

Oracle® Grid Infrastructure

Installation Guide

11g Release 2 (11.2) for IBM AIX on POWER Systems (64-Bit)

E17210-03

October 2010

Oracle Grid Infrastructure Installation Guide, 11g Release 2 (11.2) for IBM AIX on POWER Systems (64-Bit)
E17210-03

Copyright © 2007, 2010, Oracle and/or its affiliates. All rights reserved.

Primary Author: Douglas Williams

Contributing Authors: Mark Bauer, Jonathan Creighton, Prakash Jashnani, Barb Lundhild, Saar Maoz, Sundar Matpadi, Markus Michalewicz, Hanlin Qian, Dipak Saggi, Ara Shakian

Contributors: Alex Huang, Rache Wang

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	xi
Intended Audience.....	xi
Documentation Accessibility	xi
Related Documents	xii
Conventions	xiii
What's New in Oracle Grid Infrastructure Installation and Configuration?	xv
Desupported Options	xv
New Features for Release 2 (11.2.0.2)	xv
New Features for Release 2 (11.2)	xvii
New Features for Release 1 (11.1)	xix
1 Typical Installation for Oracle Grid Infrastructure for a Cluster	
Typical and Advanced Installation	1-1
Preinstallation Steps Completed Using Typical Installation	1-1
Preinstallation Steps Requiring Manual Tasks	1-2
Verify System Requirements	1-2
Check Network Requirements	1-3
Single Client Access Name (SCAN) for the Cluster.....	1-3
IP Address Requirements	1-4
Redundant Interconnect Usage.....	1-4
Intended Use of Network Interfaces	1-4
Check Operating System Packages.....	1-5
Create Groups and Users	1-5
Configure Oracle Installation Owner Shell Limits.....	1-5
Check Storage	1-6
Prepare Storage for Oracle Automatic Storage Management	1-6
Install Oracle Grid Infrastructure Software.....	1-6
2 Advanced Installation Oracle Grid Infrastructure for a Cluster Preinstallation Tasks	
Reviewing Upgrade Best Practices.....	2-1
Installation Fixup Scripts.....	2-2
Logging In to a Remote System as root Using X Terminal.....	2-3
Creating Groups, Users and Paths for Oracle Grid Infrastructure	2-4

Determining If the Oracle Inventory and Oracle Inventory Group Exists	2-5
Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist.....	2-5
Creating the Oracle Grid Infrastructure User	2-6
Understanding Restrictions for Oracle Software Installation Owners	2-6
Determining if an Oracle Software Owner User Exists.....	2-7
Creating or Modifying an Oracle Software Owner User for Oracle Grid Infrastructure.	2-7
Creating the Oracle Base Directory Path	2-9
Creating Job Role Separation Operating System Privileges Groups and Users	2-9
Overview of Creating Job Role Separation Groups and Users	2-10
Creating Database Groups and Users with Job Role Separation	2-12
Example of Creating Standard Groups, Users, and Paths	2-17
Example of Creating Role-allocated Groups, Users, and Paths	2-18
Checking the Hardware Requirements.....	2-19
Checking the Network Requirements.....	2-21
Network Hardware Requirements	2-21
IP Address Requirements	2-23
IP Address Requirements with Grid Naming Service	2-24
IP Address Requirements for Manual Configuration	2-24
Broadcast Requirements for Networks Used by Oracle Grid Infrastructure.....	2-25
Multicast Requirements for Networks Used by Oracle Grid Infrastructure.....	2-25
DNS Configuration for Domain Delegation to Grid Naming Service	2-25
Grid Naming Service Configuration Example.....	2-26
Manual IP Address Configuration Example.....	2-27
Network Interface Configuration Options	2-28
Identifying the Software Requirements	2-29
Checking the Software Requirements	2-33
Verifying UDP and TCP Kernel Parameters	2-34
Enabling Core File Creation.....	2-34
Tuning AIX System Environment	2-35
Checking Asynchronous Input Output Processes for AIX 5L.....	2-36
Tuning Virtual Memory Manager (VMM).....	2-37
Increasing System Block Size Allocation	2-37
Configuring Shell Limits	2-37
Configuring User Process Parameters	2-38
Configuring Network Tuning Parameters	2-38
Network Time Protocol Setting	2-40
Automatic SSH Configuration During Installation	2-41
Configuring Grid Infrastructure Software Owner User Environments.....	2-42
Environment Requirements for Oracle Grid Infrastructure Software Owner	2-42
Environment Requirements for Oracle Database and Oracle ASM Owners	2-42
Procedure for Configuring Oracle Software Owner Environments.....	2-43
Setting Display and X11 Forwarding Configuration	2-44
Preventing Installation Errors Caused by stty Commands	2-45
Running the rootpre.sh Script	2-45
Requirements for Creating an Oracle Grid Infrastructure Home Directory	2-46

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

Reviewing Oracle Grid Infrastructure Storage Options	3-1
Overview of Oracle Clusterware and Oracle RAC Storage Options.....	3-1
Quorum Disk Location Restriction with Existing 9.2 Clusterware Installations.....	3-2
After You Have Selected Disk Storage Options	3-2
General Information About Oracle ACFS	3-3
General Storage Considerations for Oracle Grid Infrastructure and Oracle RAC	3-3
General Storage Considerations for Oracle Clusterware.....	3-3
General Storage Considerations for Oracle RAC	3-3
Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC	3-4
Supported Storage Options	3-4
After You Have Selected Disk Storage Options	3-5
Shared File System Storage Configuration	3-6
Requirements for Using a Shared File System.....	3-6
Deciding to Use a Cluster File System for Oracle Clusterware Files.....	3-8
Deciding to Use Direct NFS for Data Files	3-8
About Direct NFS Storage	3-8
Using the Oranfstab File with Direct NFS.....	3-8
Mounting NFS Storage Devices with Direct NFS	3-9
Specifying Network Paths with the Oranfstab File	3-10
Deciding to Use NFS for Data Files	3-10
Configuring Storage NFS Mount and Buffer Size Parameters.....	3-10
Enabling Direct NFS Client Oracle Disk Manager Control of NFS	3-12
Configuring HACMP Multinode Disk Heartbeat (MNDHB) for Oracle Clusterware	3-13
Overview of Requirements for Using HACMP with Oracle Clusterware	3-14
Deploying HACMP and MDNDHB for Oracle Clusterware.....	3-14
Upgrading an Existing Oracle Clusterware and HACMP Installation.....	3-16
Configuring Raw Logical Volumes for Oracle Clusterware.....	3-17
Configuring Raw Logical Volumes in the New Oracle Clusterware Volume Group	3-18
Creating a Volume Group for Database Files	3-19
Creating a Volume Group for Oracle Clusterware	3-21
Importing the Volume Group on the Other Cluster Nodes.....	3-23
Activating the Volume Group in Concurrent Mode on All Cluster Nodes	3-24
Creating Directories for Oracle Clusterware Files on Shared File Systems.....	3-24
Creating Directories for Oracle Database Files on Shared File Systems	3-26
Disabling Direct NFS Client Oracle Disk Management Control of NFS.....	3-27
Oracle Automatic Storage Management Storage Configuration	3-27
Configuring Storage for Oracle Automatic Storage Management	3-28
Identifying Storage Requirements for Oracle ASM.....	3-28
Creating Files on a NAS Device for Use with Oracle ASM	3-31
Using an Existing Oracle ASM Disk Group	3-32
Configuring Disk Devices for Oracle ASM	3-34
Using Disk Groups with Oracle Database Files on Oracle ASM	3-36
Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM.....	3-36
Creating Disk Groups for Oracle Database Data Files	3-36
Configuring Oracle Automatic Storage Management Cluster File System.....	3-36

Upgrading Existing Oracle ASM Instances	3-38
Desupport of Raw Disks	3-38

4 Installing Oracle Grid Infrastructure for a Cluster

Preparing to Install Oracle Grid Infrastructure with OUI	4-1
Installing Grid Infrastructure	4-6
Running OUI to Install Grid Infrastructure	4-6
Installing Grid Infrastructure Using a Cluster Configuration File	4-7
Installing Grid Infrastructure Using a Software-Only Installation	4-7
Installing the Software Binaries	4-8
Configuring the Software Binaries	4-8
Configuring the Software Binaries Using a Response File.....	4-9
Restrictions for Oracle Berkeley DB.....	4-9
Confirming Oracle Clusterware Function	4-10
Confirming Oracle ASM Function for Oracle Clusterware Files	4-10
Confirming Oracle Clusterware Function	4-10

5 Oracle Grid Infrastructure Postinstallation Procedures

Required Postinstallation Tasks	5-1
Download and Install Patch Updates	5-1
Recommended Postinstallation Tasks.....	5-2
Back Up the root.sh Script.....	5-2
Tune Semaphore Parameters.....	5-3
Create a Fast Recovery Area Disk Group	5-3
About the Fast Recovery Area and the Fast Recovery Area Disk Group.....	5-3
Creating the Fast Recovery Area Disk Group	5-4
Using Older Oracle Database Versions with Grid Infrastructure	5-4
General Restrictions for Using Older Oracle Database Versions.....	5-4
Using ASMCA to Administer Disk Groups for Older Database Versions.....	5-5
Pinning Cluster Nodes for Oracle Database Release 10.x or 11.x	5-5
Enabling The Global Services Daemon (GSD) for Oracle Database Release 9.2.....	5-6
Using the Correct LSNRCTL Commands.....	5-6
Modifying Oracle Clusterware Binaries After Installation	5-6

6 How to Modify or Deinstall Oracle Grid Infrastructure

Deciding When to Deinstall Oracle Clusterware	6-1
Migrating From Oracle Restart to Oracle Clusterware	6-1
Relinking Oracle Grid Infrastructure for a Cluster Binaries.....	6-2
Changing the Oracle Grid Infrastructure Home Path.....	6-3
Deconfiguring Oracle Clusterware Without Removing Binaries.....	6-4
Removing Oracle Clusterware and Oracle ASM	6-5
About the Deinstallation Tool	6-5
Downloading The Deinstall Tool for Use with Failed Installations	6-7
Deinstall Command Example for Oracle Clusterware and Oracle ASM.....	6-7
Deinstallation Parameter File Example for Grid Infrastructure for a Cluster.....	6-7

A	Troubleshooting the Oracle Clusterware Installation Process	
	General Installation Issues.....	A-1
	Interpreting CVU "Unknown" Output Messages Using Verbose Mode.....	A-3
	Interpreting CVU Messages About Oracle Grid Infrastructure Setup.....	A-4
	About the Oracle Clusterware Alert Log.....	A-6
	Troubleshooting Issues on AIX.....	A-6
	Performing Cluster Diagnostics During Oracle Clusterware Installations.....	A-7
	Interconnect Configuration Issues.....	A-7
	Storage Configuration Issues.....	A-7
	Recovery from Losing a Node Filesystem or Grid Home.....	A-8
B	Installing and Configuring Oracle Database Using Response Files	
	How Response Files Work.....	B-1
	Reasons for Using Silent Mode or Response File Mode.....	B-2
	General Procedure for Using Response Files.....	B-2
	Preparing a Response File.....	B-3
	Editing a Response File Template.....	B-3
	Recording a Response File.....	B-4
	Running the Installer Using a Response File.....	B-5
	Running Net Configuration Assistant Using a Response File.....	B-6
	Postinstallation Configuration Using a Response File.....	B-7
	About the Postinstallation Configuration File.....	B-7
	Running Postinstallation Configuration Using a Response File.....	B-7
C	Oracle Grid Infrastructure for a Cluster Installation Concepts	
	Understanding Preinstallation Configuration.....	C-1
	Understanding Oracle Groups and Users.....	C-1
	Understanding the Oracle Inventory Group.....	C-1
	Understanding the Oracle Inventory Directory.....	C-2
	Understanding the Oracle Base Directory Path.....	C-3
	Overview of the Oracle Base directory.....	C-3
	Understanding Oracle Base and Grid Infrastructure Directories.....	C-3
	Understanding Network Addresses.....	C-3
	About the Public IP Address.....	C-3
	About the Private IP Address.....	C-4
	About the Virtual IP Address.....	C-4
	About the Grid Naming Service (GNS) Virtual IP Address.....	C-4
	About the SCAN.....	C-4
	Understanding Network Time Requirements.....	C-5
	Understanding Storage Configuration.....	C-6
	Understanding Oracle Automatic Storage Management Cluster File System.....	C-6
	About Migrating Existing Oracle ASM Instances.....	C-6
	About Converting Standalone Oracle ASM Installations to Clustered Installations.....	C-7
	Understanding Server Pools.....	C-7
	Overview of Server Pools and Policy-based Management.....	C-7
	How Server Pools Work.....	C-8

The Free Server Pool.....	C-8
The Generic Server Pool.....	C-8
Understanding Out-of-Place Upgrade	C-9

D How to Complete Installation Prerequisite Tasks Manually

Configuring SSH Manually on All Cluster Nodes.....	D-1
Checking Existing SSH Configuration on the System	D-1
Configuring SSH on Cluster Nodes	D-2
Create SSH Directory, and Create SSH Keys On Each Node	D-2
Add All Keys to a Common authorized_keys File	D-3
Enabling SSH User Equivalency on Cluster Nodes	D-4

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

Back Up the Oracle Software Before Upgrades.....	E-1
Unset Oracle Environment Variables	E-1
About Oracle ASM and Oracle Grid Infrastructure Installation and Upgrade	E-2
Restrictions for Clusterware and Oracle ASM Upgrades.....	E-2
Preparing to Upgrade an Existing Oracle Clusterware Installation.....	E-4
Verifying System Readiness for Upgrades.....	E-4
Checks to Complete Before Upgrade an Existing Oracle Clusterware Installation.....	E-4
Performing Rolling Upgrades From an Earlier Release	E-5
Performing a Rolling Upgrade of Oracle Clusterware	E-5
Performing a Rolling Upgrade of Oracle Automatic Storage Management	E-6
Preparing to Upgrade Oracle ASM	E-7
Upgrading Oracle ASM	E-7
Updating DB Control and Grid Control Target Parameters	E-8
Unlocking the Existing Oracle Clusterware Installation.....	E-8
Downgrading Oracle Clusterware After an Upgrade	E-8
Checking Cluster Health Monitor Repository Size After Upgrading	E-9

Index

List of Tables

2-1	Grid Naming Service Example Network.....	2-27
2-2	Manual Network Configuration Example	2-28
2-3	AIX Operating System Kernel Requirements	2-29
2-4	AIX APAR and Other Operating System Fixes	2-32
2-5	Recommended Values for Virtual Memory Manager	2-37
3-1	Supported Storage Options for Oracle Clusterware and Oracle RAC.....	3-5
3-2	Oracle Clusterware Shared File System Volume Size Requirements.....	3-7
3-3	Oracle RAC Shared File System Volume Size Requirements.....	3-7
3-4	Total Oracle Clusterware Storage Space Required by Redundancy Type	3-29
3-5	Total Oracle Database Storage Space Required by Redundancy Type.....	3-30
B-1	Response Files for Oracle Database.....	B-3
B-2	Response files for Oracle Grid Infrastructure.....	B-3

Preface

Oracle Grid Infrastructure Installation Guide for IBM AIX on POWER Systems (64-Bit) explains how to configure a server in preparation for installing and configuring an Oracle Grid Infrastructure installation (Oracle Clusterware and Oracle Automatic Storage Management). It also explains how to configure a server and storage in preparation for an Oracle Real Application Clusters (Oracle RAC) installation.

