grid infrastructure:

- [Back Up the root.sh Script](#)
- [Install Cluster Health Management](#)
- [Tune Semaphore Parameters](#)
- [Create a Fast Recovery Area Disk Group](#)

5.2.1 Back Up the root.sh Script

Oracle recommends that you back up the `root.sh` script after you complete an installation. If you install other products in the same Oracle home directory, then the installer updates the contents of the existing `root.sh` script during the installation. If you require information contained in the original `root.sh` script, then you can recover it from the `root.sh` file copy.

5.2.2 Install Cluster Health Management

To address troubleshooting issues, Oracle recommends that you install either Instantaneous Problem Detection OS Tool (IPD/OS), or OS Watcher and RACDDT, depending on your Linux kernel version.

5.2.2.1 Installing Instantaneous Problem Detection OS Tool (IPD/OS)

On Linux systems with Linux kernels greater than or equal to 2.6.9, install the Oracle Instantaneous problem Detection Operating System Tool (IPD/OS).

The IPD/OS tool is designed to detect and analyze operating system and cluster resource-related degradation and failures. The tool can provide better explanations for many issues that occur in clusters where Oracle Clusterware, Oracle ASM and Oracle RAC are running, such as node evictions. It tracks the operating system resource consumption at each node, process, and device level continuously. It collects and analyzes clusterwide data. In real time mode, when thresholds are reached, an alert is shown to the operator. For root cause analysis, historical data can be replayed to understand what was happening at the time of failure.

You can download the tool at the following URL:

http://www.oracle.com/technology/products/database/clustering/ipd_download_homepage.html

5.2.2.2 Installing OS Watcher and RACDDT

On Linux systems with Linux kernels earlier than 2.6.9, install OS Watcher to help resolve operating system issues with your cluster. If you intend to install an Oracle RAC database, then also install RACDDT. You must have access to My Oracle Support to download OS Watcher and RACDDT.

OS Watcher (OSW) is a collection of UNIX/Linux shell scripts that collect and archive operating system and network metrics to aid Oracle Support in diagnosing various issues related to system and performance. OSW operates as a set of background processes on the server and gathers operating system data on a regular basis. The scripts use common utilities such as `vmstat`, `netstat` and `iostat`.

RACDDT is a data collection tool designed and configured specifically for gathering diagnostic data related to Oracle RAC technology. RACDDT is a set of scripts and configuration files that is run on one or more nodes of an Oracle RAC cluster. The main script is written in Perl, while a number of proxy scripts are written using Korn shell. RACDDT will run on all supported UNIX and Linux platforms, but is not supported on any Windows platforms.

OSW is also included in the RACDDT script file, but is not installed by RACDDT. OSW must be installed on each node where data is to be collected.

To download binaries for OS Watcher and RACDDT, go to the following URL:

<https://metalink.oracle.com>

Download OSW by searching for OS Watcher, and downloading the binaries from the User Guide bulletin. Installation instructions for OSW are provided in the user guide. Download RACDDT by searching for RACDDT, and downloading the binaries from the RACDDT User Guide bulletin.

5.2.3 Tune Semaphore Parameters

Refer to 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 `semtns` (total semaphores systemwide) to this total.
3. Set `semtns1` (semaphores for each set) to 256.
4. Set `semtnsi` (total semaphore sets) to `semtns` divided by `semtns1`, rounded up to the nearest multiple of 1024.

5.2.4 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.

5.2.4.1 About the Fast Recovery Area and the Fast Recovery Area Disk Group

The Fast Recovery Area is a unified storage location for all Oracle Database files related to recovery. 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 Flash 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 flash recovery files in the same disk group. However, Oracle recommends that you create a separate Flash Recovery 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 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 database1 is your least important database, database 2 is of greater importance and database 3 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 database 1, 50 GB for database 2, and 70 GB for database 3.

See Also: *Oracle Database Storage Administrator's Guide*

5.2.4.2 Creating the Fast Recovery Area Disk Group

To create a flash recovery file disk group:

1. Navigate to the Grid home bin directory, and start ASM Configuration Assistant (`asmca`). For example:

```
$ cd /u01/app/11.2.0/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**.

5.3 Using Older Oracle Database Versions with Grid Infrastructure

Review the following sections for information about using older Oracle Database releases with 11g release 2 (11.2) grid infrastructure installations:

- [General Restrictions for Using Older Oracle Database Versions](#)
- [Pinning Cluster Nodes for Oracle Database Release 10.x or 11.x](#)
- [Enabling The Global Services Daemon \(GSD\) for Oracle Database Release 9.2](#)
- [Using the Correct LSNRCTL Commands](#)

5.3.1 General Restrictions for Using Older Oracle Database Versions

You can use Oracle Database release 9.2, release 10.x and release 11.1 with Oracle Clusterware release 11.2.

If you upgrade an existing version of Oracle Clusterware, then required configuration of existing databases is completed automatically. However, if you complete a new installation of Oracle grid infrastructure for a cluster, and then want to install a version of Oracle Database prior to 11.2, then you must complete additional manual configuration tasks.

Note: Before you start an Oracle RAC or Oracle Database installation on an Oracle Clusterware release 11.2 installation, if you upgraded from releases 11.1.0.7, 11.1.0.6, and 10.2.0.4, then you must install the one-off patch required for that release. See [Section 5.1.1, "Download and Install Patch Updates"](#) on page 5-1.

5.3.2 Pinning Cluster Nodes for Oracle Database Release 10.x or 11.x

When Oracle Database version 10.x or 11x is installed on a new Oracle grid infrastructure for a cluster configuration, it is configured for dynamic cluster configuration, in which some or all IP addresses are provisionally assigned, and other cluster identification information is dynamic. This configuration is incompatible with older database releases, which require fixed addresses and configuration.

You can change the nodes where you want to run the older database to create a persistent configuration. Creating a persistent configuration for a node is called pinning a node.

To pin a node in preparation for installing an older 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/11.2.0/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/11.2.0/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

5.3.3 Enabling The Global Services Daemon (GSD) for Oracle Database Release 9.2

When Oracle Database 9i release 2 (9.2) is installed on an 11g release 2 (11.2) Oracle grid infrastructure for a cluster configuration, the Global Services daemon (GSD) is disabled by default. Use the following commands to enable the GSD before you install a release 9.2 Oracle Database:

```
srvctl enable nodeapps -g
srvctl start nodeapps
```

5.3.4 Using the Correct LSNRCTL Commands

To administer 11g release 2 local and scan listeners using the `lsnrctl` command, set your `$ORACLE_HOME` environment variable to the path for the 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.

5.4 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 Exadata configuration to run IPC traffic over RDS on the interconnect instead of using the default UDP, then you must unlock the Grid home.

Caution: Before relinking executables, you must shut down all executables that run in the Oracle home directory that you are relinking. In addition, shut down applications linked with Oracle shared libraries.

Unlock the home using the following procedure:

1. Change directory to the path `Grid_home/crs/install`, where `Grid_home` is the path to the Grid home, and unlock the Grid home using the command `rootcrs.pl -unlock -crshome Grid_home`, where `Grid_home` is the path to your Grid infrastructure home. For example, with the grid home `/u01/app/11.2.0/grid`, enter the following command:

```
# cd /u01/app/11.2.0/grid/crs/install
# perl rootcrs.pl -unlock -crshome /u01/app/11.2.0/grid
```

2. Change user to the grid infrastructure software owner, and relink binaries using the command syntax `make -f Grid_home/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/lib/ins_rdbms.mk ipc_rds ioracle
```

Note: To relink binaries, you can also change to the grid installation owner and run the command `Grid_home/bin/relink`.

3. Relock the Grid home and restart the cluster using the following command:

```
# perl rootcrs.pl -patch
```
4. Repeat steps 1 through 3 on each cluster member node.

How to Modify or Deinstall Oracle Grid Infrastructure

This chapter describes how to remove Oracle Clusterware and Oracle ASM.

This chapter contains the following topics:

- [Deciding When to Deinstall Oracle Clusterware](#)
- [Migrating From Oracle Restart to Oracle Clusterware](#)
- [Adding Standalone Grid Infrastructure Servers to a Cluster](#)
- [Relinking Oracle Grid Infrastructure for a Cluster Binaries](#)
- [Removing Oracle Clusterware and ASM](#)

See Also: Product-specific documentation for requirements and restrictions to remove an individual product

6.1 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 Clusterware installation, either in an educational environment, or a test environment.
- You have successfully installed Oracle Clusterware, but you want to downgrade to a previous release.
- 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.

6.2 Migrating From Oracle Restart to Oracle Clusterware

If you have Oracle Restart (Oracle grid infrastructure for a standalone server) installed on a server you want to use for an Oracle Clusterware installation, then you must remove the Oracle Restart installation, and install Oracle grid infrastructure for a cluster, which includes Oracle Clusterware. You cannot install Oracle Clusterware on a node where Oracle Restart is configured, and you cannot install Oracle Clusterware on a cluster where any of the prospective nodes have Oracle Restart configured.

Use the following procedure to deconfigure Oracle Restart and install Oracle Clusterware:

1. Run `srvctl config` commands to inspect the configuration of databases and listeners using Oracle Restart. Write down the configuration information. For example:

```
$ srvctl config listener
```
2. Deconfigure Oracle Restart.
3. Prepare the server for Oracle Clusterware configuration, as described in this document.
4. Install and configure Oracle grid infrastructure, with Oracle Clusterware and Oracle ASM.
5. Use `srvctl add` commands to add the databases, listeners, and other configuration information previously supported by Oracle Restart.