Intended Audience

Oracle Grid Infrastructure Installation Guide for IBM AIX on POWER Systems (64-Bit) provides configuration information for network and system administrators, and database installation information for database administrators (DBAs) who install and configure Oracle Clusterware and Oracle Automatic Storage Management in 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 Oracle Grid Infrastructure installation and configuration is completed successfully, a system administrator should only need to provide configuration information and to grant access to the database administrator to run scripts as `root` during an Oracle RAC installation.

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

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/support/contact.html> or visit <http://www.oracle.com/accessibility/support.html> if you are hearing impaired.

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 Oracle Grid Infrastructure installation, and review those installation guides for additional information. You can install either Oracle databases for a standalone server on an Oracle 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 IBM AIX on POWER Systems (64-Bit)*

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.

Installation Guides

- *Oracle Diagnostics Pack Installation Guide*
- *Oracle Database Installation Guide for IBM AIX on POWER Systems (64-Bit)*
- *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 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*

Generic Documentation

- *Getting Started with the Oracle Diagnostics Pack*
- *Oracle Database 2 Day DBA*
- *Oracle Database Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle Database New Features Guide*
- *Oracle Database Net Services Administrator's Guide*
- *Oracle Database Reference*

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

<https://shop.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://www.oracle.com/technology/membership/index.html>

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

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

Conventions

The following text conventions are used in this document:

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

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 Oracle Automatic Storage Management), and Oracle Real Application Clusters (Oracle RAC). This guide replaces *Oracle Clusterware Installation Guide*. The topics in this section are:

- [Desupported Options](#)
- [New Features for Release 2 \(11.2.0.2\)](#)
- [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 Oracle Automatic Storage Management diskgroups, or in a supported shared file system.

New Features for Release 2 (11.2.0.2)

The following is a list of new features for Release 2 (11.2.0.2):

Enhanced Patch Set Installation

Starting with the release of the 11.2.0.2 patch set for Oracle Grid Infrastructure 11g Release 2 (Oracle Clusterware and Oracle Automatic Storage Management), Oracle Grid Infrastructure patch sets are full installations of the Oracle Grid Infrastructure software. Note the following changes with the new patch set packaging:

- Direct upgrades from previous releases (11.x, 10.x) to the most recent patch set are supported.
- New installations consist of installing the most recent patch set, rather than installing a base release and then upgrading to a patch release.
- Out-of-place patch set upgrades only are supported. An out-of-place upgrade is one in which you install the patch set into a new, separate home.

See Also: My Oracle Support note 1176177, "Important Changes to Oracle Database Patch Sets Starting With 11.2.0.2", available from the following URL:

<https://support.oracle.com>

Oracle ACFS and ADVM Support for AIX

Oracle ASM 11g release 2 (11.2.0.2) and later for AIX provides support for Oracle Automatic Storage Management Cluster File System (Oracle ACFS), including ACFS Snapshots, and Oracle ASM Dynamic Volume Manager (ADVM).

ACFS (including ACFS Snapshots) and ADVM are supported only on AIC 6.1 TL4 SP2 and later updates to AIX 6.1 (on PPC64 only).

Cluster Health Monitor Included with Oracle Clusterware

Cluster Health Monitor gathers operating system metrics in real time, and stores them in its repository for later analysis, so that it can determine the root cause of many Oracle Clusterware and Oracle RAC issues with the assistance of Oracle Support.

Cluster Health Monitor also works in conjunction with Oracle Database Quality of Service Management (QoS) by providing metrics to detect memory over-commitment on a node. QoS Management can shut down services on overloaded nodes to relieve stress, to and preserve existing workloads.

To support QoS Management, Oracle Database Resource Manager and metrics have been enhanced to support fine-grained performance metrics, and also can manage workloads with user-defined performance classes.

Grid Installation Owner and ASMOPER

During installation, in the Privileged Operating System Groups window, it is now optional to designate a group as the OSOPER for ASM group. If you choose to create an OSOPER for ASM group, then you can enter a group name configured on all cluster member nodes for the OSOPER for ASM group. In addition, the Oracle Grid Infrastructure installation owner no longer is required to be a member.

New Software Updates Option

Use the Software Updates feature to dynamically download and apply software updates as part of the Oracle Database installation. You can also download the updates separately using the download Updates option and later apply them during the installation by providing the location where the updates are present.

Redundant Interconnect Usage

In previous releases, to make use of redundant networks for the interconnect, bonding, trunking, teaming, or similar technology was required. Oracle Grid Infrastructure and Oracle RAC can now make use of redundant network interconnects, without the use of other network technology, to enhance optimal communication in the cluster. This functionality is available starting with Oracle Database 11g Release 2 (11.2.0.2).

Redundant Interconnect Usage enables load-balancing and high availability across multiple (up to four) private networks (also known as interconnects).

Oracle Database Quality of Service Management

The Database Quality of Service (QoS) Management Server allows system administrators to manage application service levels hosted in Oracle Database clusters by correlating accurate run-time performance and resource metrics and analyzing with

an expert system to produce recommended resource adjustments to meet policy-based performance objectives.

See Also: *Oracle Database Quality of Service Management User's Guide*

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), Oracle Automatic Storage Management (Oracle 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 configures Oracle ASM and Oracle Clusterware.

The installation of the combined products is called Oracle Grid Infrastructure. However, Oracle Clusterware and Oracle Automatic Storage Management remain separate products.

See Also: *Oracle Database Installation Guide* for IBM AIX on POWER Systems (64-Bit) for information about how to install Oracle Grid Infrastructure (Oracle ASM and Oracle Clusterware binaries) for a standalone server. This feature helps to ensure high availability for single-instance servers

Automatic Storage Management and Oracle Clusterware Files

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

This feature enables Oracle 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 Oracle 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.

Oracle 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 Oracle 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 delete nodes is minimized.

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 Oracle 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 Oracle 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 Oracle Enterprise Manager deployed, you can provision additional nodes added to the cluster using Oracle Enterprise Manager.

SCAN for Simplified Client Access

With this release, the single client access name (SCAN) is the host name 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 Oracle Automatic Storage Management (Oracle 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 Oracle ASM administration and database administration.

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

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](#)

Typical and Advanced Installation

There are 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, or more granularity in specifying Oracle Automatic Storage Management roles.

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.

Fixup scripts are generated during installation. You are prompted to run scripts as `root` in a separate terminal session. When you run scripts, they complete the following configuration tasks:

- If necessary sets kernel parameters required for installation and runtime to at least the minimum value.
- Reconfigures primary and secondary group memberships for the installation owner, if necessary, for the Oracle Inventory directory and the operating system privileges groups.

- Sets shell limits if necessary to required values.

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 Oracle 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

Verify System Requirements

Enter the following commands to check available memory:

```
# /usr/sbin/lsattr -E -l sys0 -a realmem
# /usr/sbin/lspcs -a
```

The minimum required RAM is 1.5 GB for Oracle Grid Infrastructure for a cluster, or 2.5 GB for Oracle 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. If the swap space and the Grid home are on the same filesystem, then add together their respective disk space requirements for the total minimum space required.

Verify the space available for Oracle Clusterware files. For example:

GPFS:

```
/usr/bin/df -k
```

To check raw device volumes in preparation for installing Oracle ASM disk groups, use the following checks:

Raw Logical Volumes in Concurrent VG (HACMP): In the following example, the variable *lv_name* is the name of the raw logical volume whose space you want to verify:

```
lslv lv_name
```

Raw hard disks: In the following example, the variable *rhdisk#* is the raw hard disk number that you want to verify, and the variable *size_mb* is the size in megabytes of the partition that you want to verify:

```
lsattr -El rhdisk# -a size_mb
```

If you use normal redundancy for OCR and voting disk files, which is 3 Oracle Cluster Registries (OCR) and 3 voting disks, ideally, in different file systems on independent disks, then you should have at least 1 GB of disk space available on separate physical disks reserved for Oracle Clusterware files. Each file system for the Oracle Clusterware files should be at least 280 MB in size.

Note: You cannot install OCR or voting disk files on raw partitions. You can install only on Oracle ASM, or on supported network-attached storage or cluster file systems. The only use for raw devices is as ASM disks.

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 12 GB of space for the Oracle Grid Infrastructure for a cluster home (Grid home) This includes Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) files and log files, and includes the Cluster Health Monitor repository.

```
/usr/bin/df -k /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`.

Check Network Requirements

Ensure that you have the following available:

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

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:

```
NE-Sa89
```

If you require a SCAN that is longer than 15 characters, then be aware that the cluster name defaults to the first 15 characters of the SCAN.

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.

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 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 cluster alias 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 SCANS.

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

Redundant Interconnect Usage

In previous releases, to make use of redundant networks for the interconnect, bonding, trunking, teaming, or similar technology was required. Oracle Grid Infrastructure and Oracle RAC can now make use of redundant network interconnects, without the use of other network technology, to enhance optimal communication in the cluster. This functionality is available starting with Oracle Database 11g Release 2 (11.2.0.2).

Redundant Interconnect Usage enables load-balancing and high availability across multiple (up to 4) private networks (also known as interconnects).

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, or as "do not use." 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.

Redundant Interconnect Usage cannot protect interfaces used for public communication. If you require high availability or load balancing for public interfaces, then use a third party solution. Typically, bonding, trunking or similar technologies can be used for this purpose.

You can enable Redundant Interconnect Usage for the private network by selecting multiple interfaces to use as private interfaces. Redundant Interconnect Usage creates a redundant interconnect when you identify more than one interface as private. This functionality is available starting with Oracle Grid Infrastructure 11g Release 2 (11.2.0.2).

Check Operating System Packages

Refer to the tables listed in [Section , "Checking the Software Requirements"](#) for the list of required packages for your operating system.

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):

```
# mkgroup -'A' id='1000' adms='root' oinstall
# mkgroup -'A' id='1031' adms='root' dba
# mkuser id='1101' pgrp='oinstall' groups='dba' home='/home/grid' grid
# mkuser id='1101' pgrp='oinstall' groups='dba' home='/home/oracle' 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/
```

This set of commands creates a single installation owner, with required system privileges groups to grant the OraInventory system privileges (`oinstall`), and to grant the OSASM/SYSASM and OSDBA/SYSDBA system privileges. It also creates the Oracle base for both Oracle Grid Infrastructure and Oracle RAC, `/u01/app/oracle`. It creates the Grid home (the location where Oracle Grid Infrastructure binaries are stored), `/u01/app/11.2.0/grid`.

Ensure that the Oracle Grid Infrastructure installation owner account has the capabilities `CAP_NUMA_ATTACH`, `CAP_BYPASS_RAC_VMM`, and `CAP_PROPAGATE`.

To check existing capabilities, enter the following command as `root`; in this example, the Grid installation user account is `grid`:

```
# /usr/bin/lsuser -a capabilities grid
```

To add capabilities, enter a command similar to the following:

```
# /usr/bin/chuser
capabilities=CAP_NUMA_ATTACH,CAP_BYPASS_RAC_VMM,CAP_PROPAGATE grid
```

Set the password on the grid installation owner account:

```
passwd grid
```

Repeat this process for each cluster member node.

Configure Oracle Installation Owner Shell Limits

Set shell limits for the Oracle Grid Infrastructure installation owner and for `root` to unlimited. Verify that unlimited is set for both accounts either by using the `smit` utility or by editing the `/etc/security/limits` file. The `root` user requires these

settings because the crs daemon (`crsd`) runs as root. Add the following lines to the limits file:

```
default:  
    fsize = -1  
    core = 2097151  
    cpu = -1  
    data = -1  
    rss = -1  
    stack = -1  
    nofiles = -1
```

Check Storage

You must have space available either on a supported file system, or on Oracle 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.

Prepare Storage for Oracle Automatic Storage Management

Review the relevant sections in Chapter 3 for the installation option you want to configure.

See Also: [Chapter 3, "Configuring Storage for Grid Infrastructure for a Cluster and Oracle Real Application Clusters \(Oracle RAC\)"](#)

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, "Installing Oracle Grid Infrastructure for a Cluster"](#)

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 for a cluster, and that you may need to complete if you intend to install Oracle Real Application Clusters (Oracle RAC) on the cluster.

This chapter contains the following topics:

- [Reviewing Upgrade Best Practices](#)
- [Installation Fixup Scripts](#)
- [Logging In to a Remote System as root Using X Terminal](#)
- [Creating Groups, Users and Paths for Oracle Grid Infrastructure](#)
- [Checking the Hardware Requirements](#)
- [Checking the Network Requirements](#)
- [Identifying the Software Requirements](#)
- [Checking the Software Requirements](#)
- [Verifying UDP and TCP Kernel Parameters](#)
- [Enabling Core File Creation](#)
- [Tuning AIX System Environment](#)
- [Network Time Protocol Setting](#)
- [Automatic SSH Configuration During Installation](#)
- [Configuring Grid Infrastructure Software Owner User Environments](#)
- [Running the rootpre.sh Script](#)
- [Requirements for Creating an Oracle Grid Infrastructure Home Directory](#)

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

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 of time. 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://support.oracle.com>

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.
- Reconfigures primary and secondary group memberships for the installation owner, if necessary, for the Oracle Inventory directory and the operating system privileges groups.
- Sets shell limits if necessary to required values.

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
```

Logging In to a Remote System as root Using X Terminal

Before you install the Oracle software, you must complete several tasks as the `root` user on the system where you install Oracle software. To complete tasks as the `root` user on a remote server, you need to enable remote display as root.

Note: If you log in as another user (for example, `oracle`), then you need to 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 not installing the software on the local system, then enter a command using the following syntax to enable remote hosts to display X applications on the local X server:

```
$ xhost + remote_host
```

where *remote_host* 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 remote_host
```

where *remote_host* is the fully qualified remote hostname. For example:

```
$ ssh 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 and start a terminal session on that system, for example, an X terminal (`xterm`).
4. If you are not logged in as the `root` user on the remote system, then enter the following command to switch user to `root`:

```
$ su - root
password:
#
```

Creating Groups, Users and Paths for Oracle Grid Infrastructure

Log in as `root`, and use the following instructions to locate or create groups and users required for installation.

Note: Ensure that all group and user numbers are identical on all cluster member nodes.

- [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 an Oracle Grid Infrastructure installation, both Oracle Clusterware and Oracle Automatic Storage Management are installed. You no longer can have separate Oracle Clusterware installation owners and Oracle Automatic Storage Management installation owners.

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:grid,oracle
```

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:

```
# mkgroup id=1000 adms='root' oinstall
```

The preceding command creates the `oraInventory` group `oinstall`, with the group ID number 1000. Members of the `oraInventory` 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 Oracle Grid Infrastructure for a Cluster software as the `oraInventory` group. Ensure that this group is available as a primary group for all planned Oracle software installation owners.

Note: Group and user IDs must be identical on all nodes in the cluster. Check to make sure that the group and user IDs you want to use are available on each cluster member node, and confirm that the primary group for each grid infrastructure for a cluster installation owner has the same name and group ID.

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.

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 Oracle 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 Oracle 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.

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, 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).

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 , "Creating Job Role Separation Operating System Privileges Groups and Users."](#)

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.

Creating a New User use the following procedure to create a new user:

1. Enter the following command:

```
# smit security
```

2. On the **Security & Users** menu, select **Users**.
3. On the **Users** menu, select **Add a User**.
4. Choose the appropriate menu items to create the Oracle Grid Infrastructure software installation owner (`grid`). In the **Primary GROUP** field, specify the Oracle Inventory group. Make a note of the information you provide in the entry fields, so that you can provide the same value on other nodes.

Note: The UID and GID for the Oracle Clusterware user must be less than 65536.

5. Press Enter, then Esc+0 (Exit) to exit.
6. Set the password of the Oracle Grid Infrastructure software installation owner (`grid`). For example:

```
# passwd grid
```

7. Ensure that the Oracle Grid Infrastructure software installation owner (`grid`) has the capabilities `CAP_NUMA_ATTACH`, `CAP_BYPASS_RAC_VMM`, and `CAP_PROPAGATE`.

To check existing capabilities, enter the following command as `root`:

```
# /usr/bin/luser -a capabilities grid
```

To add capabilities, enter a command similar to the following:

```
# /usr/bin/chuser
capabilities=CAP_NUMA_ATTACH,CAP_BYPASS_RAC_VMM,CAP_PROPAGATE grid
```

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

Modifying an Existing User use the following procedure to modify an existing user:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to modify the `grid` installation owner user.
3. In the **Primary GROUP** field, specify the Oracle Inventory group, for example `oinstall`.

4. Press Enter, then Esc+0 (Exit) to exit.
5. Repeat this procedure on all of the other nodes in the cluster.

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 `u[00-99]/app`, and it must be writable by any member of the `oraInventory` (`oinstall`) group. The OFA path for the Oracle base is `/u01/app/user`, where `user` is the name of the software installation owner.

Oracle recommends that you create an Oracle Grid Infrastructure Grid home and Oracle base homes manually, particularly if you have separate Oracle Grid Infrastructure for a cluster and Oracle Database software owners, so that you can separate log files.

For example:

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

Note: Placing Oracle Grid Infrastructure for a cluster binaries on a cluster file system is not supported.

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, Oracle 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.

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

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

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 Oracle Grid Infrastructure 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`.

Oracle 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 Oracle 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 Oracle ASM Database Administrator group (OSDBA for ASM, typically `asmdba`)

Members of the Oracle ASM Database Administrator group (OSDBA for ASM) are granted read and write access to files managed by Oracle ASM. The Oracle Grid Infrastructure installation owner and all Oracle Database software owners must be a member of this group, and all users with OSDBA membership on databases that have access to the files managed by Oracle ASM must be members of the OSDBA group for ASM.

- Members of the Oracle 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 Oracle 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`.

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](#)

Creating the OSDBA Group to Prepare for Database Installations If you intend to install Oracle Database to use with the Oracle Grid Infrastructure installation, then you must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist; for example, if this is the first installation of Oracle Database software on the system
- An OSDBA group exists, but you want to give a different group of operating system users database administrative privileges for a new Oracle Database installation

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

```
# mkgroup -'A' id='1031' adms='root' dba
```

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 either by using `smit` or by using shell command lines. Use the group name `oper` unless a group with that name already exists.

```
# mkgroup -'A' id='1032' adms='root' oper1
```

Creating the OSASM Group If the OSASM group does not exist or if you require a new OSASM group, then create it either by using `smit` or by using shell command lines. Use the group name `asmadmin` unless a group with that name already exists:

```
# mkgroup -'A' id='1020' adms='root' asmadmin
```

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 either by using `smit` or by using shell command lines. Use the group name `asmoper` unless a group with that name already exists:

```
# mkgroup -'A' id='1022' adms='root' asmoper
```

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 either by using `smit` or by using shell command lines. Use the group name `asmdba` unless a group with that name already exists. For example:

```
# mkgroup -'A' id='1021' adms='root' asmdba
```

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`.

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 [Section , "Modifying an Existing Oracle Software Owner User"](#).
- 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.

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, use `smit` or enter a command similar to the following:

```
# mkuser id='1101' pgrp='oinstall' groups='dba,asmdba' home='/home/oracle'
oracle
```

2. Set the password of the `oracle` user:

```
# passwd oracle
```


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 update it using `smit`:

1. Start SMIT by entering `smit` at the shell prompt.
2. From the Main Menu, make the following selections:
 - Security and Users
 - Groups
 - Change/Show Characteristics of a Group

The utility displays a form in which you can type the name of a specific group.

3. Either fill in the group name or use the F4 key to highlight a group name and press the **Enter** key.

The utility displays a form that provides the following fields:

- Group NAME, which is the group name for the account.
 - Group ID, which is the Group Identification Number (GID).
 - User list, which is a list of users who are members of the group. Use the F4 key to display a list of available users and the F7 key to mark the users you want to add.
4. Modify these fields as needed and press the Enter key to exit the form.

Creating Identical Database Users and Groups on Other Cluster Nodes Oracle software owner users and the Oracle Inventory, OSDBA, and OSOPER groups must exist and be identical on all cluster nodes. To create these identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, and then create the user and groups with the same name and ID on the other cluster nodes.

Note: You must complete the following procedures only if you are using local users and groups. If you are using users and groups defined in a directory service such as NIS, then they are already identical on each cluster node.

Identifying Existing User and Group IDs

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

1. Enter a command similar to the following (in this case, to determine a user ID for the `oracle` user):

```
# id oracle
```

The output from this command is similar to the following:

```
uid=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 (`gid`) 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. Create groups and users as needed, either by using `smit` or by entering command lines. To use command line entries, 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 `id` option to specify the correct `gid` for each group.

```
# mkgroup -'A' id='1000' adms='root' oinstall
# mkgroup -'A' id='1020' adms='root' asmadmin
# mkgroup -'A' id='1021' adms='root' asmdba
# mkgroup -'A' id='1022' adms='root' asmoper
# mkgroup -'A' id='1031' adms='root' dba
# mkgroup -'A' id='1032' adms='root' oper
```

Note: If the group already exists, then use `smit` 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 Oracle Grid Infrastructure (`grid`) user, use `smit` or enter a command similar to the following (in this example, to create the `oracle` user):

```
# mkuser id='1001' pgrp='oinstall' groups='dba,asmdba' home='/home/oracle'
oracle
```

In the preceding command:

- The `id` option specifies the user ID, which must be the user ID that you identified in the previous subsection
- The `pgrp` option specifies the primary group, which must be the Oracle Inventory group, for example `oinstall`
- The `groups` 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 `smit` 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 [Section , "Environment Requirements for Oracle Grid Infrastructure Software Owner."](#)

Example of Creating Standard Groups, Users, and Paths

The following is an example of how to use command lines 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 Oracle 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:

```
# mkgroup -'A' id='1000' adms='root' oinstall
# mkgroup -'A' id='1031' adms='root' dba
# mkuser id='1100' pgrp='oinstall' groups='dba' home='/home/grid' grid
# mkuser id='1101' pgrp='oinstall' groups='dba' home='/home/oracle' oracle
# mkdir -p /u01/app/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, with its Oracle base directory `/u01/app/grid`.
- An Oracle Database owner (`oracle`) with the oraInventory group as its primary group, and the OSDBA group as its secondary group, with its Oracle base directory `/u01/app/oracle`.
- `/u01/app` owned by `grid:oinstall` with 775 permissions before installation, and by root after the `root.sh` script is run during installation. This ownership and permissions enables OUI to create the Oracle Inventory directory, in the path `/u01/app/oraInventory`.
- `/u01` owned by `grid:oinstall` before installation, and by root after the `root.sh` script is run during installation.
- `/u01/app/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/grid` owned by `grid:oinstall` with 775 permissions before installation, and 755 permissions after installation.
- `/u01/app/oracle` owned by `oracle:oinstall` with 775 permissions.

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:

```
# mkgroup -'A' id='1000' adms='root' oinstall
# mkgroup -'A' id='1020' adms='root' asmadmin
# mkgroup -'A' id='1021' adms='root' asmdba
# mkgroup -'A' id='1031' adms='root' dba1
# mkgroup -'A' id='1041' adms='root' dba2
# mkgroup -'A' id='1022' adms='root' asmoper
# mkuser id='1100' pgrp='oinstall' groups='asmadmin' adms='root' home='/home/grid'
grid
# mkuser id='1101' pgrp='oinstall' groups='dba1,asmdba' adms='root'
home='/home/oracle' oracle1
# mkuser id='1102' pgrp='oinstall' groups='dba2,asmdba' adms='root'
home='/home/oracle' oracle2
# mkdir -p /u01/app/11.2.0/grid
# mkdir -p /u01/app/grid
# chown -R grid:oinstall /u01/app/grid
# 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 this 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 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) group as a secondary group.
- 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 for the grid installation owner /u01/app/grid owned by grid:oinstall with 775 permissions, and changed during the installation process to 755 permissions.
- 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`).
- `/u01/app/oraInventory`. This path remains owned by `grid:oinstall`, to enable other Oracle software owners to write to the central inventory.

Checking the Hardware Requirements

- Select servers with the same chip architecture.
- Ensure that the server is started with run level 2 (default or Normal Multi-User mode).
- Ensure servers run the same operating system level, APARs and filesets. 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 4 GB of physical RAM for Oracle Grid Infrastructure for a cluster installations without Oracle RAC; at least an additional 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:

Available RAM	Swap Space Required
Between 4 GB and 16 GB	Equal to the size of RAM
More than 16 GB	16 GB

Note: On AIX systems with 4 GB or more of memory, Oracle recommends that you set the paging space to an initial setting of half the size of RAM plus 4 GB, with an upper limit of 32 GB. During installation, to optimize paging, monitor the paging space use in a separate window. Use the command `chps` to increase or decrease the paging space size. The output of `chps` should indicate paging space use of less than 25 percent on a properly configured system. Refer to *Oracle Database Administrator's Reference for AIX* for more information about configuring paging space.

- 1 GB of disk space in the `/tmp` directory
- 12 GB of space for the Oracle Grid Infrastructure for a cluster home (Grid home) This includes Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) files and log files, and includes the Cluster Health Monitor repository.

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

- 7.5 GB of disk space for the Oracle Database files (Oracle base)
- 2 GB of disk space for a preconfigured database that uses file system storage (optional)

If you choose to configure automated backups, then you require additional disk space, either on a file system or in an Oracle Automatic Storage Management disk group.

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:

```
# /usr/sbin/lssattr -E -l sys0 -a realmem
```

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 available RAM and swap space, enter the following command:

```
# /usr/sbin/lssps -s
```

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

```
# /usr/sbin/lssps -a
```

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

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

```
# /usr/bin/df -k /tmp
```

If there is less than 1 GB of disk space available in the `/tmp` directory, then complete one of the following steps:

- Delete unnecessary files from the `/tmp` directory to make available the disk space required.
- Set the `TEMP` and `TMPDIR` environment variables when setting the `oracle` user's environment.
- Extend the file system that contains the `/tmp` directory. If necessary, contact your system administrator for information about extending file systems.

See Also: [Section , "Configuring Grid Infrastructure Software Owner User Environments"](#)

- To determine the amount of free disk space on the system, use one of the following commands, depending on where you intend to place Oracle Clusterware files:

GPFS:

```
# /usr/bin/df -k
```

Raw Logical Volumes in Concurrent VG (HACMP): in the following example, the variable *lv_name* is the name of the raw logical volume whose space you want to verify:

```
# lslv lv_name
```

Raw hard disks; in the following example, the variable *rhdisk#* is the raw hard disk number that you want to verify, and the variable *size_mb* is the size in megabytes of the partition that you want to verify:

```
# lsattr -El rhdisk# -a size_mb
```

- To determine if the system is started in 64-bit mode, enter the following command:

```
# bootinfo -K
```

The result of this command should be 64, indicating that the 64-bit kernel is enabled.

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](#)
- [Broadcast Requirements for Networks Used by Oracle Grid Infrastructure](#)
- [Multicast Requirements for Networks Used by Oracle Grid Infrastructure](#)
- [DNS Configuration for Domain Delegation to Grid Naming Service](#)
- [Grid Naming Service Configuration Example](#)
- [Manual IP Address Configuration Example](#)
- [Network Interface Configuration Options](#)

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

<https://support.oracle.com>

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).

With Redundant Interconnect Usage, you can identify multiple interfaces to use for the cluster private network, without the need of using bonding or other technologies. This functionality is available starting with Oracle Database 11g Release 2 (11.2.0.2).

When you define multiple interfaces, Oracle Clusterware creates from one to four highly available IP (HAIP) addresses. Oracle RAC and Oracle ASM instances use these interface addresses to ensure highly available, load-balanced interface communication between nodes. The installer enables Redundant Interconnect Usage to provide a high availability private network. Oracle Database, CSS, OCR, CRS, CTSS, and EVM components in 11.2.0.2 employ Redundant Interconnect Usage.

By default, Oracle Grid Infrastructure software uses all of the HAIP addresses for private network communication, providing load-balancing across the set of interfaces you identify for the private network. If a private interconnect interface fails or become non-communicative, then Oracle Clusterware transparently moves the corresponding HAIP address to one of the remaining functional interfaces.

Note: If you define more than four interfaces as private network interfaces, be aware that Oracle Clusterware activates only four of the interfaces at a time. However, if one of the four active interfaces fails, then Oracle Clusterware transitions the HAIP addresses configured to the failed interface to one of the reserve interfaces in the defined set of private interfaces.

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

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

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

Note: NIC bonding is not required to use multiple NICs. During installation, you are given the opportunity to specify the planned use for all interfaces detected by the Installer. Oracle Clusterware uses all NICs that you identify as Private for the private interconnect.

- If you install Oracle Clusterware using OUI, then 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 adapters should be the same on all nodes. This restriction does not apply if you use cloning, either to create a new cluster, or to add nodes to an existing cluster.

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`.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about how to add nodes using cloning

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

Note: UDP is the default interface protocol for Oracle RAC, and 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`.

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 manage IP addresses and name resolution in the cluster in one of the following ways:

- Dynamic IP address assignment using Oracle Grid Naming Service (GNS). If you select this option, then network administrators assign static IP address for the physical host name and dynamically allocated IPs for the Oracle Clusterware managed VIP addresses. In this case, IP addresses for the VIPs are assigned by a DHCP 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 IP address assignment. If you select this option, then network administrators assign a fixed IP address for each physical host name in the cluster and for IPs for the Oracle Clusterware managed VIPs. In addition, domain name server (DNS) based static name resolution is used for each node. Selecting this option requires that you request network administration updates when you modify the cluster.

Note: Oracle recommends that you use a static host name for all server node public hostnames.

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

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.

Note: If you have vendor clusterware installed, then you cannot choose to use GNS, because the vendor clusterware does not support it.

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 , "Understanding Network Addresses"](#) for more information about network addresses

Broadcast Requirements for Networks Used by Oracle Grid Infrastructure

Broadcast communications (ARP and UDP) must work properly across all the public and private interfaces configured for use by Oracle Grid Infrastructure release 2 patchset 1 (11.2.0.2) and later releases.

The broadcast must work across any configured VLANs as used by the public or private interfaces.

Multicast Requirements for Networks Used by Oracle Grid Infrastructure

With Oracle Grid Infrastructure release 2 (11.2), on each cluster member node, the Oracle mDNS daemon uses multicasting on all interfaces to communicate with other nodes in the cluster.

With Oracle Grid Infrastructure release 2 patchset 1 (11.2.0.2) and later releases, multicasting is required on the private interconnect. For this reason, at a minimum, you must enable multicasting for the cluster:

- Across the broadcast domain as defined for the private interconnect
- On the IP address subnet ranges 224.0.0.0/24 and 230.0.1.0/24

You do not need to enable multicast communications across routers.

DNS Configuration for Domain Delegation to Grid Naming Service

If you plan to use GNS, then before Oracle 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. The following is an overview of what needs to be done for domain delegation. Your actual procedure may be different from this example.

Configure the DNS to send GNS name resolution requests using delegation:

1. In the DNS, create an entry for the GNS virtual IP address, where the address uses the form `gns-server.CLUSTERNAME.DOMAINNAME`. For example, where the cluster name is `mycluster`, and the domain name is `example.com`, and the IP address is `192.0.2.1`, create an entry similar to the following:

```
mycluster-gns.example.com A 192.0.2.1
```

The address you provide must be routable.