6.3 Adding 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:

1. Inspect the Oracle 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 -d db_unique_name
srvctl config service -d db_unique_name
srvctl config listener -l lsnrname
```

Write down the configuration information for the server.

2. Change directory to *Grid home/crs/install*. For example:

```
# cd /u01/app/11.2.0/grid/crs/install
```

3. Deconfigure and deinstall the Oracle grid infrastructure installation for a standalone server, using the following command:

```
# roothas.pl -deconfig
```

4. Clone the Oracle grid infrastructure for a cluster software from an existing node.
5. 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 -d db_unique_name -o $ORACLE_HOME -x nodename
```

For example, with the database name *mydb_node1*, and the *nodename* *node1*, enter the following command:

```
srvctl add database -d mydb_node1 -o $ORACLE_HOME -x node1
```

6. Add each service to the database, using the command `srvctl add service`.

6.4 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 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.

As root:

```
# cd Grid_home/crs/install
# perl rootcrs.pl -unlock
```

As the grid infrastructure for a cluster owner:

```
$ export ORACLE_HOME=Grid_home
$ Grid_home/bin/relink
```

As root again:

```
# cd Grid_home/crs/install
# perl rootcrs.pl -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.

6.5 Removing Oracle Clusterware and ASM

The `deinstall` command removes Oracle Clusterware and ASM from your server. The following sections describe the command, and provide information about additional options to use the command:

- [About the Deinstallation Tool](#)
- [Example of Running the Deinstall Command for Oracle Clusterware and ASM](#)
- [Example of a Deinstallation Parameter File for Grid Infrastructure for a Cluster](#)

6.5.1 About the Deinstallation Tool

The Deinstallation Tool (`deinstall`) is available in the installation media before installation, and is available in Oracle home directories after installation. It is located in the path `$ORACLE_HOME/deinstall`.

The `deinstall` command stops Oracle software, and removes Oracle software and configuration files on the operating system.

The command uses the following syntax, where variable content is indicated by italics:

```
deinstall -home complete path of Oracle home [-silent] [-checkonly] [-local]
[-paramfile complete path of input parameter property file] [-params name1=value
name2=value . . .] [-o complete path of directory for saving files] [-help | -h]
```

The options are:

- `-home`

Use this flag to indicate the home path of the Oracle home that you want to check or deinstall. To deinstall Oracle software using the `deinstall` command in the Oracle home you plan to deinstall, provide a parameter file in another location, and do not use the `-home` flag.

- `-silent`

Use this flag to run the command in noninteractive mode. This option requires a properties file that contains the configuration values for the Oracle home that is being deinstalled or deconfigured.

To create a properties file and provide the required parameters, refer to the template file `deinstall.rsp.tmpl`, located in the `response` folder. If you prefer, instead of using the template file, you can generate a properties file by using the `-checkonly` option to have `deconfig` discover information from the Oracle home that you want to deinstall and deconfigure. It generates the properties file, which you can then use with the `-silent` option.

- `-checkonly`

Use this flag to check the status of the Oracle software home configuration. Running the command with the `checkonly` flag does not remove the Oracle configuration.

- `-local`

Use this flag on a multinode environment to deconfigure Oracle software in a cluster.

When you run `deconfig` with this flag, it deconfigures and deinstalls the Oracle software on the local node (the node where `deconfig` is run). On remote nodes, it deconfigures Oracle software, but does not deinstall the Oracle software.

- `-paramfile` *complete path of input parameter property file*

Use this flag to run `deconfig` with a parameter file in a location other than the default. When you use this flag, provide the complete path where the parameter file is located.

The default location of the parameter file depends on the location of `deconfig`:

- From the installation media or stage location: `$ORACLE_HOME/inventory/response`.
- From a unzipped archive file from OTN: `/ziplocation/response`.
- After installation from the installed Oracle home: `$ORACLE_HOME/deinstall/response`.

- `-params` [*name1=value name2=value name3=value ...*]

Use this flag with a parameter file to override one or more values in a parameter file you have already created.

- `-o` *complete path of directory for saving response files*

Use this flag to provide a path other than the default location where the properties file (`deinstall.rsp.tmpl`) is saved.

The default location of the parameter file depends on the location of `deconfig`:

- From the installation media or stage location before installation: `$ORACLE_HOME/`
- From a unzipped archive file from OTN: `/ziplocation/response/`.

- After installation from the installed Oracle home: `$ORACLE_HOME/deinstall/response`.
- `-help | -h`
Use the help option (`-help` or `-h`) to obtain additional information about the command option flags.

6.5.2 Example of Running the Deinstall Command for Oracle Clusterware and ASM

As the `deinstall` command runs, you are prompted to provide the home directory of the Oracle software that you want to remove from your system. Provide additional information as prompted.

To run the `deinstall` command from an Oracle grid infrastructure for a cluster home in the path `/u01/app/11.2.0/grid`, where you are running the command using the parameter file in the software owner location `/home/usr/grid`, enter the following command:

```
$ cd /u01/app/11.2.0/grid/deinstall/
$ ./deinstall -paramfile /home/usr/grid/myparamfile.tmpl
```

You can generate the parameter file by running the `deinstall` command using the `-checkonly` flag before you run the command to deinstall the home, or you can use the response file template and manually edit it to create the parameter file to use with the `deinstall` command.

6.5.3 Example of a Deinstallation Parameter File for Grid Infrastructure for a Cluster

You can run the `deinstall` command with the `-paramfile` option to use the values you specify in the parameter file. The following is an example of a parameter 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/11.2.0/grid`, the Oracle base (the Oracle base for 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 are running the deinstallation session from) is `node1`:

```
#Copyright (c) 2005, 2006 Oracle Corporation. All rights reserved.
#Fri Feb 06 00:08:58 PST 2009
LOCAL_NODE=node1
HOME_TYPE=CRS
ASM_REDUNDANCY=\
ORACLE_BASE=/u01/app/grid/
VIP1_MASK=255.255.252.0
NEW_NODEVIPs='node1-vip/255.255.252.0/eth0,node2-vip/255.255.252.0/eth0'
VOTING_DISKS=/u02/storage/grid/vdsk
SCAN_PORT=1522
silent=true
ASM_UPGRADE=false
ORA_CRS_HOME=/u01/app/11.2.0/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/11.2.0/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=nodescn
SHAREJAR_NAME=share.jar
HELPIJAR_NAME=help4.jar
SILENT=false
local=false
INVENTORY_LOCATION=/u01/app/oraInventory
GNS_CONF=false
JEWTJAR_NAME=jwt4.jar
OCR_LOCATIONS=/u02/storage/grid/ocr
EMBASEJAR_NAME=oemlt.jar
ORACLE_HOME=/u01/app/11.2.0/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/11.2.0/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
```


A

Troubleshooting the Oracle Grid Infrastructure Installation Process

This appendix provides troubleshooting information for installing Oracle grid infrastructure.

See Also: The Oracle Database 11g 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:

- [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](#)
- [Interconnect Configuration Issues](#)

A.1 General Installation Issues

The following is a list of examples of types of errors that can occur during installation. It contains the following issues:

- [An error occurred while trying to get the disks](#)
- [Could not execute auto check for display colors using command /usr/X11R6/bin/xdpyinfo](#)
- [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](#)
- [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](#)
- [YPBINDPROC_DOMAIN: Domain not bound](#)

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.

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 grid infrastructure installation owner, then log out of the server, and log back in as the grid installation owner.

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 this fails to work, 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 will be available, and entering the `xclock` command will display the X clock.

When the X clock appears, then 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/nodenumbr/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 older 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 mountpoint `/u02`:

```
#umount /u02
#mount -a -t nfs
#cd $ORA_CRS_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 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.

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=4g 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 `$ORA_CRS_HOME/log/hostname/client/ocrconfig_pid.log`, where `pid` stands for the process id.

Action: Identify a storage device that has 280 MB or more available space. Oracle recommends that you allocate the entire disk to Oracle ASM.

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 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.3 Interpreting CVU Messages About Oracle Grid Infrastructure Setup

If the Cluster Verification Utility 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 Cluster Verification Utility 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 "[Setting Display and X11 Forwarding Configuration](#)" on page 2-51 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 refer to [Chapter 2, "Preventing Installation Errors Caused by stty Commands"](#) on page 2-51 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 refer to 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: ["Creating Groups, Users and Paths for Oracle Grid Infrastructure" in Chapter 2](#) for instructions about how to create required groups, and how to configure the installation owner user

A.4 About the Oracle Clusterware Alert Log

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 might see alert messages from the Cluster Ready Services (CRS) daemon process when it starts, if it aborts, if the failover process fails, or if automatic restart of a CRS resource failed.

Enterprise Manager monitors the Clusterware log file and posts an alert on the Cluster Home page if an error is detected. For example, if a voting disk 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 `CRS_home/log/hostname/alerthostname.log`, where `CRS_home` is the directory in which Oracle Clusterware was installed and `hostname` is the host name of the local node.

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

A.5 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 older 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/technology/products/database/clustering/cvu/cvu_download_homepage.html

A.6 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.7 Interconnect Configuration Issues

If you use multiple network interface cards (NICs) for the interconnect, then the NICs should be bonded 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, then they must be on the same subnet.