2. Set up forwarding of the GNS subdomain to the GNS virtual IP address, so that GNS resolves addresses to the GNS subdomain. To do this, create a BIND configuration entry similar to the following for the delegated domain, where `cluster01.example.com` is the subdomain you want to delegate:

```
cluster01.example.com NS mycluster-gns.example.com
```

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

```
options attempts: 2
options timeout: 1

search cluster01.example.com example.com
nameserver xxx.xxx.xxx.42
nameserver xxx.xxx.xxx.15
```

`/etc/nsswitch.conf` controls name service lookup order. In some system configurations, the Network Information System (NIS) can cause problems with 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
```

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

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 2–1 Grid Naming Service Example Network

Identity	Home Node	Host Node	Given Name	Type	Address	Address Assigned By	Resolved By
GNS VIP	None	Selected by Oracle Clusterware	mycluster-gns.example.com	virtual	192.0.2.1	Fixed by net administrator	DNS
Node 1 Public	Node 1	node1	node1 ¹	Public	192.0.2.101	Fixed	GNS
Node 1 VIP	Node 1	Selected by Oracle Clusterware	node1-vip	Virtual	192.0.2.104	DHCP	GNS
Node 1 Private	Node 1	node1	node1-priv	Private	192.168.0.1	Fixed or DHCP	GNS
Node 2 Public	Node 2	node2	node2 ¹	Public	192.0.2.102	Fixed	GNS
Node 2 VIP	Node 2	Selected by Oracle Clusterware	node2-vip	Virtual	192.0.2.105	DHCP	GNS
Node 2 Private	Node 2	node2	node2-priv	Private	192.168.0.2	Fixed or DHCP	GNS
SCAN VIP 1	none	Selected by Oracle Clusterware	mycluster-scan.cluster01.example.com	virtual	192.0.2.201	DHCP	GNS
SCAN VIP 2	none	Selected by Oracle Clusterware	mycluster-scan.cluster01.example.com	virtual	192.0.2.202	DHCP	GNS
SCAN VIP 3	none	Selected by Oracle Clusterware	mycluster-scan.cluster01.example.com	virtual	192.0.2.203	DHCP	GNS

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

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 2–2 Manual Network Configuration Example

Identity	Home Node	Host Node	Given Name	Type	Address	Address Assigned By	Resolved By
Node 1 Public	Node 1	node1	node1 ¹	Public	192.0.2.101	Fixed	DNS
Node 1 VIP	Node 1	Selected by Oracle Clusterware	node1-vip	Virtual	192.0.2.104	Fixed	DNS and hosts file
Node 1 Private	Node 1	node1	node1-priv	Private	192.168.0.1	Fixed	DNS and hosts file, or none
Node 2 Public	Node 2	node2	node2 ¹	Public	192.0.2.102	Fixed	DNS
Node 2 VIP	Node 2	Selected by Oracle Clusterware	node2-vip	Virtual	192.0.2.105	Fixed	DNS and hosts file
Node 2 Private	Node 2	node2	node2-priv	Private	192.168.0.2	Fixed	DNS and hosts file, or none
SCAN VIP 1	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.201	Fixed	DNS
SCAN VIP 2	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.202	Fixed	DNS
SCAN VIP 3	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.203	Fixed	DNS

¹ Node 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 (`en1`, 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.

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.

Identifying the 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 [Section , "Checking the Software Requirements."](#)

Requirements listed here are current as of the initial release date. To obtain the most current information about kernel requirements, refer to the online version on the Oracle Technology Network (OTN) at the following URL:

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

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.

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.

Table 2–3 AIX Operating System Kernel Requirements

Item	Requirement
Operating systems	AIX 6.1 TL 09 SP1 ("6100-02-01), 64-bit kernel (Note: Ensure that the operating system level is "Technology Level 02 Service Pack 01 or higher") AIX 5L V5.3 TL 09 SP1 ("5300-09-01"), 64 bit kernel or later
AIX 6L operating system filesets	The following operating system filesets are required: bos.adt.base bos.adt.lib bos.adt.libm bos.perf.libperfstat bos.perf.perfstat bos.perf.proctools rsct.basic.rte rsct.compat.clients.rte xlC.aix61.rte 10.1.0.0 (or later) You must have the IBM XL C/C++ runtime filesets for installation, but you do not require the C/C++ compilers. You do not require a license for the XL C/C++ runtime filesets. Version: IBM XL C/C++ Enterprise Edition for AIX, V9.0 September 2008 PTF

Table 2–3 (Cont.) AIX Operating System Kernel Requirements

Item	Requirement
AIX 5L operating system filesets	<p>The following operating system filesets are required:</p> <ul style="list-style-type: none"> bos.adt.base bos.adt.lib bos.adt.libm bos.perf.libperfstat bos.perf.perfstat bos.perf.proctools rsct.basic.rte rsct.compat.clients.rte xlC.aix50.rte 10.1.0.0 (or later)
Obtaining C/C++ Compilers	<p>You must have the IBM XL C/C++ runtime filesets for installation, but you do not require the C/C++ compilers. You do not require a license for the XL C/C++ runtime filesets. Version: IBM XL C/C++ Enterprise Edition for AIX, V9.0 September 2008 PTF</p> <p>To obtain the XLC runtime filesets, Oracle Database users who do not install the IBM XL C/C++ Enterprise Edition V8.0 or V9.0 compiler should install the IBM XL C/C++ Enterprise Edition V8.0 for AIX Runtime Environment Component. For AIX5L, this contains xLC.aix.rte 8.0.0.8 and xLC.rte 8.0.0.8.</p> <p>Download the C/C++ compiler from the following Web site: http://www-1.ibm.com/support/docview.wss?uid=swg24015075</p> <p>Download the C runtime environment file sets, with no license requirement, from the following Web site: http://www-1.ibm.com/support/docview.wss?uid=swg24015077</p>

Table 2–3 (Cont.) AIX Operating System Kernel Requirements

Item	Requirement
Oracle RAC	<p>Review the following additional requirements if needed for your configuration:</p> <p>High Availability Cluster Multi-Processing (HACMP) 5.4.1</p> <p>Note: HACMP is required only if you want to use raw logical volumes for Oracle Clusterware or database file storage. However, it is supported for all installations. You cannot use raw devices for OCR or voting disk files.</p> <p>If you do not want to use HACMP, then you must not have HACMP installed on your system. If you have previously installed HACMP, then you must remove the following:</p> <ul style="list-style-type: none"> ■ HACMP filesets (<code>cluster.es.*</code>) ■ <code>rsct.hacmp.rte</code> ■ <code>rsct.compat.basic.hacmp.rte</code> ■ <code>rsct.compat.clients.hacmp.rte</code> <p>If you want to use HACMP, then review patch sets to ensure that you have required patches. Changes in the fileset packaging of HACMP 5.4, including 5.4.1, require updates to the Oracle <code>rootpre.sh</code> script. Download and install patch 6718715 before installing Oracle Clusterware.</p>
	General Parallel File System (GPFS):
	<p>AIX 6L:</p> <p><code>gpfs.base 3.2.1.8</code> or later.</p>
	<p>AIX 5L:</p> <p><code>gpfs.base 3.2.1.8</code> or later</p>
	<p>Note: GPFS is required only if you want to use a cluster file system for Oracle Clusterware or database files.</p>
ADA	OC Systems PowerAda 5.4d
JDK	<p>Use one of the following Java versions:</p> <p>Java 6 64-bit 6.0.0.50 IZ30726 (SR2)</p> <p>Java 5 64-bit 5.0.0.250 IZ55274 (SR8a)</p>
Pro*FORTRAN	IBM XL Fortran v. 11.1 April 2008 PTF for AIX
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK), GNU Compiler Collection (GCC)	<p>Note: If you do not install the C/C++ compilers, then you require the C/C++ runtime filesets for installation as described in the "Operating system filesets" row in this table.</p> <ul style="list-style-type: none"> ■ February 2007 XL C/C++ Enterprise Edition V8.0 for AIX PTF 8.0: <p>You can download the PTF from the following link:</p> <p>http://www-1.ibm.com/support/docview.wss?uid=swg24015075</p> <p>Download the C++ runtime environment file sets, with no license requirement, from the following Web site:</p> <p>http://www-1.ibm.com/support/docview.wss?uid=swg24015077</p>
	<ul style="list-style-type: none"> ■ <code>gcc 3.45</code>
Pro*COBOL	<ul style="list-style-type: none"> ■ IBM COBOL for AIX version 3.1 ■ Micro Focus Server Express 5.1

Table 2–3 (Cont.) AIX Operating System Kernel Requirements

Item	Requirement
Oracle Messaging Gateway	IBM WebSphere MQ V6.0.2.0, client and server: mqm.Client.Bnd mqm.Server.Bnd

Verify that the following patches are installed on the system. The procedure following the table describes how to check these requirements

Note: There may be more recent versions of the patches listed installed on your system. If a listed patch is not installed, then determine if you have a more recent patch installed that includes the listed patch before you install the patch version listed.

Table 2–4 AIX APAR and Other Operating System Fixes

Installation Type or Product	Requirement
AIX 6L installations	<p>If you are using a later TL level than the minimum level listed for this release, then contact IBM to determine if the required APARs listed here are included in the TL level that you have on your system. If they are included, then you do not need to install them. If they are not included, then you must install the equivalent APAR for the appropriate TL level.</p> <p>If you are using the minimum operating system TL level for AIX 6L listed above, then install all AIX 6L 6.1 Authorized Problem Analysis Reports (APARs) for AIX 6.1 TL 09 SP1, and the following AIX fixes:</p> <p>IZ41855 IZ51456 IZ52319</p> <p>These 6.1 fixes are already present in the following TL levels:</p> <ul style="list-style-type: none"> ■ AIX 6.1 TL-02 SP-04 and later ■ AIX 6.1 TL-03 SP-02 and later ■ AIX 6.1 TL-04
AIX 5L installations	<p>If you are using a later TL level than the minimum level listed for this release, then check with IBM to determine if the required APARs listed here are included in the TL level that you have on your system. If they are included, then you do not need to install them. If they are not included, then you must install the equivalent APAR for the appropriate TL level.</p> <p>If you are using the minimum operating system TL level for AIX 5L listed above, then install all AIX 5L 5.3 Authorized Problem Analysis Reports (APARs) for AIX 5L v. 5.3 ML06, and the following AIX fixes:</p> <p>IZ42940 IZ49516 IZ52331</p> <p>These fixes are already present in the following TL levels:</p> <ul style="list-style-type: none"> ■ AIX 5.3 TL-09 SP-05 and later ■ AIX 5.3 TL-10 SP-02 and later ■ AIX 5.3 TL-11

Table 2–4 (Cont.) AIX APAR and Other Operating System Fixes

Installation Type or Product	Requirement
Oracle JDBC/OCI Drivers AIX 5L v5.3	<p>Note: These APARs are required only if you are using the associated JDK version.</p> <p>APAR required for JDK 1.4.2 (64-bit):</p> <ul style="list-style-type: none"> ■ IY63533: JDK 1.4.2 64-bit SR1 caix64142-20040917
Oracle ODBC Drivers	<p>You do not require ODBC drivers for Oracle Clusterware or Oracle Database.</p> <p>To use ODBC with Oracle Database, you must also install gcc 3.4.5 or later.</p> <p>Oracle ODBC driver on AIX is certified with ODBC Driver Manager 2.2.12. You can download and install the Driver Manager from the following link:</p> <p>http://www.unixodbc.org</p>

Checking the Software Requirements

To ensure that the system meets these requirements:

1. To determine the version of AIX installed, enter the following command:

```
# oslevel -r
```

If the operating system version is lower than AIX 5.3, then upgrade your operating system to at least this maintenance level. AIX 5L version 5.3 maintenance packages are available from the following Web site:

<http://www-933.ibm.com/support/fixcentral/>

2. To determine whether the required filesets are installed and committed, enter a command similar to the following:

```
# lsllpp -l bos.adt.base bos.adt.lib bos.adt.libm bos.perf.libperfstat \
bos.perf.perfstat bos.perf.proctools rsct.basic.rte rsct.compat.clients.rte \
xlC.aix61.rte 10.1.0.0
```

If a fileset is not installed and committed, then install it. Refer to your operating system or software documentation for information about installing filesets.

To ensure that the system meets these requirements:

1. To determine if required APARs are installed, enter a command similar to the following:

```
# instfix -i -k "IZ41855 IZ51456 IZ52319"
```

If an APAR is not installed, then download it from the following Web site and install it:

<http://www-933.ibm.com/support/fixcentral/>

2. To determine whether a PTF is installed, enter a command similar to the following:

```
# lsllpp -l -B U489726 U485561 ...
```

If a PTF is not installed, then download it from the following Web site and install it:

<http://www-933.ibm.com/support/fixcentral/>

3. If you require a CSD for WebSphere MQ, then refer to the following Web site for download and installation information:

<http://www-01.ibm.com/software/>

Verifying UDP and TCP Kernel Parameters

Use NDD to ensure that the AIX kernel TCP/IP ephemeral port range parameters are set to 9000-65500. This range is larger than the default (32768-65536). For example:

```
# /usr/sbin/no -a | fgrep ephemeral
tcp_ephemeral_low = 32768
tcp_ephemeral_high = 65535
udp_ephemeral_low = 32768
udp_ephemeral_high = 65535
```

In the preceding example, the TCP and UDP ephemeral ports are set to the default range.

If necessary, update the UDP and TCP ephemeral port range to 9000-65500. For example:

```
# /usr/sbin/no -p -o tcp_ephemeral_low=9000 -o tcp_ephemeral_high=65500
# /usr/sbin/no -p -o udp_ephemeral_low=9000 -o udp_ephemeral_high=65500
```

Enabling Core File Creation

During installation, the installer checks the system configuration file that sets core dump preferences to see if core dumps are enabled.

To check if core dumps are enabled, enter the following command:

```
# coreadm
```

Review the output to see if core dumps are disabled. For example:

```
# coreadm
global core file pattern:
global core file content: default
init core file pattern: core
init core file content: default
global core dumps: disabled
per-process core dumps: enabled
global setid core dumps: disabled
per-process setid core dumps: disabled
global core dump logging: disabled
```

In the preceding output example:

- `global core dumps: disabled` indicates that the operating system is not configured to generate global core dump.
- `per-process core dumps: enabled` indicates that the operating system is configured to generate a core dump for each abnormal processes
- `init core file pattern` indicates that the contents of a core dump are gathered from the live process for a process core dump.

To enable core file creation:

1. Set the `ulimit` setting for core dumps to `unlimited`:

```
# ulimit -c unlimited
```

2. Set the ulimit setting for core files to unlimited:

```
# ulimit -f unlimited
```

3. Log in as root, and enter the following commands:

```
# mkdir -p /var/cores
# coreadm -g /var/cores/%f.%n.%p.%t.core -e global -e global-setid -e
# log -d process -d proc-setid
```

In this command:

- `-g` specifies the global core file name pattern. Unless a per-process pattern or setting overrides it, core files are stored in the specified directory with a name in the format *program.node.pid.time.core*. For example:
mytest.myhost.1234.1102010309.core.
- `-e` specifies options to enable. This command example enables the following:
 - Use of a global (that is, system-wide) core file name pattern, ensuring a consistent location of core files
 - Capability of `setuid` programs to also dump core files using the same pattern
 - Generation by the operating system of a syslog message during an attempt to dump core, whether or not this attempt was successful
- `-d` specifies options to disable. This command example disables the following:
 - Core dumps listed using a process core file pattern
 - core dumps listed by process ID of `setuid` programs

In this command example, the command configures the operating system to store all core dumps in a central location with names identifying which process dumped core, and when core was dumped. These changes only affect processes that you start after you run the `coreadm` command. If you want to ensure that you capture core dumps from existing processes, then enter the command `coreadm -u` to apply the settings to all existing processes.

Tuning AIX System Environment

Perform the following system tuning and configuration all cluster nodes.

- [Checking Asynchronous Input Output Processes for AIX 5L](#)
- [Tuning Virtual Memory Manager \(VMM\)](#)
- [Increasing System Block Size Allocation](#)
- [Configuring Shell Limits](#)
- [Configuring User Process Parameters](#)
- [Configuring Network Tuning Parameters](#)

Note: The parameter and shell limit values shown in this section are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

Checking Asynchronous Input Output Processes for AIX 5L

On AIX 5, run the `rootpre.sh` script to enable the Asynchronous Input Output (AIO) device drivers. On AIX 6, the AIO device drivers are enabled by default. For both AIX 5 and AIX 6, increase the number of `aioserver` processes from the default value. The recommended value for `aio_maxreqs` is 64k (65536). Confirm this value for both AIX 5 and AIX 6.

Confirm the `aio_maxreqs` value using the procedure for your release:

AIX 6.1

```
# ioo -o aio_maxreqs
aio_maxreqs = 65536
```

AIX 5.3

```
# lsattr -El aio0 -a maxreqs
maxreqs 65536 Maximum number of REQUESTS True
```

When performing an asynchronous I/O to a file system, note that each asynchronous I/O operation is tied to an asynchronous I/O server. Thus, the number of asynchronous I/O servers limits the number of concurrent asynchronous I/O operations in the system.

The initial number of servers that are started during a system restart is determined by the `minservers` parameter. As concurrent asynchronous I/O operations occur, additional asynchronous I/O servers are started, up to a maximum of the value set in the `maxservers` parameter.

On AIX 5.3, if you are using Oracle Database with data files on a file system, then increase the default values for `minservers` and `maxservers`, as the default values for these parameters are too small. Increase the `minservers` and `maxservers` values based on I/O `kprocs` for each processor.

In general, to set the number of asynchronous I/O servers, complete the following procedure:

1. Adjust the initial value of `maxservers` to 10 times the number of disks that are to be used concurrently but no more than 80.
2. Monitor the performance effects on the system during periods of high I/O activity. If all AIO server processes are started, then increase the `maxservers` value. Also, continue to monitor the system performance during peak I/O activity to determine if there was a benefit from the additional AIO servers. Too many asynchronous I/O servers increase memory and processor overload of additional processes, but this disadvantage is small.

To monitor the number of AIO server processes that have started, enter the following command:

```
# ps -ek|grep -v grep|grep -v posix_aioserver|grep -c aioserver
```

See Also: [Section , "Running the rootpre.sh Script."](#)

Tuning Virtual Memory Manager (VMM)

Oracle recommends that you use the `vmo` command to tune virtual memory using the following values:

Table 2–5 Recommended Values for Virtual Memory Manager

Parameter	Value
<code>minperm%</code>	3 (AIX 5.3 default is 20)
<code>maxperm%</code>	90 (AIX 5.3 default is 80)
<code>maxclient% = 90</code>	90 (AIX 5.3 default is 80)
<code>lru_file_repage</code>	0 (AIX 5.3 default is 1)
<code>strict_maxclient</code>	1 (AIX 5.3 default is 1)
<code>strict_maxperm</code>	0 (AIX 5.3 default is 0)

For example:

```
vmo -p -o minperm%=3
vmo -p -o maxperm%=90
vmo -p -o maxclient%=90
vmo -p -o lru_file_repage=0
vmo -p -o strict_maxclient=1
vmo -p -o strict_maxperm=0
```

You must restart the system for these changes to take effect.

Increasing System Block Size Allocation

Oracle recommends that you increase the space allocated for ARG/ENV list to 128. The size is specified by number of 4K blocks.

For example:

```
/usr/sbin/chdev -l sys0 -a ncargs='128'
```

Configuring Shell Limits

Set shell limits for the Oracle Grid Infrastructure installation owner and for root. Verify that unlimited is set for both accounts either by using the `smit` utility or by editing the `/etc/security/limits` file. The `root` user requires these settings because the `crs` daemon (`crsd`) runs as `root`.

Shell Limit	Item in <code>limits.conf</code>	Hard Limit
Maximum number of open file descriptors	<code>nofile</code>	-1 (unlimited)
Maximum number of processes available to a single user	<code>maxuproc</code>	16384
Maximum size of the stack segment of the process	<code>stack</code>	-1 (unlimited)

To increase the shell limits:

1. Add the following lines to the `/etc/security/limits` file:

```
default:
    fsize = -1
    core = 2097151
    cpu = -1
    data = -1
    rss = -1
    stack = -1
    nofiles = -1
```

2. Enter the following command to list the current setting for the maximum number of process allowed by the Oracle software user:

```
/usr/sbin/lsattr -E -l sys0 -a maxuproc
```

If necessary, change the maxuproc setting using the following command:

```
/usr/sbin/chdev -l sys0 -a maxuproc = 16384
```

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

Caution: se shell programs supported by your operating system vendor. If you use a shell program that is not supported by your operating system, then you can encounter errors during installation.

Configuring User Process Parameters

Verify that the maximum number of processes allowed for each user is set to 2048 or greater:

Note: For production systems, this value should be at least 128 plus the sum of the PROCESSES and PARALLEL_MAX_SERVERS initialization parameters for each database running on the system.

1. Enter the following command:


```
# smit chgsys
```
2. Verify that the value shown for **Maximum number of PROCESSES allowed for each user** is greater than or equal to 2048.

If necessary, edit the existing value.
3. When you have finished making changes, press Enter, then Esc+0 (Exit) to exit.

Configuring Network Tuning Parameters

Verify that the network tuning parameters shown in the following table are set to the values shown or higher values. The procedure following the table describes how to verify and set the values.

Network Tuning Parameter	Recommended Value
ipqmaxlen	512
rfc1323	1
sb_max	2*65536

Network Tuning Parameter	Recommended Value
tcp_recvspace	65536
tcp_sendspace	65536
udp_recvspace	655360 Note: The recommended value of this parameter is 10 times the value of the udp_sendspace parameter. The value must be less than the value of the sb_max parameter.
udp_sendspace	65536 Note: This value is suitable for a default database installation. For production databases, the minimum value for this parameter is 4 KB plus the value of the database DB_BLOCK_SIZE initialization parameter multiplied by the value of the DB_MULTIBLOCK_READ_COUNT initialization parameter: $(DB_BLOCK_SIZE * DB_MULTIBLOCK_READ_COUNT) + 4\text{ KB}$

To view the current value specified for these parameters, and to change them if necessary:

1. To check the current values of the network tuning parameters, enter commands similar to the following:

```
# no -a | more
```

2. If you must change the value of any parameter, then enter the following command to determine whether the system is running in compatibility mode:

```
# lsattr -E -l sys0 -a pre520tune
```

If the system is running in compatibility mode, then the output is similar to the following, showing that the value of the pre520tune attribute is enabled:

```
pre520tune enable Pre-520 tuning compatibility mode True
```

3. If the system is running in compatibility mode, then follow these steps to change the parameter values:

- a. Enter commands similar to the following to change the value of each parameter:

```
# no -o parameter_name=value
```

For example:

```
# no -o udp_recvspace=655360
```

- b. Add entries similar to the following to the `/etc/rc.net` file for each parameter that you changed in the previous step:

```
if [ -f /usr/sbin/no ] ; then
  /usr/sbin/no -o udp_sendspace=65536
  /usr/sbin/no -o udp_recvspace=655360
  /usr/sbin/no -o tcp_sendspace=65536
  /usr/sbin/no -o tcp_recvspace=65536
  /usr/sbin/no -o rfc1323=1
  /usr/sbin/no -o sb_max=2*655360
  /usr/sbin/no -o ipqmaxlen=512
fi
```

By adding these lines to the `/etc/rc.net` file, the values persist when the system restarts.

4. If the system is not running in compatibility mode, then enter commands similar to the following to change the parameter values:

- `ipqmaxlen` parameter:

```
/usr/sbin/no -r -o ipqmaxlen=512
```

- Other parameter:

```
/usr/sbin/no -p -o parameter=value
```

Note: If you modify the `ipqmaxlen` parameter, then you must restart the system.

These commands modify the `/etc/tunables/nextboot` file, causing the attribute values to persist when the system restarts.

Network Time Protocol Setting

Oracle Clusterware requires the same time zone setting on all cluster nodes. During installation, the installation process picks up the time zone setting of the Grid installation owner on the node where OUI runs, and uses that on all nodes as the default TZ setting for all processes managed by Oracle Clusterware. This default is used for databases, Oracle ASM, and any other managed processes.

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

Note: Before starting the installation of Oracle Grid Infrastructure, Oracle recommends that you ensure the clocks on all nodes are set to the same time.

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

To disable the NTP service, run the following command as the `root` user

```
# stopsrc -s xntpd
```

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 stat resource ora.ctssd -t -init
```

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 enable slewing, which prevents time from being adjusted backward. Restart the network time protocol daemon after you complete this task.

To do this on AIX, configure the XNTP daemon to start at each system restart by editing the file `/etc/rc.tcpip`:

1. Open the `/etc/rc.tcpip` file, and locate the following line:

```
start /usr/sbin/xntpd "$src_running"
```

2. Change the line to the following:

```
start /usr/sbin/xntpd "$src_running" "-x"
```

3. Save the file.

To enable XNTP after it has been disabled, enter the following command on each cluster member node:

```
# startsrc -s xntpd -a "-x"
```

Automatic SSH Configuration During Installation

To install Oracle software, Secure Shell (SSH) connectivity should be set up between all cluster member nodes. OUI uses the `ssh` and `scp` commands during installation to run remote commands on and copy files to the other cluster nodes. You must configure SSH so that these commands do not prompt for a password.

Note: SSH is used by Oracle configuration assistants for configuration operations from local to remote nodes. It is also used by Oracle 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: [Section , "Preventing Installation Errors Caused by stty Commands"](#) for information about how to remove `stty` commands in user profiles

By default, OUI searches for SSH public keys in the directory `/usr/local/etc/`, and `ssh-keygen` binaries in `/usr/local/bin`. However, on AIX, SSH public keys typically are located in the path `/etc/ssh`, and `ssh-keygen` binaries are located in the path `/usr/bin`. To ensure that OUI can set up SSH, use the following command to create soft links:

```
# ln -s /etc/ssh /usr/local/etc
# ln -s /usr/bin /usr/local/bin
```

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](#)
- [Environment Requirements for Oracle Database and Oracle ASM Owners](#)
- [Setting Display and X11 Forwarding Configuration](#)
- [Preventing Installation Errors Caused by `stty` Commands](#)

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

Environment Requirements for Oracle Database and Oracle ASM Owners

If you intend to install Oracle Database or Oracle ASM, then complete the following additional tasks. If you plan to install other software using the role-based privileges method, then complete the following tasks for the Oracle Database software owner (`oracle`) and Oracle ASM software owner (`asm`).

- Create an Oracle Base path. The Optimal Flexible Architecture path for the Oracle Base is `/u01/app/user`, where `user` is the name of the user account that you want to own the Oracle Database software. For example: `/u01/app/oracle`.

Note: Do not create the Oracle Clusterware home under Oracle base. Creating an Oracle Clusterware installation in an Oracle base directory path will cause succeeding Oracle installations to fail.

- Set the installation software owner user (`asm, 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 the software owners' environment variable `DISPLAY` environment variables in preparation for the Oracle ASM or Oracle Database installation

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:

- 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 variable is set in the file, then remove the appropriate 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:

- 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 that you want to use to display OUI (your workstation or PC).

12. If you determined that the `/tmp` directory has less than 1 GB MB of free disk 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 -k` 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
```

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 host name 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
```

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.

Running the rootpre.sh Script

Note: Do not run the `rootpre.sh` script if you have a later release of the Oracle Database software already installed on this system.

Run the `rootpre.sh` script:

1. Switch user to root:

```
$ su - root
```

2. Complete one of the following steps, depending on the location of the installation

If the installation files are on disc, enter a command similar to the following, where *directory_path* is the disc mount point directory or the path of the database directory on the DVD:

```
# /directory_path/rootpre.sh
```

If the installation files are on the hard disk, change directory to the Disk1 directory and enter the following command:

```
# ./rootpre.sh
```

3. Exit from the root account:

```
# exit
```

4. Repeat steps 1 through 3 on all nodes of the cluster.

Note: Do not run the `rootpre.sh` script if you have a later release of Oracle Database software already installed on this system.

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 Clusterware binaries. 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.

Oracle recommends that you install Oracle Grid Infrastructure on local homes, rather than using a shared home on shared storage.

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 Oracle Automatic Storage Management (Oracle ASM), and that you must complete before adding an Oracle Real Application Clusters (Oracle RAC) installation to the cluster.

This chapter contains the following topics:

- [Reviewing Oracle Grid Infrastructure Storage Options](#)
- [Shared File System Storage Configuration](#)
- [Oracle Automatic Storage Management Storage Configuration](#)
- [Desupport of Raw Disks](#)

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](#)
- [Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC](#)
- [Supported Storage Options](#)
- [After You Have Selected Disk Storage Options](#)

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

<https://support.oracle.com>

Overview of Oracle Clusterware and Oracle RAC Storage Options

There are two ways of storing Oracle Clusterware files:

- **Oracle 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.

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

- **A supported shared file system:** Supported file systems include the following:
 - **General Parallel File System (GPFS):** A cluster file system for AIX that provides concurrent file access
 - **A supported cluster file system.** Note that if you intend to use a cluster file system for your data files, then you should create partitions large enough for the database files when you create partitions for Oracle Clusterware.

See Also: The Certification page on My Oracle Support for supported cluster file systems

- **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: You can no longer use OUI to install Oracle Clusterware or Oracle Database files on raw disks.

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

Quorum Disk Location Restriction with Existing 9.2 Clusterware Installations

When upgrading your Oracle9i release 9.2 Oracle RAC environment to Oracle Database 11g Release 2 (11.2), you are prompted to specify one or more voting disks during the Oracle Clusterware installation. You must specify a new location for the voting disk in Oracle Database 11g Release 1 (11.1). You cannot reuse the old Oracle9i release 9.2 quorum disk for this purpose.

After You Have Selected Disk Storage Options

When you have determined your disk storage options, you must perform the following tasks in the order listed:

1: Configure shared storage for Oracle Clusterware files

- **To use a file system (local or GPFS) for Oracle Clusterware files,** refer to [Shared File System Storage Configuration](#) on page 3-6

2: Configure storage for Oracle Database files and recovery files

- **To use a file system for database or recovery file storage,** refer to [Section , "Shared File System Storage Configuration"](#) on page 3-6, and ensure that in addition to the volumes you create for Oracle Clusterware files, you also create additional volumes with sizes sufficient to store database files.

See Also: [Section , "Creating Directories for Oracle Database Files on Shared File Systems"](#)

General Information About Oracle ACFS

Oracle Automatic Storage Management Cluster File System (Oracle 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 Oracle ACFS. Note the following about Oracle ACFS:

- Oracle Restart does not support root-based Oracle Clusterware resources. For this reason, the following restrictions apply if you run Oracle ACFS on an Oracle Restart Configuration
 - You must manually load and unload Oracle ACFS drivers.
 - You must manually mount and unmount Oracle ACFS file systems, after the Oracle ASM instance is running
 - You can place Oracle ACFS database home file systems into the Oracle ACFS mount registry, along with other registered Oracle ACFS file systems.
- You cannot put Oracle Clusterware binaries and files on Oracle ACFS.
- You cannot put Oracle Database files on Oracle ACFS.
- You can put Oracle Database binaries and administrative files (for example, trace files) on Oracle ACFS.
- Oracle ACFS provides a general purpose file system for other files.