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.

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](#)
- [Creating the oraInst.loc File](#)
- [Preparing a Response File](#)
- [Running the Installer Using a Response File](#)
- [Running Net Configuration Assistant Using a Response File](#)
- [Running Database Configuration Assistants Using Response Files](#)
- [Postinstallation Configuration Using a Response File](#)

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" ...
```

This method is particularly useful if you do not want to embed sensitive information, such as passwords, in the response file. For example:

```
-silent "s_dlgRBOPassword=binks342" ...
```

Ensure that you enclose the variable and its setting in quotes.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for more information about response files

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>

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. Create the `oraInst.loc` file if it is not present on the server.
2. Prepare a response file.
3. Run the installer in silent or response file mode.

4. 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.

Creating the oraInst.loc File

If you plan to install Oracle products using the installer in silent or response file mode, and an `oraInst.loc` file does not already exist, then you must manually create the `oraInst.loc` file. This file specifies the location of the Oracle Inventory directory, which is where the installer creates the central inventory of Oracle products installed on the system.

Note: If Oracle software has been installed previously on the system, then the `oraInst.loc` file should already exist. If the file does exist, then you do not need to create this file.

To create the `oraInst.loc` file, follow these steps:

1. Switch user to root:

```
$ su - root
```

2. Change directory:

```
# cd /etc/
```

3. Use a text editor to create the `oraInst.loc` file, containing the following lines:

```
inventory_loc=$ORACLE_BASE/oraInventory
inst_group=oinstall
```

This example assumes that the `$ORACLE_BASE` environment variable for the Oracle software installation owner is set to the path of the Oracle base directory, such as `/u01/app/oracle`.

4. Set the ownership of the `oraInst.loc` file to an Oracle software installation owner, and to members of the `oraInventory` group, and change permissions to 664. For example, if the installation owner is `oracle`, and the `oraInventory` group is `oinstall`, then enter the following commands:

```
# chown oracle:oinstall oraInst.loc
# chmod 664 oraInst.loc
```

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](#)

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 B-1 lists the response files provided with this software:

Table B-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 B-2 Response files for Oracle Grid Infrastructure

Response File	Description
<code>crs_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 if you prefer.

2. Open the response file in a text editor:

```
$ vi /local_dir/response_file.rsp
```

Remember that you can specify sensitive information, such as passwords, at the command line rather than within the response file. ["How Response Files Work"](#) on page B-1 explains this method.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for detailed information on creating response files

3. Follow the instructions in the file to edit it.

Note: The installer or configuration assistant fails if you do not correctly configure the response file.

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.

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: You cannot use record mode to create a response file during an installation that uses the Typical installation method.

To record a response file:

1. Complete preinstallation tasks as for a normal installation.

When you run the installer to record a response file, it checks the system to verify that it meets the requirements to install the software. For this reason, Oracle recommends that you complete all of the required preinstallation tasks and record the response file while completing an installation.

2. If you have not installed Oracle software on this system previously, create the `oraInst.loc` file as described in [Creating the oraInst.loc File](#).
3. Ensure that the Oracle software owner user (typically `oracle`) has permissions to create or write to the Oracle home path that you will specify when you run the installer.
4. On each installation screen, specify the required information.
5. When the installer displays the Summary screen, perform the following:
 - a. Click **Save Response File** and specify a file name and location to save the values for the response file.
 - b. Click **Finish** to create the response file and continue with the installation.

Click **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.

6. If you do not complete the installation, then delete the Oracle home directory that the installer created using the path you specified in the Specify File Locations screen.
7. Before you use the saved response file on another system, edit the file and make any required changes.

Use the instructions in the file as a guide when editing it.

Running the Installer Using a Response File

Now, you are ready to run Oracle Universal Installer at the command line, specifying the response file you created, to perform the installation. 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 after some time.

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, set the `DISPLAY` environment variable.

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 `root.sh` script. For example

```
$ su root
```

```
password:
# /oracle_home_path/root.sh
```

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 if you prefer.

2. Open the response file in a text editor:

```
$ vi /local_dir/netca.rsp
```

3. Follow the instructions in the file to edit it.

Note: Net Configuration Assistant fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user, and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.
5. Enter a command similar to the following to run Net Configuration Assistant in silent mode:

```
$ $ORACLE_HOME/bin/netca /silent /responsefile /local_dir/netca.rsp
```

In this command:

- The `/silent` option indicates 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.

Running Database Configuration Assistants Using Response Files

You can run configuration assistants in response file or silent mode to configure and start Oracle software after it is installed on the system. To run configuration assistants in response file or silent mode, you must copy and edit a response file template.

Note: If you copied the software to a hard disk, then the response file template is located in the `/response` directory.

This section contains the following topics:

- [About the Database Configuration Assistant in Response File Mode](#)
- [Running Database Configuration Assistant in Response File or Silent Mode](#)

About the Database Configuration Assistant in Response File Mode

In the response file mode, Database Configuration Assistant uses values that you specify, in the response file or as command line options, to create a database. As it configures and starts the database, it displays a window that contains status messages and a progress bar. The window that it displays is the same window that is displayed when you choose to create a preconfigured database during an Enterprise Edition or Standard Edition installation.

To run Database Configuration Assistant in response file mode, you must use a graphical display and set the `DISPLAY` environment variable. Use `-progressOnly` flag to set the run mode to response file.

Oracle provides a response file template named `dbca.rsp` in the `/response` directory on the installation media.

Running Database Configuration Assistant in Response File or Silent Mode

To run Database Configuration Assistant in response file or silent mode:

1. Copy the `dbca.rsp` response file template from the response file directory to a directory on your system:

```
$ cp /directory_path/response/dbca.rsp local_directory
```

In this example, `directory_path` is the path of the `database` directory on the DVD. If you have copied the software to a hard drive, you can edit the file in the `response` directory if you prefer.

Note: As an alternative to editing the response file template, you can also create a database by specifying all required information as command line options when you run Database Configuration Assistant. For information about the list of options supported, enter the following command:

```
$ $ORACLE_HOME/bin/dbca -help
```

2. Open the response file in a text editor:

```
$ vi /local_dir/dbca.rsp
```

3. Edit the file, following the instructions in the file.

Note: Database Configuration Assistant fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user, and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.
5. If you intend running Database Configuration Assistant in response file mode, set the `DISPLAY` environment variable.
6. Use the following command syntax to run Database Configuration Assistant in silent or response file mode using a response file:

```
$ORACLE_HOME/bin/dbca {-progressOnly | -silent} -responseFile \  
/local_dir/dbca.rsp
```

In this example:

- The `-silent` option runs Database Configuration Assistant in silent mode.
- The `-progressOnly` option runs Database Configuration Assistant in response file mode.
- `local_dir` is the full path of the directory where you copied the `dbca.rsp` response file template.

Postinstallation Configuration Using a Response File

Use the following sections to create and run a response file configuration after installing Oracle software.

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.

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 B-1 Password response file for Oracle grid infrastructure installation for a cluster

Oracle grid infrastructure requires passwords for 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.

Example B-2 Password response file for Oracle Real Application Clusters

Oracle Database configuration requires the `SYS`, `SYSTEM`, `SYSMAN`, and `DBSNMP` passwords for use with Database Configuration Assistant (DBCA). The `S_ASMSNMPPASSWORD` response is necessary only if the database is using Oracle ASM for storage. Also, if you selected to configure Oracle Enterprise Manager, then you must provide the password for the Oracle software installation owner for the `S_HOSTUSERPASSWORD` response.

```
oracle.assistants.server|S_SYSPASSWORD=password
oracle.assistants.server|S_SYSTEMPASSWORD=password
oracle.assistants.server|S_SYSMANPASSWORD=password
oracle.assistants.server|S_DBSNMPPASSWORD=password
oracle.assistants.server|S_HOSTUSERPASSWORD=password
oracle.assistants.server|S_ASMSNMPPASSWORD=password
```

If you do not want to enable Oracle Enterprise Manager or 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](#)

C.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 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 Technical Network (OTN) and My Oracle Support for further information.

C.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 2MB to 256MB, depending on the kernel version and the hardware architecture. For Oracle Databases, using HugePages reduces the operating system maintenance of page states, and increases TLB (Translation Lookaside Buffer) hit ratio.

Without HugePages, the operating system keeps each 4 KB of memory as a page, and when it is allocated to the SGA, then 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, since 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 $256 \text{ MB} * 1024 \text{ KB} / 4 \text{ KB} = 65536$ PTEs.

C.2 Restrictions for HugePage Configurations

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.

HugePages are 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.

Oracle Grid Infrastructure for a Cluster Installation Concepts

This appendix explains the reasons for preinstallation tasks that you are asked to perform, and other installation concepts.

This appendix contains the following sections:

- [Understanding Preinstallation Configuration](#)
- [Understanding Storage Configuration](#)
- [Understanding Out-of-Place Upgrade](#)

Understanding Preinstallation Configuration

This section reviews concepts about grid infrastructure for a cluster preinstallation tasks. It contains the following sections:

- [Understanding Oracle Groups and Users](#)
- [Understanding the Oracle Base Directory Path](#)
- [Understanding Network Addresses](#)
- [Understanding Network Time Requirements](#)

Understanding Oracle Groups and Users

This section contains the following topics:

- [Understanding the Oracle Inventory Group](#)
- [Understanding the Oracle Inventory Directory](#)

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, this group then must be the primary group for all users granted administrative 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. Refer to "Determining If the Oracle Inventory and Oracle Inventory Group Exists" to identify an existing Oracle Inventory group.

Understanding the Oracle Inventory Directory

The Oracle Inventory directory (`oraInventory`) is the central inventory location for all Oracle software installed on a server.

The first time you install Oracle software on a system, the installer checks to see if you have created an Optimal Flexible Architecture (OFA) compliant path in the format `u[01-09]/app`, such as `/u01/app`, and that the user running the installation has permissions to write to that path. If this is true, then the installer creates the Oracle Inventory directory in the path `/u[01-09]/app/oraInventory`. For example:

```
/u01/app/oraInventory
```

When you provide an Oracle base path when prompted during installation, or you have set the environment variable `$ORACLE_BASE` for the user performing the Oracle grid infrastructure installation, then OUI creates the Oracle Inventory directory in the path `$ORACLE_BASE/.. /oraInventory`. For example, if `$ORACLE_BASE` is set to `/opt/oracle/11`, then the Oracle Inventory directory is created in the path `/opt/oracle/oraInventory`, one directory level above Oracle base.

If you have created neither an OFA-compliant path nor set `$ORACLE_BASE`, then the Oracle Inventory directory is placed in the home directory of the user that is performing the installation. For example:

```
/home/oracle/oraInventory
```

As this placement can cause permission errors during subsequent installations with multiple Oracle software owners, Oracle recommends that you either create an OFA-compliant installation path, or set an `$ORACLE_BASE` environment path.

For new installations, Oracle recommends that you allow OUI to create the Oracle Inventory directory (`oraInventory`). By default, if you create an Oracle path in compliance with OFA structure, such as `/u01/app`, that is owned by an Oracle software owner, then the Oracle Inventory is created in 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. 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.

Understanding the Oracle Base Directory Path

This section contains information about preparing an Oracle base directory.

Overview of the Oracle Base directory

During installation, you are prompted to specify an Oracle base location, which is owned by the user performing the installation. 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.

Understanding Oracle Base and Grid Infrastructure 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 grid infrastructure for a standalone database--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

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 Clusterware to use. Public and virtual IP addresses are configured on public interfaces. Private addresses are configured on private interfaces.

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

- [About the Public IP Address](#)
- [About the Private IP Address](#)
- [About the Virtual IP Address](#)
- [About the Grid Naming Service \(GNS\) Virtual IP Address](#)
- [About the SCAN](#)

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).

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 marked as private as the cluster interconnects. 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, nor marked as a public subnet by `oifcfg`. Oracle does not support changing the interconnect to an interface using a subnet that you have designated as a public subnet.

See Also: *Oracle Clusterware Administration and Deployment Guide* for further information about setting up and using bonded multiple interfaces

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

About the Virtual IP Address

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

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 hostnames 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

About the SCAN

Oracle Database 11g release 2 clients connect to the 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 reflecting multiple listeners 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 contracted on a client's behalf. 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 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 prior to Oracle Database 11g release 2 can continue to use their existing connection addresses; using SCANS is not required. When you upgrade to Oracle Clusterware 11g release 2 (11.2), 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 SCAN is optional for most deployments. However, clients using Oracle Database release 11g release 2 and later policy-managed databases using server pools must access the database using the SCAN. This is required because policy-managed databases can run on different servers at different times, so connecting to a particular node virtual IP address for a policy-managed database is not possible.

Understanding Network Time Requirements

Oracle Clusterware 11g release 2 (11.2) 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 an existing cluster synchronization service, such as NTP, then it will start in an observer mode. Otherwise, it will start in an active mode to ensure that time is synchronized between cluster nodes. CTSS will 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.

Understanding Storage Configuration

[Understanding Automatic Storage Management Cluster File System \(ACFS\)](#)

[About Migrating Existing Oracle ASM Instances](#)

[About Converting Standalone Oracle ASM Installations to Clustered Installations](#)

Understanding Automatic Storage Management Cluster File System (ACFS)

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 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.

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 Automatic Storage Management Configuration Assistant (ASMCA, located in the path `Grid_home/bin`) to upgrade the existing Oracle ASM instance to Oracle ASM 11g release 2 (11.2), and subsequently configure failure groups, ASM volumes and Automatic Storage Management Cluster File System (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 ASM home, then after installing the Oracle ASM 11g release 2 (11.2) binaries, you can start ASMCA to upgrade the existing Oracle ASM instance. You can then choose to configure an ACFS deployment by creating ASM volumes and using the upgraded Oracle ASM to create the ACFS.

On an existing Oracle Clusterware or Oracle RAC installation, if the prior version of Oracle ASM instances on all nodes is Oracle ASM 11g release 1, then you are provided with the option to perform a rolling upgrade of Oracle ASM instances. If the prior version of Oracle ASM instances on an Oracle RAC installation are from an Oracle ASM release prior to Oracle ASM 11g release 1, then rolling upgrades cannot be performed. Oracle ASM is then upgraded on all nodes to 11g release 2 (11.2).

About Converting Standalone Oracle ASM Installations to Clustered Installations

If you have an 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 disks) 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 11g release 2 (11.2) 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, diskgroup names on the cluster-enabled Oracle ASM instances must be different from existing standalone diskgroup names.

See Also: *Oracle Database Storage Administrator's Guide*

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 Clusterware 11g release 2 is not supported.

See Also: [Appendix F, "How to Upgrade to Oracle Grid Infrastructure 11g Release 2"](#) for instructions on completing rolling upgrades

How to Complete Installation Prerequisite Tasks Manually

This appendix provides instructions for how to complete configuration tasks manually that Cluster Verification Utility (CVU) and the installer (OUI) normally complete during installation. Use this appendix as a guide if you cannot use the fixup script.

This appendix contains the following information:

- [Configuring SSH Manually on All Cluster Nodes](#)
- [Configuring Kernel Parameters](#)
- [Checking OCFS2 Version Manually](#)

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](#)

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.

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:

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=502(grid) gid=501(oinstall) groups=501(oinstall),502(grid,asmadmin,asmdba)
$ id grid
uid=502(grid) gid=501(oinstall) groups=501(oinstall),502(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.

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_rsa.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
```

3. 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.

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`:

1. On the system where you want to run OUI, log in as the `grid` user.
2. Use the following command syntax, where `hostname1`, `hostname2`, and so on, are the public hostnames (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 hostname 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 ["Setting Display and X11 Forwarding Configuration"](#) on page 2-51.

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.

Configuring Kernel Parameters

This section contains the following:

- [Minimum Parameter Settings for Installation](#)
- [Additional Parameter and Kernel Settings for SUSE Linux](#)
- [Checking OCFS2 Version Manually](#)

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.

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 packages if necessary to minimum values.

If you cannot use the `fixup` scripts, then review the following table to set values manually:

Parameter	Value	File
<code>semmsl</code>	250	<code>/proc/sys/kernel/sem</code>
<code>semnms</code>	32000	
<code>semopm</code>	100	
<code>semmni</code>	128	

Parameter	Value	File
shmmax	Either 4 GB - 1 byte, or half the size of physical memory (in bytes), whichever is lower. Default: 536870912	/proc/sys/kernel/shmmax
shmmni	4096	/proc/sys/kernel/shmmni
shmall	2097152	/proc/sys/kernel/shmall
file-max	6815744	/proc/sys/fs/file-max
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
aio-max-nr	1048576	/proc/sys/fs/aio-max-nr

Note: If the current value for any parameter is greater than the value listed in this table, then the Fixup scripts do not change the value of that parameter.

Additional Parameter and Kernel Settings for SUSE Linux

On SUSE 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. On SUSE 10 systems only, use a text editor to change the `/etc/boot.sysctl` parameter `RUN_PARALLEL` flag from `yes` to `no`.
3. 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 501:

```
# echo 501 > /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=501
```

Note: Only one group can be defined as the `vm.hugetlb_shm_group`.

4. Repeat steps 1 through 3 on all other nodes in the cluster.

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 Web site:

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

How to Upgrade to Oracle Grid Infrastructure 11g Release 2

This appendix describes how to perform Oracle Clusterware and Oracle Automatic Storage Management 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 Automatic Storage Management 11g release 2 (11.2) 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](#)
- [Restrictions for Clusterware and ASM Upgrades to Grid Infrastructure](#)
- [Verify System Readiness for Upgrades](#)
- [Upgrading an Existing Oracle Clusterware Installation](#)
- [Performing Rolling Upgrades From an Earlier Release](#)
- [Updating DB Control and Grid Control Target Parameters](#)
- [Downgrading Oracle Clusterware After an Upgrade](#)

F.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.

F.2 Restrictions for Clusterware and ASM Upgrades to Grid Infrastructure

Be aware of the following restrictions and changes for upgrades to Oracle grid infrastructure installations, which consists of Oracle Clusterware and Automatic Storage Management (Oracle ASM):

- To upgrade existing Oracle Clusterware installations to Oracle grid infrastructure 11g, your release must be greater than or equal to 10.1.0.3, 10.2.0.3, or 11.1.0.6.
- To upgrade existing Oracle ASM installations to Oracle grid infrastructure 11g release 2 (11.2) in a rolling fashion, your release must be at least 11.1.0.6.

See Also: Oracle Upgrade Companion" Note 785351.1 on My Oracle Support:

<https://metalink.oracle.com>

- Oracle Clusterware and Oracle ASM upgrades are always out-of-place upgrades. With 11g release 2 (11.2), 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 grid infrastructure for a cluster home for Oracle Clusterware and Oracle ASM 11g release 2 (11.2).
- 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`. Starting with Oracle Database 11g, the same user that owned the Oracle Clusterware 10g software must perform the Oracle Clusterware 11g upgrade.
- Oracle ASM and Oracle Clusterware both run in the Oracle grid infrastructure home.
- During a major version upgrade to 11g release 2 (11.2), the software in the 11g release 2 (11.2) grid infrastructure home is not fully functional until the upgrade is completed. Running `srvctl`, `crsctl`, and other commands from the 11g release 2 (11.2) home is not supported until the final `rootupgrade.sh` script is run and the upgrade is complete across all nodes.

To manage databases in the existing earlier version (release 10.x or 11.1) database homes during the grid infrastructure upgrade, use the `srvctl` from the existing database homes.

- During Oracle Clusterware installation, if there is a standalone Oracle ASM version on the local node, then it is converted to a clustered Oracle ASM 11g release 2 (11.2) installation, and Oracle ASM runs in the Oracle grid infrastructure home on all nodes.
- If a standalone (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 standalone Oracle ASM installation. However, during installation, if you select to place the Oracle Cluster Registry (OCR) and voting disk files on Oracle ASM, then a clustered Oracle ASM installation is created on all nodes in the cluster, and the standalone Oracle ASM installation on the remote node will become nonfunctional.

See Also: *Oracle Database Upgrade Guide*

F.3 Verify System Readiness for Upgrades

Use the Cluster Verification Utility to assist you with system checks in preparation for starting a database upgrade.

See Also: *Oracle Database Upgrade Guide*

F.4 Upgrading 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.

F.4.1 Preparing to Upgrade an Existing Oracle Clusterware Installation

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
 - A SCAN address
 - Privileged user operating system groups to grant access to Oracle ASM data files (the OSDBA for ASM group), to grant administrative privileges to the Oracle ASM instance (OSASM group), and to grant a subset of administrative privileges to the Oracle ASM instance (OSOPER for ASM group)
 - `root` user access, to run scripts as `root` during installation
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
```

F.5 Performing Rolling Upgrades From an Earlier Release

Use the following procedures to upgrade Oracle Clusterware or Automatic Storage Management:

- [Verify System Readiness for Upgrades](#)
- [Performing a Rolling Upgrade of Oracle Clusterware](#)
- [Performing a Rolling Upgrade of Automatic Storage Management](#)

Note: When you upgrade to Oracle Clusterware 11g release 2 (11.2), Automatic Storage Management (Oracle ASM) is installed in the same home as Oracle Clusterware. In Oracle documentation, this home is called the "grid infrastructure home," or Grid home. Also note that Oracle does not support attempting to add additional nodes to a cluster during a rolling upgrade.

F.5.1 Verify System Readiness for Upgrades

Use Cluster Verification Utility to assist you with system checks in preparation for starting a database upgrade. The installer runs the appropriate CVU 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.

With Oracle Clusterware 11g release 2 (11.2), you can perform upgrades on a shared Oracle Clusterware home.

F.5.2 Performing a Rolling Upgrade of Oracle Clusterware

Use the following procedure to upgrade Oracle Clusterware from an earlier release to a later release:

Note: Oracle recommends that you leave Oracle RAC instances running. When you start the root script on each node, that node's instances are shut down and then started up again by the `rootupgrade.sh` script.

For standalone Oracle Databases on the cluster, only those that use Oracle ASM need to be shut down. Listeners do not need to be shut down.

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.

Note: In contrast with releases prior to Oracle Clusterware 11g release 2, all upgrades are rolling upgrades, even if you select all nodes for the upgrade.

Oracle recommends that you select all cluster member nodes for the upgrade, and then shut down database instances on each node before you run the upgrade `root` script, starting the database instance up again on each node after the upgrade is complete. You can also use this procedure to upgrade a subset of nodes in the cluster.

3. Select installation options as prompted.

When prompted, run the `rootupgrade.sh` script on each node in the cluster that you want to upgrade. 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 `rootupgrade.sh` script is run on a node, the upgraded Oracle Clusterware stack and AUTOSTART resources are started on the node.

4. Run the `rootupgrade.sh` script on each node on which you are performing the rolling upgrade. Run the script on the local node first. 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.

5. After running the `rootupgrade.sh` script on the last node in the cluster, ASM Configuration Assistant runs automatically, and the Oracle Clusterware upgrade is complete.

If an earlier version of Automatic Storage Management is installed, then the installer starts ASM Configuration Assistant to upgrade Oracle ASM to 11g release 2 (11.2). 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 the Oracle Clusterware binaries. Until ASM is upgraded, Oracle databases that use ASM can't be created. Until ASM is upgraded, the 11g release 2 (11.2) ASM management tools in the Grid home (for example, `srvctl`) will not work.

Note: At the end of the upgrade, if you set the OCR backup location manually to the older release Oracle Clusterware home (CRS home), then you must change the OCR backup location to the Oracle grid infrastructure home (Grid home). If you did not set the OCR backup location manually, then this issue does not concern you.

Because upgrades of Oracle Clusterware are out-of-place upgrades, the previous release Oracle Clusterware home cannot be the location of the OCR backups. Backups in the old Oracle Clusterware home could be deleted.

F.5.3 Performing a Rolling Upgrade of Automatic Storage Management

After you have completed the Oracle Clusterware 11g release 2 (11.2) upgrade, if you did not choose to upgrade Oracle ASM when you upgraded Oracle Clusterware, then you can do it separately using the Automatic Storage Management Configuration Assistant (`asmca`) to perform rolling upgrades.

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` will not work until Oracle ASM is upgraded.

Note: ASMCA performs a rolling upgrade only if the earlier version of Oracle ASM is either 11.1.0.6 or 11.1.0.7. Otherwise, ASMCA performs a normal upgrade, in which ASMCA brings down all Oracle ASM instances on all nodes of the cluster, and then brings them all up in the new Grid home.

F.5.3.1 Preparing to Upgrade Oracle ASM

Note the following if you intend to perform rolling upgrades of Oracle ASM:

- The active version of Oracle Clusterware must be 11g release 2 (11.2). To determine the active version, enter the following command:


```
$ crsctl query crs activeversion
```
- You can upgrade a standalone Oracle ASM installation to a clustered Oracle ASM installation. However, you can only upgrade an existing standalone Oracle ASM installation if you run the installation from the node on which the Oracle ASM installation is installed. You cannot upgrade a single instance Oracle ASM installation on a remote node.

- You must ensure that any rebalance operations on your existing Oracle ASM installation are completed before starting the upgrade process.
- During the upgrade process, you alter the Oracle ASM instances to an upgrade state. 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`)

F.5.3.2 Upgrading Oracle ASM

Complete the following procedure to upgrade Oracle ASM:

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 11g release 2 (11.2) home, start ASMCA. For example:

```
$ cd /u01/11.2/grid/bin
$ ./asmca
```

3. Select **Upgrade**.

ASM Configuration Assistant upgrades Oracle ASM in succession for all nodes in the cluster.

See Also: *Oracle Database Upgrade Guide* and *Oracle Database Storage Administrator's Guide* for additional information about preparing an upgrade plan for Oracle ASM, and for starting, completing, and stopping Oracle ASM upgrades

F.6 Updating DB Control and Grid Control Target Parameters

Because Oracle Clusterware release 2 (11.2) is an out-of-place upgrade of the Oracle Clusterware home in a new location (the 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 dbcontrol or Grid control.

Use the following procedure to resolve this issue:

1. Log in to `dbconsole` or `gridconsole`.
2. Navigate to the **Cluster** tab.
3. Click **Monitoring Configuration**
4. Update the value for **Oracle Home** with the new Grid home path.

F.7 Downgrading Oracle Clusterware After an Upgrade

After a successful or a failed upgrade to Oracle Clusterware 11g release 2 (11.2), you can restore Oracle Clusterware to the previous version.

The restoration procedure in this section restores the Clusterware configuration to the state it was in before the Oracle Clusterware 11g release 2 (11.2) upgrade. Any configuration changes you performed during or after the 11g release 2 (11.2) upgrade are removed and cannot be recovered.

To restore Oracle Clusterware to the previous release:

1. On all remote nodes, use the command syntax `Grid_home/crs/install/rootcrs.pl -downgrade [-force]` to stop the 11g release 2 (11.2) resources, shut down the 11g release 2 (11.2) stack.

Note: This command does not reset the OCR, or delete `ocr.loc`.

For example:

```
# /u01/app/grid/11.2.0/crs/install/rootcrs.pl -downgrade
```

If you want to stop a partial or failed 11g release 2 (11.2) installation and restore the previous release Oracle Clusterware, then use the `-force` flag with this command.

2. After the `rootcrs.pl -downgrade` script has completed on all remote nodes, on the local node use the command syntax `Grid_home/crs/install/rootcrs.pl -downgrade -lastnode -oldcrshome pre11.2_crs_home -version pre11.2_crs_version [-force]`, where `pre11.2_crs_home` is the home of the earlier Oracle Clusterware installation, and `pre11.2_crs_version` is the release number of the earlier Oracle Clusterware installation.

For example:

```
# /u01/app/grid/11.2.0/crs/install/rootcrs.pl -downgrade -lastnode -oldcrshome
/u01/app/crs -version 11.1.0.6.0
```

This script downgrades the OCR, and removes binaries from the Grid home. If you want to stop a partial or failed 11g release 2 (11.2) installation and restore the previous release Oracle Clusterware, then use the `-force` flag with this command.

3. After the local node script completes, you are prompted to run `root.sh` from the earlier release Oracle Clusterware installation home in sequence on each member node of the cluster. 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 older version, and configures the old initialization scripts to run the earlier release Oracle Clusterware stack.

Numerics

- 32-bit and 64-bit
 - software versions in the same cluster not supported, 2-21

A

- accounts configured by, 2-3
- ACFS. *See* Oracle ACFS.
- aio-max-nr, E-6
- AMD 64
 - software requirements for, 2-33
- and asmcmd errors, 2-9
- architecture
 - checking system architecture, 2-23
- Asianux
 - operating system requirements, 2-30, 2-34
- ASM
 - and multiple databases, 2-14
 - and rolling upgrade, F-5
 - ASM library driver (ASMLIB), 4-7
 - block device names, 3-25
 - candidate disks, 1-6, 3-23
 - characteristics of failure groups, 3-19, 3-28
 - checking disk availability, 3-25
 - configuring disks for ASM, 3-23
 - creating the OSDBA for ASM group, 2-16
 - disk groups, 3-17
 - displaying attached disks, 3-25
 - failure groups, 3-17
 - examples, 3-19, 3-28
 - identifying, 3-19, 3-28
 - identifying available disks, 3-25
 - number of instances on each node, 1-5, 3-2
 - OSASM or ASM administrator, 2-14
 - OSDBA for ASM group, 2-14
 - recommendations for disk groups, 3-17
 - required for Standard Edition Oracle RAC, 3-1
 - required for Typical install type, 3-1
 - rolling upgrade of, 4-2
 - space required for Oracle Clusterware files, 3-17
 - space required for preconfigured database, 3-17
 - storage option for data files, 3-2
 - storing Oracle Clusterware files on, 3-4
- ASM disks

- marking, 4-7
- ASM group
 - creating, 2-16
- ASM library driver (oracleasm)
 - installing, 3-23
- ASMLIB, 4-7
- Automatic Storage Management Cluster File System.
 - See* Oracle ACFS.
- Automatic Storage Management. *See* ASM.

B

- Bash shell
 - default user startup file, 2-48
 - setting shell limits, 2-50
- .bash_profile file, 2-48
- binaries
 - relinking, 5-7
- block devices
 - and upgrades, 3-3
 - creating permissions file for Oracle Clusterware files, 3-27
 - creating permissions file for Oracle Database files, 3-31
 - desupport of, 3-31
 - desupported, 3-3
 - device name, 3-25
 - for upgrades only, 3-31
- block devices desupported, xiv
- BMC
 - configuring, 2-44
- BMC interface
 - preinstallation tasks, 2-42
- Bourne shell
 - default user startup file, 2-48
 - setting shell limits on Linux x86, 2-50

C

- C shell
 - default user startup file, 2-48
 - setting shell limits, 2-50
- candidate disks
 - marking, 1-6, 3-23
- central inventory, 2-13
 - about, D-1

- central inventory. *See Also* OINSTALL group, and Oracle Inventory group
- changing host names, 4-3
- checking Linux distribution, 2-39
- chmod command, 3-15
- chown command, 3-15
- clients
 - connecting to SCANs, D-5
- cluster configuration file, 4-7
- cluster file system
 - storage option for data files, 3-2
- cluster file systems
 - grid home not supported on, 3-2
- cluster name
 - requirements for, 4-4
- cluster nodes
 - private node names, 4-4
 - public node names, 4-3
 - specifying uids and gids, 2-17
 - virtual node names, 4-4
- Cluster Time Synchronization Service, 2-40
- Cluster Verification Utility
 - cvuqdisk, 2-41
 - fixup scripts, 2-5
 - user equivalency troubleshooting, A-5
- commands, 2-50
 - asmca, 3-29, 4-6, 5-5, F-5
 - asmcmd, 2-9
 - cat, 2-39
 - chmod, 3-15
 - chown, 3-15
 - crsctl, 4-10, 5-6, F-2, F-5
 - dd, xvi
 - df, 1-2, 2-22
 - env, 2-50
 - fdisk, 3-25, 3-27
 - free, 2-23
 - groupadd, 2-18
 - id, 2-18
 - ipmitool, 2-44
 - limit, 2-51
 - lsdev, 3-25
 - lsmod, 2-43
 - mkdir, 3-15
 - modinfo, 2-39
 - modprobe, 2-43
 - nscd, 2-29
 - ntpd, 2-40
 - partprobe, 3-28, 3-31
 - passwd, 2-19
 - ping, 2-24
 - rootcrs.pl, 5-7
 - rootupgrade.sh, F-2
 - rpm, 2-39, 2-41
 - sqlplus, 2-9
 - srvctl, F-2
 - swap, 2-22
 - swapon, 2-22
 - ulimit, 2-50
 - umask, 2-48

- uname, 2-23, 2-39, 3-23
- unset, F-3
- useradd, 2-10, 2-17, 2-18
- usermod, 2-17
- xhost, 2-6
- xterm, 2-6
- configuring kernel parameters, E-5
- cron jobs, 4-5, A-8
- crs_install.rsp file, B-4
- csh.login file, 2-50
- csh.login.local file, 2-50
- ctsdd, 2-40
- custom database
 - failure groups for ASM, 3-19, 3-28
 - requirements when using ASM, 3-17
- Custom installation type
 - reasons for choosing, 2-13
- cvuqdisk, 2-41

D

- data files
 - creating separate directories for, 3-13, 3-14
 - setting permissions on data file directories, 3-15
 - storage options, 3-2
- data loss
 - minimizing with ASM, 3-19, 3-28
- Database Configuration Assistant
 - running in silent mode, B-7
- database files
 - supported storage options, 3-4
- databases
 - ASM requirements, 3-17
- dba group. *See* OSDBA group
- dbca.rsp file, B-4
- default file mode creation mask
 - setting, 2-48
- default Linux installation
 - recommendation for, 2-2
- desupported
 - block devices, xiv
 - raw devices, xiv
- device names
 - IDE disks, 3-25
 - SCSI disks, 3-25
- /dev/shm
 - and initialization parameters, A-3
- df command, 2-22, 2-49
- Direct NFS
 - disabling, 3-15
 - enabling, 3-12
 - for data files, 3-7
 - minimum write size value for, 3-8
- directory
 - creating separate data file directories, 3-13, 3-14
 - permission for data file directories, 3-15
- disk group
 - ASM, 3-17
 - recommendations for Oracle ASM disk groups, 3-17

- disk groups
 - recommendations for, 3-17
- disk space
 - checking, 2-23
 - requirements for preconfigured database in ASM, 3-17
- disks
 - checking availability for ASM, 3-25
 - displaying attached disks, 3-25
- disks. *See Also* ASM disks
- DISPLAY environment variable
 - setting, 2-49

E

- emulator
 - installing from X emulator, 2-6
- Enterprise Linux
 - operating system requirements, 2-30, 2-34
- enterprise.rsp file, B-4
- env command, 2-50
- environment
 - checking settings, 2-50
 - configuring for oracle user, 2-47
- environment variables
 - DISPLAY, 2-49
 - ORACLE_BASE, 2-48, B-3, D-2
 - ORACLE_HOME, 2-9, 2-48, F-3
 - ORACLE_SID, 2-48, F-3
 - removing from shell startup file, 2-48
 - SHELL, 2-48
 - TEMP and TMPDIR, 2-22, 2-49
- errata
 - Linux kernel errata, 2-39
- errors
 - X11 forwarding, 2-51, E-5
 - /etc/csh.login file, 2-50
 - /etc/csh.login.local file, 2-50
 - /etc/pam.d/login file, 2-50
 - /etc/profile file, 2-50
 - /etc/profile.local file, 2-50
 - /etc/security/limits.so file, 2-50
- Exadata
 - relinking binaries example for, 5-7
- examples
 - ASM failure groups, 3-19, 3-28

F

- failure group
 - ASM, 3-17
 - characteristics of ASM failure group, 3-19, 3-28
 - examples of ASM failure groups, 3-19, 3-28
- fdisk command, 3-25
- features, new, xviii
- fencing, xvii
 - and IPMI, 2-42, 4-5
- file mode creation mask
 - setting, 2-48
- file system

- storage option for data files, 3-2
- file systems, 3-4
- file-max, E-6
- files
 - \$ORACLE_HOME/lib/libnfsodm11.so, 3-13
 - \$ORACLE_HOME/lib/libodm11.so, 3-13
 - .bash_profile, 2-48
 - dbca.rsp, B-4
 - default shell startup file, 2-50
 - editing shell startup file, 2-48
 - enterprise.rsp, B-4
 - /etc/csh.login, 2-50
 - /etc/csh.login.local, 2-50
 - /etc/pam.d/login, 2-50
 - /etc/profile, 2-50
 - /etc/profile.local, 2-50
 - /etc/security/limits.so, 2-50
 - .login, 2-48
 - oraInst.loc, 2-8
 - oraInst.loc file, B-3
 - /proc/sys/fs/file-max, E-5
 - /proc/sys/kernel/sem, E-5
 - /proc/sys/kernel/shmall, E-5
 - /proc/sys/kernel/shmmax, E-5