Note: On AIX, Oracle ACFS has the following installation requirements:

- The AIX version must be AIX 6.1 TL4 SP2, or later updates to AIX 6.1 on PPC64.
 - The system must be running in the RBAC mode, which is the default.
 - The owner of the Oracle Grid Infrastructure installation must be a local user.
-
-

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.

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.

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 disks are supported only when upgrading an existing installation using the disks already configured. On new installations, using disks is not supported by Oracle 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.

Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC

Oracle Grid Infrastructure and Oracle RAC only support cluster-aware volume managers. Volume manager vendors provide cluster solutions that may or may not be cluster-aware. To confirm that a volume manager you want to use is supported, look under the Certifications tab on My Oracle Support to confirm if the vendor's cluster solution is certified for Oracle Grid Infrastructure and Oracle RAC. My Oracle Support is available at the following URL:

<https://support.oracle.com>

Supported Storage Options

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

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
Oracle Automatic Storage Management Note: Loopback devices are not supported for use with Oracle ASM	Yes	No	No	Yes	Yes
Oracle Automatic Storage Management Cluster File System (Oracle ACFS)	No	No	Yes	No	No
General Parallel File System (GPFS) <ul style="list-style-type: none"> ▪ Note: You cannot place ASM files on GPFS. ▪ Oracle does not recommend the use of GPFS for voting disks if HACMP is used. 	Yes	Yes	Yes	Yes	Yes
Local file system	No	Yes	Yes	No	No
NFS file system on a certified NAS filer Note: Requires a certified NAS device. Oracle does not recommend the use of NFS for voting disks if HACMP is used.	Yes	Yes	Yes	Yes	Yes
Shared disk partitions (raw disks), including raw logical volumes managed by HACMP	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) and later 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.

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 [Section , "Shared File System Storage Configuration."](#)
- **To use Oracle Automatic Storage Management**, refer to [Section , "Oracle Automatic Storage Management Storage Configuration."](#)

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](#)
- [Enabling Direct NFS Client Oracle Disk Manager Control of NFS](#)
- [Configuring HACMP Multinode Disk Heartbeat \(MNDHB\) for Oracle Clusterware](#)
- [Configuring Raw Logical Volumes for Oracle Clusterware](#)
- [Configuring Raw Logical Volumes in the New Oracle Clusterware Volume Group](#)
- [Creating a Volume Group for Database Files](#)
- [Creating a Volume Group for Oracle Clusterware](#)
- [Importing the Volume Group on the Other Cluster Nodes](#)
- [Activating the Volume Group in Concurrent Mode on All Cluster Nodes](#)
- [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](#)

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 a cluster file system, it must be a supported cluster file system. Refer to My Oracle Support (<https://support.oracle.com>) for a list of supported cluster file systems.
- To use an NFS file system, it must be on a supported NAS device. Log in to My Oracle Support, and click the Certification tab to find the most current information about supported NAS devices:
<https://support.oracle.com/>
- If you choose to place your Oracle Cluster Registry (OCR) files on a shared file system, then Oracle recommends that 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 on separate disks, 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 on separate physical disks, 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 disk 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 must extend these partitions to at least 300 MB. Oracle recommends that you do not use partitions, but instead place OCR and voting disks in disk groups marked as QUORUM disk groups.

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 300 MB for each voting disk volume.
Oracle Cluster Registry (OCR) with external redundancy	1	At least 300 MB for each OCR volume
Oracle Clusterware files (OCR and voting disks) with redundancy provided by Oracle software.	1	At least 300 MB for each OCR volume At least 300 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.

Deciding to Use a Cluster File System for Oracle Clusterware Files

For new installations, Oracle recommends that you use Oracle Automatic Storage Management (Oracle ASM) to store voting disk and OCR files.

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](#)
- [Specifying Network Paths with the Oranfstab File](#)

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 supported for Oracle RAC. Refer to the following URL for support information:

<https://support.oracle.com>

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 `oralfstab` 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/oralfstab` file. In non-shared Oracle RAC installs, `oralfstab` must be replicated on all nodes.

When the `oralfstab` 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 `oralfstab` file is placed in `/etc`, then you must replicate the file `/etc/oralfstab` file on all nodes, and keep each `/etc/oralfstab` file synchronized on all nodes, just as you must with the `/etc/fstab` file.

See Also: [Section , "Mounting NFS Storage Devices with Direct NFS"](#) 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 (`wtmax`) less than 32768.

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/oralfstab`
2. `/etc/oralfstab`
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 `oralfstab`, then it first verifies kernel NFS mounts by cross-checking entries in `oralfstab` 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 [Section , "Configuring Raw Logical Volumes for Oracle Clusterware."](#) Additionally, an informational message is logged into the Oracle alert and trace files indicating that Direct NFS could not be established.

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

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.

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

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.

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.

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`.

If you are using NFS, then you must set the values for the NFS buffer size parameters `rsize` and `wsize` to 32768. The NFS mount options for binaries are:

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

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

```
cio,rw,bg,hard,intr,rsize=32768,wsize=32768,tcp,noac,vers=3,timeo=600
```

The NFS client-side mount options for Oracle Database datafiles are:

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

Update the `/etc/filesystems` file on each node with entries similar to the following:

```
/NFS_mount:
dev = "/vol/gridhome"
vfs = nfs
nodename = /vol/gridhome
mount = true
options = rw,bg,hard,nointr,rsize=32768,wsiz=32768,proto=tcp,vers=3,timeo=600
account = false
/NFS_mount:
dev = "/vol/CWfiles"
vfs = nfs
nodename = /vol/CWfiles
mount = true
options = cio,rw,bg,hard,intr,rsize=32768,wsiz=32768,tcp,noac,vers=3,timeo=600
account = false
/NFS_mount:
dev = "/vol/datafiles"
vfs = nfs
nodename = /u02/app/oracle/data
mount = true
options =
cio,rw,bg,hard,nointr,rsize=32768,wsiz=32768,proto=tcp,noac,vers=3,timeo=600
account = false
```

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

If you want to create a mount point for binaries only, then enter the following line for a binaries mount point:

```
nfs_server:/vol/grid /u01/oracle/grid nfs -yes
rw,bg,hard,nointr,rsize=32768,wsiz=32768,proto=tcp,vers=3,timeo=600
```

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
```

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

<https://support.oracle.com>

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

Enabling Direct NFS Client Oracle Disk Manager Control of NFS

Complete the following procedure to enable Direct NFS:

1. Create an `orantstab` 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 `orantstab`. A single `orantstab` 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. By default, Direct NFS is installed in a disabled state. To enable Direct NFS, complete the following steps on each node. If you use a shared Grid home for the cluster, then complete the following steps in the shared Grid home:

- a. Log in as the Oracle Grid Infrastructure installation owner.
- b. Change directory to *Grid_home/lib*.
- c. Enter the following commands:

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

Configuring HACMP Multinode Disk Heartbeat (MNDHB) for Oracle Clusterware

This section contains the following topics:

See Also: My Oracle Support for additional information about HACMP deployment and HACMP certification

- [Overview of Requirements for Using HACMP with Oracle Clusterware](#)
- [Deploying HACMP and MDNDHB for Oracle Clusterware](#)

- [Upgrading an Existing Oracle Clusterware and HACMP Installation](#)

Overview of Requirements for Using HACMP with Oracle Clusterware

You must define one Multi-node Disk Heartbeat (MNDHB) network for each Oracle Clusterware voting disk. Each MNDHB and voting disk pair must be located on a single hard disk, separate from the other pairs. You must also configure MNDHB so that the node is halted if access is lost to a quorum of the MNDHB networks in the enhanced concurrent volume group.

To reduce the likelihood of a cluster partition, IBM recommends that HACMP is deployed with multiple IP networks and at least one non-IP network. The non-IP networks can be implemented using RS232 or disk heart-beating. For systems using Oracle RAC and HACMP enhanced concurrent resources (enhanced concurrent logical volumes) for database storage, you must configure MNDHB networks.

Install, configure and have HACMP running before installing Oracle Clusterware. For an Oracle RAC configuration, do not use HACMP for IP failovers on the Oracle RAC network interfaces (public, VIP or private). These network interfaces should not be configured to use HACMP IP failover, as Oracle Clusterware manages VIP failovers for Oracle RAC. The RAC network interfaces are bound to individual nodes and RAC instances. Problems can occur with Oracle Clusterware if HACMP reconfigures IP addresses over different interfaces, or fails over addresses across nodes. You only can use HACMP for failover of IP address on Oracle RAC nodes if Oracle RAC does not use these addresses.

Deploying HACMP and MDNDHB for Oracle Clusterware

Complete the following tasks, replacing each term in italics with the appropriate response for your system, or carrying out the action described and entering the appropriate response for your image:

1. Start HACMP.
2. Enter the following command to ensure that the HACMP `clcomdES` daemon is running:

```
# lssrc -s clcomdES
```

If the daemon is not running, then start it using the following command:

```
# startsrc -s clcomdES
```

3. Ensure that your versions of HACMP and AIX meet the system requirements listed in [Section , "Identifying the Software Requirements"](#).
4. Create HACMP cluster and add the Oracle Clusterware nodes. For example:

```
# smitty cm_add_change_show_an_hacmp_cluster.dialog
* Cluster Name [mycluster]
```

5. Create an HACMP cluster node for each Oracle Clusterware node. For example:

```
# smitty cm_add_a_node_to_the_hacmp_cluster_dialog
* Node Name [mycluster_node1]
Communication Path to Node []
```

6. Create HACMP Ethernet heartbeat networks. The HACMP configuration requires network definitions. Select NO for the IP address takeover for these networks, since they are used by Oracle Clusterware.

Create at least two network definitions: one for the Oracle public interface and a second one for the Oracle private (cluster interconnect) network. Additional Ethernet heartbeat networks can be added if desired.

For example:

```
# smitty cm_add_a_network_to_the_hacmp_cluster_select
- select ether network
* Network Name [my_network_name]
* Network Type ether
* Netmask [my.network.netmask.here]
* Enable IP Address Takeover via IP Aliases [No]
IP Address Offset for Heart beating over IP Aliases []
```

7. For each of the networks added in the previous step, define all of the IP names for each Oracle Clusterware node associated with that network, including the public, private and VIP names for each Oracle Clusterware node. For example:

```
# smitty cm_add_communication_interfaces_devices.select
- select: Add Pre-defined Communication Interfaces and Devices / Communication
Interfaces / desired network
* IP Label/Address [node_ip_address]
* Network Type ether
* Network Name some_network_name
* Node Name [my_node_name]
Network Interface []
```

8. Create an HACMP resource group for the enhanced concurrent volume group resource with the following options:

```
# smitty config_resource_group.dialog.custom
* Resource Group Name [my_resource_group_name]
* Participating Nodes (Default Node Priority) [mynode1,mynode2,mynode3]
Startup Policy Online On All Available Nodes
Failover Policy Bring Offline (On Error Node Only)
Fallback Policy Never Fallback
```

9. Create an AIX enhanced concurrent volume group (Big VG, or Scalable VG) using either the command `smitty mkvg`, or using command lines. The VG must contain at least one hard disk for each voting disk. You must configure at least three voting disks.

In the following example, where you see *default*, accept the default response:

```
# smitty mksvg
VOLUME GROUP name [my_vg_name] PP SIZE in MB
* PHYSICAL VOLUME names [mydisk1,mydisk2,mydisk3]
Force the creation of a volume group? no
Activate volume group AUTOMATICALLY no at system restart?
Volume Group MAJOR NUMBER []
Create VG Concurrent Capable? enhanced concurrent
Max PPs per VG in kilobytes default
Max Logical Volumes default
```

10. Under "Change/Show Resources for a Resource Group (standard)", add the concurrent volume group to the resource group added in the preceding steps.

For example:

```
# smitty cm_change_show_resources_std_resource_group_menu_dmn.select
- select_resource_group_from_step_6
Resource Group Name shared_storage
```

```

Participating Nodes (Default Node Priority) mynode1,mynode2,mynode3
Startup Policy Online On All Available Nodes
Failover Policy Bring Offline (On Error Node Only)
Fallback Policy Never Fallback
Concurrent Volume Groups [enter_VG_from_step_7]
Use forced varyon of volume groups, if necessary false
Application Servers []

```

11. Using the following command, ensure that one MNDHB network is defined for each Oracle Clusterware voting disk. Each MNDHB and voting disk pair must be collocated on a single hard disk, separate from the other pairs. The MNDHB network and Voting Disks exist on shared logical volumes in an enhanced concurrent logical volume managed by HACMP as an enhanced concurrent resource. For each of the hard disks in the VG created in step 6 on which you want to place a voting disk logical volume (LV), create a MNDHB LV.

```

# smitty cl_add_mndhb_lv
- select_resource_group_defined_in_step_6
* Physical Volume name enter F4, then select a hard disk
Logical Volume Name []
Logical Volume Label []
Volume Group name ccvg
Resource Group Name shared_storage
Network Name [n]

```

Note: When you define the LVs for the Oracle Clusterware voting disks, they should be defined on the same disks: one for each disk, as used in this step for the MNDHB LVs.

12. Configure MNDHB so that the node is halted if access is lost to a quorum of the MNDHB networks in the enhanced concurrent volume group. For example:

```

# smitty cl_set_mndhb_response
- select_the_VG_created_in_step_7
On loss of access Halt the node
Optional notification method []
Volume Group ccvg

```

13. Verify and Synchronize HACMP configuration. For example:

```

# smitty cm_initialization_and_standard_config_menu_dmn
- select "Verify and Synchronize HACMP Configuration"

```

Enter Yes if prompted: "Would you like to import shared VG: *ccvg*, in resource group *my_resource_group* onto node: *mynode* to node: *racha702* [Yes / No]:"

14. Add the Add the HACMP cluster node IP names to the file `/usr/es/sbin/cluster/etc/rhosts`.

Upgrading an Existing Oracle Clusterware and HACMP Installation

Complete the following procedure:

1. Back up all databases, and back up the Oracle Cluster Registry (OCR)
2. Shut down on all nodes all Oracle RAC databases, all node applications, and Oracle Clusterware.
3. Enter the following command to disable Oracle Clusterware from starting when nodes are restarted:

```
# crsctl disable crs
```

4. Shut down HACMP on all nodes.
5. Install HACMP APAR IZ01809, following the directions in the README included with that APAR.
6. Determine if the existing voting disk LVs are already on separate hard disks, and if each of these disks have sufficient space (at least 256 MB for the MNDHB LVs. If this is true, then create a MNDHB LV on each of the hard disks. If this is not true, then create new MNDHB LVs and new voting disk LVs, located on separate hard disks using the following command, responding to the sections in italics with the appropriate information for your system:

```
# smitty cl_add_mndhb_lv
- Select_resource_group
* Physical Volume name Enter F4, then select disk for the MNDHB and Voting Disk pair
Logical Volume Name []
Logical Volume Label []
Volume Group name ccvg
Resource Group Name shared_storage
Network Name [net_diskhbmulti_01]
```

7. Verify and Synchronize HACMP configuration.
8. Start HACMP on all nodes.
9. If you added new LVs for voting disks in step 5, then replace each of the existing voting disks with the new ones.
10. Enter the following command to re-enable Oracle Clusterware:

```
# crsctl enable crs
```
11. Start Oracle Clusterware on all nodes, and verify that all resources start correctly.

Configuring Raw Logical Volumes for Oracle Clusterware

Note: To use raw logical volumes for Oracle Clusterware, HACMP must be installed and configured on all cluster nodes.

This section describes how to configure raw logical volumes for Oracle Clusterware and database file storage. The procedures in this section describe how to create a new volume group that contains the logical volumes required for both types of files.

Before you continue, review the following guidelines which contain important information about using volume groups with this release of Oracle RAC:

- You must use concurrent-capable volume groups for Oracle Clusterware.
- The Oracle Clusterware files require less than 560 MB of disk space, with external redundancy. To make efficient use of the disk space in a volume group, Oracle recommends that you use the same volume group for the logical volumes for both the Oracle Clusterware files and the database files.
- If you are upgrading an existing Oracle9*i* release 2 Oracle RAC installation that uses raw logical volumes, then you can use the existing SRVM configuration repository logical volume for the OCR and create a new logical volume in the same volume group for the Oracle Clusterware voting disk. However, you must

remove this volume group from the HACMP concurrent resource group that activates it before you install Oracle Clusterware.

See Also: The HACMP documentation for information about removing a volume group from a concurrent resource group.

Note: If you are upgrading a database, then you must also create a new logical volume for the SYSAUX tablespace. Refer to the ["Configuring Raw Logical Volumes in the New Oracle Clusterware Volume Group"](#) section on page 3-18 for more information about the requirements for the Oracle Clusterware voting disk and SYSAUX logical volumes.

- You must use a HACMP concurrent resource group to activate new or existing volume groups that contain only database files (not Oracle Clusterware files).

See Also: The HACMP documentation for information about adding a volume group to a new or existing concurrent resource group.

- All volume groups that you intend to use for Oracle Clusterware must be activated in concurrent mode before you start the installation.
- The procedures in this section describe how to create basic volumes groups and volumes. If you want to configure more complex volumes, (using mirroring, for example), then use this section in conjunction with the HACMP documentation.

Configuring Raw Logical Volumes in the New Oracle Clusterware Volume Group

To create the required raw logical volumes in the new Oracle Clusterware volume group:

1. Identify the logical volumes that you must create.
2. If you prefer, you can also use the command `smit mklv` to create raw logical volumes.

The following example shows the command used to create a logical volume for the `ocr` volume group in the SYSAUX tablespace with a physical partition size of 114 MB ($1792/7 = 256$):

```
# /usr/sbin/mklv -y test_sysaux_raw_1792m -T 0 -w n -s n -r n ocr 7
```

3. Change the owner, group, and permissions on the character device files associated with the logical volumes that you created, as follows:

Note: The device file associated with the Oracle Cluster Registry must be owned by `root`. All other device files must be owned by the Oracle software owner user (`oracle`).

```
# chown oracle:dba /dev/rora_vote_raw_280m
# chmod 660 /dev/rora_vote_raw_280m
# chown root:oinstall /dev/rora_ocr_raw_280m
# chmod 640 /dev/rora_ocr_raw_280m
```

Creating a Volume Group for Database Files

To create a volume group for the Oracle Database files:

1. If necessary, install the shared disks that you intend to use.
2. To ensure that the disks are available, enter the following command on every node:

```
# /usr/sbin/lsdev -Cc disk
```

The output from this command is similar to the following:

```
hdisk0 Available 1A-09-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 1A-09-00-9,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 17-08-L      SSA Logical Disk Drive
```

3. If a disk is not listed as available on any node, then enter the following command to configure the new disks:

```
# /usr/sbin/cfgmgr
```

4. Enter the following command on any node to identify the device names and any associated volume group for each disk:

```
# /usr/sbin/lspv
```

The output from this command is similar to the following:

```
hdisk0      0000078752249812  rootvg
hdisk1      none                none
hdisk4      00034b6fd4ac1d71  ccvg1
```

For each disk, this command shows:

- The disk device name
- Either the 16 character physical volume identifier (PVID) if the disk has one, or none
- Either the volume group to which the disk belongs, or none

The disks that you want to use may have a PVID, but they must not belong to existing volume groups.

5. If a disk that you want to use for the volume group does not have a PVID, then enter a command similar to the following to assign one to it:

```
# /usr/sbin/chdev -l hdiskn -a pv=yes
```

6. To identify used device major numbers, enter the following command on each node of the cluster:

```
# ls -la /dev | more
```

This command displays information about all configured devices, similar to the following:

```
crw-rw---- 1 root system 45, 0 Jul 19 11:56 vg1
```

In this example, 45 is the major number of the vg1 volume group device.

7. Identify an appropriate major number that is unused on all nodes in the cluster.
8. To create a volume group, enter a command similar to the following, or use SMIT (`smit mkvg`):

```
# /usr/sbin/mkvg -y VGname -B -s PPSize -V majornum -n \  
-C PhysicalVolumes
```

9. The following table describes the options and variables used in this example. Refer to the `mkvg` man page for more information about these options.

Command Option	SMIT Field	Sample Value and Description
-y <i>VGname</i>	VOLUME GROUP name	oracle_vg1 Specify the name for the volume group. The name that you specify could be a generic name, as shown, or it could specify the name of the database that you intend to create.
-y <i>VGname</i>	VOLUME GROUP name	oracle_vg1 Specify the name for the volume group. The name that you specify could be a generic name, as shown, or for a database volume group, it could specify the name of the database that you intend to create.
-B	Create a big VG format Volume Group	Specify this option to create a big VG format volume group. Note: If you are using SMIT, then choose yes for this field.
-s <i>PPsize</i>	Physical partition SIZE in megabytes	32 Specify the size of the physical partitions for the database. The sample value shown enables you to include a disk up to 32 GB in size (32 MB * 1016).

Command Option	SMIT Field	Sample Value and Description
<code>-V Majornum</code>	Volume Group MAJOR NUMBER	46 Specify the device major number for the volume group that you identified in Step 7.
<code>-n</code>	Activate volume group AUTOMATICALLY at system restart	Specify this option to prevent the volume group from being activated at system restart. Note: If you are using SMIT, then choose no for this field.
<code>-C</code>	Create VG Concurrent Capable	Specify this option to create a concurrent capable volume group. Note: If you are using SMIT, then choose yes for this field.
<code>PhysicalVolumes</code>	PHYSICAL VOLUME names	hdisk3 hdisk4 Specify the device names of the disks that you want to add to the volume group.

10. Enter a command similar to the following to vary on the volume group that you created:

```
# /usr/sbin/varyonvg VGname
```

Creating a Volume Group for Oracle Clusterware

To create a volume group for the Oracle Clusterware files:

1. If necessary, install the shared disks that you intend to use.
2. To ensure that the disks are available, enter the following command on every node:

```
# /usr/sbin/lsdev -Cc disk
```

The output from this command is similar to the following:

```
hdisk0 Available 1A-09-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 1A-09-00-9,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 17-08-L SSA Logical Disk Drive
```

3. If a disk is not listed as available on any node, then enter the following command to configure the new disks:

```
# /usr/sbin/cfgmgr
```

4. Enter the following command on any node to identify the device names and any associated volume group for each disk:

```
# /usr/sbin/lspv
```

The output from this command is similar to the following:

```
hdisk0 0000078752249812 rootvg
hdisk1 none none
```

```
hdisk4      00034b6fd4ac1d71  ccvg1
```

For each disk, this command shows:

- The disk device name
- Either the 16 character physical volume identifier (PVID) if the disk has one, or none
- Either the volume group to which the disk belongs, or none

The disks that you want to use may have a PVID, but they must not belong to existing volume groups.

5. If a disk that you want to use for the volume group does not have a PVID, then enter a command similar to the following to assign one to it:

```
# /usr/sbin/chdev -l hdiskn -a pv=yes
```

6. To identify used device major numbers, enter the following command on each node of the cluster:

```
# ls -la /dev | more
```

This command displays information about all configured devices, similar to the following:

```
crw-rw----  1 root      system    45,  0 Jul 19 11:56 vg1
```

In this example, 45 is the major number of the vg1 volume group device.

7. Identify an appropriate major number that is unused on all nodes in the cluster.
8. To create a volume group, enter a command similar to the following, or use SMIT (smit mkvg):

```
# /usr/sbin/mkvg -y VGname -B -s PPSize -V majornum -n \  
-C PhysicalVolumes
```

9. The following table describes the options and variables used in this example. Refer to the mkvg man page for more information about these options.

Command Option	SMIT Field	Sample Value and Description
-y <i>VGname</i>	VOLUME GROUP name	oracle_vg1 Specify the name for the volume group. The name that you specify could be a generic name, as shown, or it could specify the name of the database that you intend to create.
-y <i>VGname</i>	VOLUME GROUP name	oracle_vg1 Specify the name for the volume group. The name that you specify could be a generic name, as shown, or for a database volume group, it could specify the name of the database that you intend to create.
-B	Create a big VG format Volume Group	Specify this option to create a big VG format volume group. Note: If you are using SMIT, then choose yes for this field.

Command Option	SMIT Field	Sample Value and Description
-s <i>PPsize</i>	Physical partition SIZE in megabytes	32 Specify the size of the physical partitions for the database. The sample value shown enables you to include a disk up to 32 GB in size (32 MB * 1016).
-V <i>Majornum</i>	Volume Group MAJOR NUMBER	46 Specify the device major number for the volume group that you identified in Step 7.
-n	Activate volume group AUTOMATICALLY at system restart	Specify this option to prevent the volume group from being activated at system restart. Note: If you are using SMIT, then choose no for this field.
-C	Create VG Concurrent Capable	Specify this option to create a concurrent capable volume group. Note: If you are using SMIT, then choose yes for this field.
<i>PhysicalVolumes</i>	PHYSICAL VOLUME names	hdisk3 hdisk4 Specify the device names of the disks that you want to add to the volume group.

10. Enter a command similar to the following to vary on the volume group that you created:

```
# /usr/sbin/varyonvg VGname
```

Importing the Volume Group on the Other Cluster Nodes

To make the volume group available to all nodes in the cluster, you must import it on each node, as follows:

1. Because the physical volume names may be different on the other nodes, enter the following command to determine the PVID of the physical volumes used by the volume group:

```
# /usr/sbin/lspv
```

2. Note the PVIDs of the physical devices used by the volume group.
3. To vary off the volume group that you want to use, enter a command similar to the following on the node where you created it:

```
# /usr/sbin/varyoffvg VGname
```

4. On each cluster node, complete the following steps:
 - a. Enter the following command to determine the physical volume names associated with the PVIDs you noted previously:

```
# /usr/sbin/lspv
```

- b. On each node of the cluster, enter commands similar to the following to import the volume group definitions:

```
# /usr/sbin/importvg -y VGname -V MajorNumber PhysicalVolume
```

In this example, *MajorNumber* is the device major number for the volume group and *PhysicalVolume* is the name of one of the physical volumes in the volume group.

For example, to import the definition of the `oracle_vg1` volume group with device major number 45 on the `hdisk3` and `hdisk4` physical volumes, enter the following command:

```
# /usr/sbin/importvg -y oracle_vg1 -V 45 hdisk3
```

- c. Change the owner, group, and permissions on the character device files associated with the logical volumes you created, as follows:

```
# chown oracle:dba /dev/rora_vote_raw_280m
# chmod 660 /dev/rora_vote_raw_280m
# chown root:oinstall /dev/rora_ocr_raw_280m
# chmod 640 /dev/rora_ocr_raw_280m
```

- d. Enter the following command to ensure that the volume group will not be activated by the operating system when the node starts:

```
# /usr/sbin/chvg -a n VGname
```

Activating the Volume Group in Concurrent Mode on All Cluster Nodes

To activate the volume group in concurrent mode on all cluster nodes, enter the following command on each node:

```
# /usr/sbin/varyonvg -c VGname
```

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 NFS or GPFS storage, you must complete this procedure only if you want to place the Oracle Clusterware files on a separate file system to 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 that you want 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. Make sure that the file systems are configured to mount automatically when a node restarts.

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

File Type	File System Requirements
Oracle Clusterware files	Choose a file system with at least 600 MB of free disk space (one OCR and one voting disk, with external redundancy).
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 more than one type of file, 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 has permissions to create directories on the storage location where you plan to install Oracle Clusterware, then OUI creates the Oracle Clusterware file directory.
 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.

Note: If you are using GPFS, then you can use the same file system for all purposes. This includes using it for the Oracle home directory and for storing data files and logs. For optimal performance, you should use a large GPFS block size (typically, at least 512 KB). GPFS is designed for scalability, and there is no requirement to create multiple GPFS file systems as long as the amount of data fits in a single GPFS file system.

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 subdirectories in each of the mount point directories, and set the appropriate owner, group, and permissions, you have completed GPFS configuration.

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 -k` 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 NFS configuration for Oracle Database shared storage.

Disabling Direct NFS Client Oracle Disk Management Control of NFS

Complete the following steps to disable the Direct NFS client:

1. Remove the `oranfstab` file.
2. Log in as the Oracle Grid Infrastructure installation owner, and disable the Direct NFS client using the following commands, where `Grid_home` is the path to the Oracle Grid Infrastructure home:

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

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

3. 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.

Oracle Automatic Storage Management Storage Configuration

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

- [Configuring Storage for Oracle Automatic Storage Management](#)
- [Using an Existing Oracle ASM Disk Group](#)
- [Configuring Disk Devices for Oracle ASM](#)

- [Using Disk Groups with Oracle Database Files on Oracle ASM](#)
- [Upgrading Existing Oracle ASM Instances](#)

Configuring Storage for Oracle Automatic Storage Management

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

- [Identifying Storage Requirements for Oracle ASM](#)
- [Creating Files on a NAS Device for Use with Oracle ASM](#)

Identifying Storage Requirements for Oracle ASM

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

1. Determine whether you want to use Oracle ASM for Oracle Clusterware files (OCR and voting 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 Oracle ASM for the other.

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

If you enable automated backups during the installation, then you can select Oracle ASM as the storage mechanism for recovery files by specifying an Oracle ASM 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, then you can decide whether you want to use the same Oracle ASM 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 Oracle ASM disk group for database files and recovery files.
2. Choose the Oracle ASM redundancy level to use for the Oracle ASM disk group.

The redundancy level that you choose for the Oracle ASM disk group determines how Oracle ASM 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 Oracle ASM does not mirror data in an external redundancy disk group, Oracle recommends that you use external redundancy with storage

devices such as RAID, or other similar devices that provide their own data protection mechanisms.

- Normal redundancy

In a normal redundancy disk group, to increase performance and reliability, Oracle ASM by default uses two-way mirroring. A normal redundancy disk group requires a minimum of two disk devices (or two failure groups). The effective disk space in a normal redundancy disk group is half the sum of the disk space in all of its devices.

For Oracle Clusterware files, 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, Oracle ASM uses three-way mirroring to increase performance and provide the highest level of reliability. A high redundancy disk group requires a minimum of three disk devices (or three failure groups). The effective disk space in a high redundancy disk group is one-third the sum of the disk space in all of its devices.

For Oracle Clusterware files, 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 3–4](#) and [Table 3–5](#) 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 3–4 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	300 MB	600 MB	900 MB
Normal	3	600 MB	900 MB	1.5 GB ¹
High	5	900 MB	1.5 GB	2.4 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 disk group 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 disk group with at least 2 GB of available space.

A quorum failure group is a special type of failure group and disks in these failure groups do not contain user data. A quorum failure group is not considered when determining redundancy requirements in respect to storing user data. However, a quorum failure group counts when mounting a disk group.

Table 3–5 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

- For Oracle Clusterware installations, you must also add additional disk space for the Oracle ASM metadata. You can use the following formula to calculate the additional disk space requirements (in MB) for OCR and voting disk files, and the Oracle ASM metadata:

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

Where:

- redundancy = Number of mirrors: external = 1, normal = 2, high = 3.
- ausize = Metadata AU size in megabytes.
- 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.

- Optionally, identify failure groups for the Oracle ASM disk group devices.

If you intend to use a normal or high redundancy disk group, then you can further protect your database against hardware failure by associating a set of disk devices in a custom failure group. By default, each device comprises its own failure group. However, if two disk devices in a normal redundancy disk group are attached to

the same SCSI controller, then the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To protect against failures of this type, you could use two SCSI controllers, each with two disks, and define a failure group for the disks attached to each controller. This configuration would enable the disk group to tolerate the failure of one SCSI controller.

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

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

For failure groups containing database files and clusterware files, including voting 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.

6. If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group. Use the following guidelines when identifying appropriate disk devices:
 - All of the devices in an Oracle ASM disk group should be the same size and have the same performance characteristics.
 - Do not specify multiple partitions on a single physical disk as a disk group device. Oracle ASM expects each disk group device to be on a separate physical disk.
 - Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing Oracle ASM from optimizing I/O across the physical devices. They are not supported with Oracle RAC.

Creating Files on a NAS Device for Use with Oracle ASM

If you have a certified NAS storage device, then you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an Oracle ASM disk group.

To create these files, follow these steps:

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

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

2. Switch user to `root`.
3. Create a mount point directory on the local system. For example:

```
# mkdir -p /mnt/asm
```

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/vfstab`.

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

<https://support.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 [Section , "Configuring Storage NFS Mount and Buffer Size Parameters"](#).

5. Enter a command similar to the following to mount the NFS file system on the local system:

```
# mount /mnt/asm
```

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/asm/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/asm/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/asm
# chmod -R 660 /mnt/asm
```

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

```
/mnt/asm/sales1/
```

Note: During installation, disk paths mounted on Oracle ASM are listed as default database storage candidate disks.