 - /proc/sys/kernel/shmmni, E-5
 - /proc/sys/net/ipv4/ip_local_port_range, E-5
 - .profile, 2-48
 - profile.local, 2-50
 - response files, B-3
- filesets, 2-29
- fixup script, 2-5
 - about, 1-1

G

- GFS, 3-4
- gid
 - identifying existing, 2-18
 - specifying, 2-18
 - specifying on other nodes, 2-17
- globalization
 - support for, 4-3
- GNS
 - about, 2-25
- GPFS, 3-4
- grid home
 - and Oracle base restriction, 2-9
 - default path for, 2-52
 - disk space for, 2-22
 - unlocking, 5-7
- grid infrastructure owner (grid), 2-13
- grid naming service. *See* GNS
- grid user, 2-13
- group IDs
 - identifying existing, 2-18
 - specifying, 2-18
 - specifying on other nodes, 2-17
- groups
 - checking for existing OINSTALL group, 2-7
 - creating identical groups on other nodes, 2-17

- creating the ASM group, 2-16
- creating the OSDBA for ASM group, 2-16
- creating the OSDBA group, 2-15
- OINSTALL, 2-7, 2-8
- OSASM (asmadmin), 2-14
- OSDBA (dba), 2-13
- OSDBA for ASM (asmdba), 2-14
- OSDBA group (dba), 2-13
- OSOPER (oper), 2-13
- OSOPER for ASM, 2-14
- OSOPER group (oper), 2-13
- required for installation owner user, 2-13
- specifying when creating users, 2-18
- using NIS, 2-12, 2-17

H

- hangcheck-timer, xvii
- hardware requirements, 2-21
- host names
 - changing, 4-3
 - legal hostnames, 4-4

I

- id command, 2-18
- IDE disks
 - device names, 3-25
- initialization parameters
 - MEMORY_TARGET, A-3
- INS-32026 error, 2-9
- installation
 - and cron jobs, 4-5
 - and globalization, 4-3
 - response files, B-3
 - oraInst.loc file, B-3
 - preparing, B-3, B-5
 - templates, B-3
 - silent mode, B-6
 - using cluster configuration file, 4-7
- installation types
 - and ASM, 3-17
- interfaces
 - requirements for private interconnect, D-4
- intermittent hangs
 - and socket files, 4-10
- ip_local_port_range, E-6
- IPMI
 - addresses not configurable by GNS, 2-42
 - configuring driver for, 2-42
 - preinstallation tasks, 2-42
 - preparing for installation, 4-5

J

- JDK requirements, 2-29
- job role separation users, 2-13

K

- kernel

- Linux errata, 2-39
- kernel parameters
 - configuring, E-5
- Korn shell
 - and LDAP, 2-32, 2-37
 - default user startup file, 2-48
 - setting shell limits, 2-50
- ksh
 - See Korn shell

L

- LDAP
 - odisrvreg
 - system requirements for, 2-32, 2-37
 - oidca
 - system requirements for, 2-32, 2-37
 - schemasynch
 - system requirements for, 2-32, 2-37
- legal hostnames, 4-4
- libnfsodm11.so, 3-13
- libodm11.so, 3-13
- lightweight directory access protocol
 - See LDAP
- limit command, 2-51
- limits.so file, 2-50
- Linux
 - checking version, 2-39
 - cvuqdisk package, 2-41
 - determining distribution of, 2-39
 - kernel errata, 2-39
 - setting shell limits, 2-50
- Linux (32-bit)
 - software requirements for, 2-30
- Linux x86-64
 - software requirements for, 2-33
- log file
 - how to access during installation, 4-6
- .login file, 2-48
- login file, 2-50
- lsdev command, 3-25
- LVM
 - recommendations for ASM, 3-17

M

- mask
 - setting default file mode creation mask, 2-48
- memory requirements, 2-21
- Memory Size (SGA and PGA), A-3
- MEMORY_TARGET, A-3
 - and /dev/shm, A-3
- minimal Linux installation
 - recommendation for, 2-2
- mixed binaries, 2-30
- mkdir command, 3-15
- mode
 - setting default file mode creation mask, 2-48
- multiple databases
 - and ASM, 2-14

multiple oracle homes, 2-9, 3-15
My Oracle Support, 5-1

N

Net Configuration Assistant (NetCA)
 response files, B-7
 running at command prompt, B-7
netca, 4-6
netca.rsp file, B-4
Network Information Services
 See NIS
new features, xiii
NFS, 3-4, 3-10
 and data files, 3-9
 and Oracle Clusterware files, 3-5
 buffer size parameters for, 3-9, 3-10
 Direct NFS, 3-7
 for data files, 3-9
 rsize, 3-10
NIS
 alternative to local users and groups, 2-12
nofile
 shell limit on Linux, 2-50
noninteractive mode. *See* response file mode
nproc
 shell limit on Linux, 2-50
NTP protocol
 and slewing, 2-40

O

OCFS2, 3-4
 checking on Linux, 2-39
 checking version of, 2-33, 2-38, E-7
 download location for Linux, 2-39
 grid home not supported on, 3-2
 status of, 2-39
OCR. *See* Oracle Cluster Registry
OINSTALL group
 about, D-1
 and oraInst.loc, 2-7
 checking for existing, 2-7
 creating on other nodes, 2-17
OINSTALL group. *See Also* Oracle Inventory group
oper group. *See* OSOPER group
operating system
 checking distribution and version of Linux, 2-39
 different on cluster members, 2-30
 limitation for Oracle ACFS, D-6
 missing packages, A-7
 requirements, 2-29
oprocd, xvii
optimal flexible architecture
 and oraInventory directory, D-2
Oracle ACFS
 about, D-6
 supported Linux versions for, 3-2
Oracle base
 grid homes not permitted under, 2-53

Oracle base directory
 about, D-3
 grid home must not be in an Oracle Database
 Oracle base, 2-9
 minimum disk space for, 2-22
Oracle Cluster Registry
 configuration of, 4-5
 mirroring, 3-5
 partition sizes, 3-6
 permissions file to own block device
 partitions, 3-27
 supported storage options, 3-4
Oracle Clusterware
 and file systems, 3-4
 and upgrading Oracle ASM instances, 1-5, 3-2
 installing, 4-1
 rolling upgrade of, 4-2
 supported storage options for, 3-4
 upgrading, 3-6
Oracle Clusterware files
 ASM disk space requirements, 3-17
Oracle Database
 creating data file directories, 3-13, 3-14
 data file storage options, 3-2
 operating system patch requirements for, 2-31,
 2-35
 privileged groups, 2-13
 requirements with ASM, 3-17
Oracle Database Configuration Assistant
 response file, B-4
Oracle Disk Manager
 and Direct NFS, 3-12
Oracle Enterprise Linux
 and Oracle Validated RPM, 2-3
Oracle grid infrastructure response file, B-4
oracle home, 2-9
 ASCII path restriction for, 4-5
 multiple oracle homes, 2-9, 3-15
Oracle Inventory group
 about, D-1
 checking for existing, 2-7
 creating, 2-8
 creating on other nodes, 2-17
 oraInst.loc file and, 2-8
Oracle Net Configuration Assistant
 response file, B-4
Oracle patch updates, 5-1
Oracle Real Application Clusters
 configuring disks for ASM, 3-25
 operating system patch requirements for, 2-31,
 2-35
Oracle Software Owner user
 configuring environment for, 2-47
 creating, 2-8, 2-9, 2-16
 creating on other nodes, 2-17
 determining default shell, 2-48
 required group membership, 2-13
Oracle software owner user
 description, 2-13
Oracle Software Owner users

- setting shell limits for, 2-50
- Oracle Universal Installer
 - response files
 - list of, B-4
- Oracle Upgrade Companion, 2-4
- oracle user
 - and Oracle Validated RPM, 2-3
 - configuring environment for, 2-47
 - creating, 2-8, 2-9, 2-10, 2-16, 2-17
 - creating on other nodes, 2-17
 - description, 2-13
 - determining default shell, 2-48
 - required group membership, 2-13
 - setting shell limits for, 2-50
- Oracle Validated Configuration RPM, xvii
 - about, 2-3
 - included with Oracle Enterprise Linux, 2-3
 - installing, 2-4
 - troubleshooting, 2-3
- ORACLE_BASE environment variable
 - removing from shell startup file, 2-48
- ORACLE_HOME environment variable
 - removing from shell startup file, 2-48
- ORACLE_SID environment variable
 - removing from shell startup file, 2-48
- oracleasm RPM
 - installing, 3-23
- oracle-validated.rpm, 2-3
- oraInst.loc
 - and central inventory, 2-7
 - contents of, 2-7
- oraInst.loc file
 - location, 2-8
 - location of, 2-8
- oraInventory, 2-13
 - about, D-1
 - creating, 2-8
- oraInventory. *See Also* Oracle Inventory group
- OS Watcher, 5-3
- OSASM group, 2-14
 - about, 2-14
 - and multiple databases, 2-14
 - and SYSASM, 2-14
 - creating, 2-16
- OSDBA for ASM group, 2-14
 - about, 2-14
- OSDBA group
 - and SYSDBA privilege, 2-13
 - creating, 2-15
 - creating on other nodes, 2-17
 - description, 2-13
- OSDBA group for ASM
 - creating, 2-16
- OSOPER for ASM group
 - about, 2-14
 - creating, 2-16
- OSOPER group
 - and SYSOPER privilege, 2-13
 - creating, 2-15
 - creating on other nodes, 2-17

- description, 2-13

P

- package cvuqdisk not installed, 2-41
- packages
 - checking on Linux, 2-39
 - default Linux installation and, 2-30
- packages. *See Also* RPMs
- partition
 - using with ASM, 3-17
- passwd command, 2-19
- passwords
 - specifying for response files, B-2
 - See also* security
- patch updates
 - download, 5-1
 - install, 5-1
 - My Oracle Support, 5-1
- PC X server
 - installing from, 2-6
- permissions
 - for data file directories, 3-15
- physical RAM requirements, 2-21
- policy-managed databases
 - and SCANS, D-5
- postinstallation
 - patch download and install, 5-1
 - root.sh back up, 5-2
- preconfigured database
 - ASM disk space requirements, 3-17
 - requirements when using ASM, 3-17
- privileged groups
 - for Oracle Database, 2-13
- processor
 - checking system architecture, 2-23
 - /proc/sys/fs/aio-max-nr, E-6
 - /proc/sys/fs/file-max file, E-6
 - /proc/sys/kernel/sem file, E-5
 - /proc/sys/kernel/shmall, E-6
 - /proc/sys/kernel/shmall file, E-6
 - /proc/sys/kernel/shmmax, E-6
 - /proc/sys/kernel/shmni file, E-6
 - /proc/sys/kernel/shmni, E-6
 - /proc/sys/net/core/rmem_default file, E-6
 - /proc/sys/net/core/rmem_max file, E-6
 - /proc/sys/net/core/wmem_default file, E-6
 - /proc/sys/net/core/wmem_max file, E-6
 - /proc/sys/net/ipv4/ip_local_port_range, E-6
- .profile file, 2-48
- profile file, 2-50
- profile.local file, 2-50
- PRVF-5436 error, 2-40

R

- RACDDT, 5-3
- RAID
 - and mirroring Oracle Cluster Registry and voting disk, 3-5

- recommended ASM redundancy level, 3-17
- RAM requirements, 2-21
- raw devices
 - and upgrades, 3-3, 3-31
 - desupport of, 3-31
 - upgrading existing partitions, 3-6
- raw devices desupported, xiv
- recovery files
 - supported storage options, 3-4
- Red Hat
 - operating system requirements, 2-30, 2-34
- Red Hat Package Manager
 - See* RPM
- redundancy level
 - and space requirements for preconfigured database, 3-17
- relinking grid infrastructure home binaries, 5-7, 6-3
- requirements, 3-17
 - hardware, 2-21
- response file installation
 - oraInst.loc file, B-3
 - preparing, B-3
- response files
 - templates, B-3
 - silent mode, B-6
- response file mode
 - about, B-1
 - reasons for using, B-2
 - See also* response files, silent mode, B-1
- response files
 - about, B-1
 - creating with template, B-3
 - crs_install.rsp, B-4
 - dbca.rsp, B-4
 - enterprise.rsp, B-4
 - general procedure, B-2
 - Net Configuration Assistant, B-7
 - netca.rsp, B-4
 - passing values at command line, B-2
 - passwords, B-2
 - security, B-2
 - specifying with Oracle Universal Installer, B-6
- response files. *See also* silent mode
- rmem_default, E-6
- rmem_max, E-6
- rolling upgrade
 - ASM, 4-2
 - of ASM, F-5
 - Oracle Clusterware, 4-2
- root user
 - logging in as, 2-6
- root.sh, 4-6
 - back up, 5-2
 - running, 4-3
- RPMs
 - checking, 2-39
 - default Linux installation and, 2-30
- rsize parameter, 3-10
- run level, 2-21

S

- SCAN listener, D-5
- SCANS, 1-3, 2-25
 - understanding, D-5
 - use of SCANS required for clients of policy-managed databases, D-5
- scripts
 - root.sh, 4-3
- SCSI disks
 - device names, 3-25
- security
 - dividing ownership of Oracle software, 2-12
 - See also* passwords
- sem file, E-5
- semnmi parameter
 - recommended value on Linux, E-5
- semnms parameter
 - recommended value on Linux, E-5
- semmsl parameter
 - recommended value on Linux, E-5
- semopm parameter
 - recommended value on Linux, E-5
- setting shell limits, 2-50
- shell
 - default shell startup file, 2-50
 - determining default shell for oracle user, 2-48
- SHELL environment variable
 - checking value of, 2-48
- shell limits
 - setting on Linux, 2-50
- shell startup file
 - editing, 2-48
 - removing environment variables, 2-48
- shmall, E-6
- shmmax, E-6
- shmmni, E-6
- silent mode
 - about, B-1
 - reasons for using, B-2
 - See also* response files., B-1
- silent mode installation, B-6
- single client access names. *See* SCAN addresses
- software requirements, 2-29
 - checking software requirements, 2-39
- ssh
 - and X11 Forwarding, 2-51
 - automatic configuration from OUI, 2-47
 - configuring, E-1
 - supported version of, E-1
 - when used, 2-47
- stack
 - shell limit on Linux, 2-50
- startup file
 - default shell startup file, 2-50
 - for shell, 2-48
- storage
 - marking ASM candidate disks, 4-7
- supported storage options
 - Oracle Clusterware, 3-4
- suppressed mode

- reasons for using, B-2
- SUSE
 - operating system requirements, 2-30, 2-34
- swap space
 - requirements, 2-21
- SYSASM, 2-14
 - and OSASM, 2-14
- SYSDBA
 - using database SYSDBA on ASM
 - deprecated, 2-14
- SYSDBA privilege
 - associated group, 2-13
- SYSOPER privilege
 - associated group, 2-13
- system architecture
 - checking, 2-23

T

- tcsh shell
 - setting shell limits, 2-50
- TEMP environment variable, 2-22
 - setting, 2-49
- temporary directory, 2-22
- temporary directory. *See* /tmp directory
- temporary disk space
 - checking, 2-22
 - freeing, 2-22
 - requirements, 2-21
- /tmp directory
 - checking space in, 2-22
 - freeing space in, 2-22
- TMPDIR environment variable, 2-22
 - setting, 2-49
- troubleshooting
 - asmcmd errors and oracle home, 2-9
 - automatic SSH configuration from OUI, 2-47
 - different group and user IDs, 2-3
 - disk space errors, 4-5
 - DISPLAY errors, 2-51
 - environment path errors, 4-5
 - error messages, A-1
 - intermittent hangs, 4-10
 - log file, 4-6
 - missing operating system packages, A-7
 - nfs mounts, 2-29
 - permissions errors and oraInventory, D-1
 - permissions errors during installation, D-2
 - public network failures, 2-29
 - run level error, 2-21
 - sqlplus errors and oracle home, 2-9
 - ssh, E-1
 - ssh configuration failure, E-2
 - ssh errors, 2-51
 - stty errors, 2-51
 - unexplained installation errors, 4-5, A-8
 - user equivalency, A-5, E-1
 - user equivalency error due to different user or group IDs, 2-11, 2-17
 - user equivalency errors, 2-8
 - voting disk backups with dd command, xvi
 - with OS Watcher and RACDDT, 5-3
 - X11 forwarding error, 2-51
- Typical installation type
 - response file installations, B-5

U

- uid
 - identifying existing, 2-18
 - specifying, 2-18
 - specifying on other nodes, 2-17
- ulimit command, 2-50
- umask, 2-50
- umask command, 2-48, 2-50
- upgrade
 - of Oracle Clusterware, 4-2
 - restrictions for, F-1
 - unseting environment variables for, F-3
- upgrades, 2-4
 - and SCANs, D-5
 - of Oracle ASM, F-5
 - using raw or block devices with, 3-3
- upgrading
 - and existing Oracle ASM instances, 1-5, 3-2
 - and OCR partition sizes, 3-6
 - and voting disk partition sizes, 3-6
 - shared Oracle Clusterware home to local grid homes, 2-53
- user equivalence
 - testing, A-5
- user equivalency errors
 - groups and users, 2-11, 2-17
- user IDs
 - identifying existing, 2-18
 - specifying, 2-18
 - specifying on other nodes, 2-17
- useradd command, 2-10, 2-17, 2-18
- users
 - creating identical users on other nodes, 2-17
 - creating the oracle user, 2-8, 2-9, 2-16
 - oracle software owner user (oracle), 2-13
 - setting shell limits for, 2-50
 - setting shell limits for users on Linux, 2-50
 - specifying groups when creating, 2-18
 - using NIS, 2-12, 2-17

V

- voting disks
 - backing up with dd command deprecated, xvi
 - configuration of, 4-4
 - mirroring, 3-5
 - partition sizes, 3-6
 - supported storage options, 3-4

W

- wmem_default, E-6
- wmem_max, E-6
- workstation

- installing from, 2-6
- wsize parameter, 3-10
- wtmax, 3-8
 - minimum value for Direct NFS, 3-8

X

- X emulator
 - installing from, 2-6
- X window system
 - enabling remote hosts, 2-6
- X11 forwarding
 - error, 2-51
- X11 forwarding errors, E-5
- xhost command, 2-6
- xterm command, 2-6