Using an Existing Oracle ASM Disk Group

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

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

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

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

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

To determine if an existing Oracle ASM disk group exists, or to determine if there is sufficient disk space in a disk group, you can use the Oracle 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 Oracle ASM instance is configured on the system:

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

If an Oracle ASM 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 Oracle ASM 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 Oracle ASM instance begins with a plus sign.

2. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the Oracle ASM instance.
3. Connect to the Oracle ASM 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.

Configuring Disk Devices for Oracle ASM

You can configure raw disks for use as Oracle Automatic Storage Management (Oracle ASM) disk groups. To use Oracle ASM with raw disks, you must create sufficient partitions for your data files, and then bind the partitions to raw disks. Make a list of the raw disk names you create for the data files, and have the list available during database installation.

In the following procedure, you are directed to set physical volume IDs (PVIDs) for raw disks. Oracle recommends that you complete the entire procedure, even if you are certain that you do not have PVIDs configured on your system, to prevent the possibility of configuration issues.

Note: If you intend to use Hitachi HDLM (`dm1 f` devices) for storage, then ASM instances do not automatically discover the physical disks, but instead discover only the logical volume manager (LVM) disks. This is because the physical disks can only be opened by programs running as root.

Physical disk paths have path names similar to the following:

```
/dev/rdlmfdrv8  
/dev/rdlmfdrv9
```

Use the following procedure to configure disks:

1. If necessary, install the disks that you intend to use for the disk group and restart the system.
2. Identify or create the disks that you want to include in the Oracle ASM disk group. As the `root` user, enter the following command on any node to identify the device names for the disk devices that you want to use:

```
# /usr/sbin/lspv | grep -i none
```

This command displays information similar to the following for each disk device that is not configured in a volume group:

```
hdisk17          0009005fb9c23648          None
```

In this example, `hdisk17` is the device name of the disk and `0009005fb9c23648` is the physical volume ID (PVID).

3. If a disk device that you want to use does not have a PVID, then enter a command similar to the following to assign one to it, where `n` is the number of the `hdisk`:

```
# chdev -l hdiskn -a pv=yes
```

Note: If you have an existing PVID, then `chdev` overwrites the existing PVID. Be aware that if you have applications depending on the previous PVID, then they will fail.

- On each of the other nodes, enter a command similar to the following to identify the device name associated with each PVID on that node:

```
# /usr/sbin/lspv | grep -i "0009005fb9c23648"
```

The output from this command should be similar to the following:

```
hdisk18          0009005fb9c23648          None
```

In this example, the device name associated with the disk device (hdisk18) is different on this node.

- If the device names are the same on all nodes, then enter commands similar to the following on all nodes to change the owner, group, and permissions on the character raw device files for the disk devices where `grid` is the Oracle Grid Infrastructure installation owner, and `asmadmin` is the OSASM group:

```
# chown grid:asmadmin /dev/rhdiskn
# chmod 660 /dev/rhdiskn
```

- To enable simultaneous access to a disk device from multiple nodes, you must set the appropriate Object Data Manager (ODM) attribute, depending on the type of reserve attribute used by your disks. The following section describes how to perform this task using `hdisk` logical names. Refer to your operating system documentation to find logical device names.

To determine the reserve setting your disks use, enter the following command, where *n* is the `hdisk` device number:

```
# lsattr -E -l hdiskn | grep reserve_
```

The response is either a `reserve_lock` setting, or a `reserve_policy` setting. If the attribute is `reserve_lock`, then ensure that the setting is `reserve_lock = no`. If the attribute is `reserve_policy`, then ensure that the setting is `reserve_policy = no_reserve`.

If necessary, change the setting with the `chdev` command using the following syntax, where *n* is the `hdisk` device number:

```
chdev -l hdiskn -a [ reserve_lock=no | reserve_policy=no_reserve ]
```

For example, to change a setting for the device `hdisk4` from `reserve_lock=yes` to `reserve_lock=no`, enter the following command:

```
# chdev -l hdisk4 -a reserve_lock=no
```

To verify that the setting is correct on all disk devices, enter the following command:

```
# lsattr -El hdiskn | grep reserve
```

- Enter commands similar to the following on any node to clear the PVID from each disk device that you want to use:

```
# /usr/sbin/chdev -l hdiskn -a pv=clear
```

When you are installing Oracle Clusterware, you must enter the paths to the appropriate device files when prompted for the path of the OCR and Oracle Clusterware voting disk. For example:

```
/dev/rhdisk10
```

Using Disk Groups with Oracle Database Files on Oracle ASM

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

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

Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM

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

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

If you intend to use a normal or high redundancy disk group, then you can further protect your database against hardware failure by associating a set of disk devices in a custom failure group. By default, each device comprises its own failure group. However, if two disk devices in a normal redundancy disk group are attached to the same SCSI controller, then the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To protect against failures of this type, you could use two SCSI controllers, each with two disks, and define a failure group for the disks attached to each controller. This configuration would enable the disk group to tolerate the failure of one SCSI controller.

Note: If you define custom failure groups, then you must specify a minimum of two failure groups for normal redundancy and three failure groups for high redundancy.

Creating Disk Groups for Oracle Database Data Files

If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group. Use the following guidelines when identifying appropriate disk devices:

- All of the devices in an Oracle 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. Oracle 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 Oracle Automatic Storage Management disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing Oracle Automatic Storage Management from optimizing I/O across the physical devices. They are not supported with Oracle RAC.

Configuring Oracle Automatic Storage Management Cluster File System

Oracle ACFS is installed as part of an Oracle Grid Infrastructure installation (Oracle Clusterware and Oracle Automatic Storage Management) for 11g release 2 (11.2). You can configure Oracle ACFS for a database home, or use ASMCA to configure ACFS as a general purpose file system.

Note: Oracle ACFS is supported only on AIX AIC 6.1 TL4 SP2 and later updates to AIX 6.1 (on PPC64 only). All other AIX releases supported with Oracle Grid Infrastructure for a Cluster 11g release 2 (11.2) are not supported for Oracle ACFS.

To configure Oracle ACFS for an Oracle Database home for an Oracle RAC database:

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


```
$ cd /u01/app/11.2.0/grid
```
3. Ensure that the Oracle Grid Infrastructure installation owner has read and write permissions on the storage mountpoint you want to use. For example, if you want to use the mountpoint `/u02/acfsmounts/`:


```
$ ls -l /u02/acfsmounts
```
4. Start Oracle ASM Configuration Assistant as the grid installation owner. For example:


```
./asmca
```
5. The Configure ASM: ASM Disk Groups page shows you the Oracle ASM disk group you created during installation. Click the **ASM Cluster File Systems** tab.
6. On the ASM Cluster File Systems page, right-click the Data disk, then select **Create ACFS for Database Home**.
7. In the Create ACFS Hosted Database Home window, enter the following information:
 - **Database Home ADVM Volume Device Name:** Enter the name of the database home. The name must be unique in your enterprise. For example: `dbase_01`
 - **Database Home Mountpoint:** Enter the directory path for the mount point. For example: `/u02/acfsmounts/dbase_01`
Make a note of this mount point for future reference.
 - **Database Home Size (GB):** Enter in gigabytes the size you want the database home to be.
 - **Database Home Owner Name:** Enter the name of the Oracle Database installation owner you plan to use to install the database. For example: `oracle1`
 - **Database Home Owner Group:** Enter the OSDBA group whose members you plan to provide when you install the database. Members of this group are given operating system authentication for the SYSDBA privileges on the database. For example: `dba1`
 - Click **OK** when you have completed your entries.
8. Run the script generated by Oracle ASM Configuration Assistant as a privileged user (`root`). On an Oracle Clusterware environment, the script registers the ACFS as a resource managed by Oracle Clusterware. Registering ACFS as a resource helps Oracle Clusterware to mount the ACFS automatically in proper order when ACFS is used for an Oracle RAC database Home.

9. During Oracle RAC installation, ensure that you or the DBA who installs Oracle RAC selects for the Oracle home the mount point you provided in the **Database Home Mountpoint** field (in the preceding example, `/u02/acfsmounts/dbase_01`).

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

Upgrading Existing Oracle ASM Instances

If you have an Oracle ASM installation from a prior release installed on your server, or in an existing Oracle Clusterware installation, then you can use Oracle Automatic Storage Management Configuration Assistant (ASMCA, located in the path `Grid_home/bin`) to upgrade the existing Oracle ASM instance to 11g release 2 (11.2), and subsequently configure failure groups, Oracle ASM volumes and Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

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

During installation, if you are upgrading from an Oracle ASM release prior to 11.2, and you chose to use Oracle ASM and ASMCA detects that there is a prior Oracle ASM version installed in another Oracle ASM home, then after installing the Oracle ASM 11g release 2 (11.2) binaries, you can start ASMCA to upgrade the existing Oracle ASM instance. You can then choose to configure an Oracle ACFS deployment by creating Oracle ASM volumes and using the upgraded Oracle ASM to create the Oracle ACFS.

If you are upgrading from Oracle ASM 11g release 2 (11.2.0.1) or later, then Oracle ASM is always upgraded with Oracle Grid Infrastructure as part of the rolling upgrade, and ASMCA is started by the root scripts during upgrade. ASMCA cannot perform a separate upgrade of Oracle ASM from release 11.2.0.1 to 11.2.0.2.

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).

Desupport of Raw Disks

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 directly on raw disks 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 a existing raw disks or raw logical volumes, and perform a rolling upgrade of your existing installation. Performing a new installation using raw disks or raw logical volumes 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 Oracle 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](#)

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. The Oracle ASM instances are restarted during the installation.

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 TNSLSD is listening to the same port on which the SCAN listens (default 1521), then the TNSLSD 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 Oracle Automatic Storage Management and Oracle Clusterware Upgrade If You Have Existing Installations**

During the Oracle Grid Infrastructure installation, existing Oracle Clusterware and clustered Oracle ASM installations are both upgraded.

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 `/var/opt/oracle/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. 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.

❑ **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 host name, unless you use a third party cluster software. In that case, the public host name 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 host name 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 host name is considered a new host, and must be added to the cluster. A node under the old name will appear to be down until it is removed from the cluster.

If you use third-party clusterware, then use your vendor documentation to complete setup of your public and private domain addresses.

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 Oracle ASM, 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 disk group. 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 disk group 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.

❑ **Ensure that the Oracle home path you select for the Oracle 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`

- ❑ **Decide if you want to use the Software Updates option.** OUI can install critical patch updates, system requirements updates (hardware, operating system parameters, and kernel packages) for supported operating systems, and other significant updates that can help to ensure your installation proceeds smoothly. Oracle recommends that you enable software updates during installation.

If you choose to enable software updates, then during installation you must provide a valid My Oracle Support user name and password during installation, so that OUI can download the latest updates, or you must provide a path to the location of an software updates packages that you have downloaded previously.

If you plan to run the installation in a secured data center, then you can download updates before starting the installation by starting OUI on a system that has Internet access in update download mode. To start OUI to download updates, enter the following command:

```
$ ./runInstaller -downloadUpdates
```

Provide the My Oracle Support user name and password, and provide proxy settings if needed. After you download updates, transfer the update file to a directory on the server where you plan to run the installation.

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](#)

Running OUI to Install Grid Infrastructure

Complete the following steps to install grid infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) on your cluster. At any time during installation, if you have a question about what you are being asked to do, click the **Help** button on the OUI page.

1. Change to the /Disk1 directory on the installation media, or where you have downloaded the installation binaries, and run 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. If `root.sh` fails on any of the nodes, then you can fix the problem and follow the steps in [Section , "Deconfiguring Oracle Clusterware Without Removing Binaries,"](#) rerun `root.sh` on that node, and continue.

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. Oracle 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 IBM AIX on POWER Systems (64-Bit)*.

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

Installing Grid Infrastructure Using a Cluster Configuration File

During interactive installations of Oracle Grid Infrastructure, you are given the option either of providing cluster configuration information manually, or of using a cluster configuration file. A cluster configuration file is a text file that you can create before starting OUI, which provides OUI with information about the cluster name and node names that it requires to configure the cluster.

During silent installations, you must provide cluster configuration information in the response file only.

Oracle suggests that you consider using a response 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 `grid_install.rsp`.
3. Follow the directions in that section for creating a cluster configuration file.

Installing Grid Infrastructure Using a Software-Only Installation

This section contains the following tasks:

- [Installing the Software Binaries](#)
- [Configuring the Software Binaries](#)
- [Configuring the Software Binaries Using a Response File](#)

Note: Oracle recommends that only advanced users should perform the software-only installation, as this installation option requires manual postinstallation steps to enable the Oracle Grid Infrastructure software.

A software-only installation consists of installing Oracle Grid Infrastructure for a cluster on one node.

If you use the Install Grid Infrastructure Software Only option during installation, then this installs the software binaries on the local node. To complete the installation for your cluster, you must perform the additional steps of configuring Oracle Clusterware and Oracle ASM, creating a clone of the local installation, deploying this clone on other nodes, and then adding the other nodes to the cluster.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about how to clone an Oracle Grid Infrastructure installation to other nodes, and then adding them to the cluster

Installing the Software Binaries

To perform a software-only installation:

1. Run 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. Make note of this information.

However, ignore the instruction to run the `roothas.pl` script, unless you intend to install Oracle Grid Infrastructure on a standalone server (Oracle Restart).

5. 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.
6. Use Oracle Universal Installer as described in steps 1 through 4 to install the Oracle Grid Infrastructure software on every remaining node that you want to include in the cluster, and complete a software-only installation of Oracle Grid Infrastructure on every node.

Configure the cluster using the full OUI configuration wizard GUI as described in [Section , "Configuring the Software Binaries,"](#) or configure the cluster using a response file as described in [section Section , "Configuring the Software Binaries Using a Response File."](#)

Configuring the Software Binaries

Configure the software binaries by starting Oracle Grid Infrastructure configuration wizard in GUI mode:

1. Log in to a terminal as the Grid infrastructure installation owner, and enter the following command:

```
$ ./config.sh
```

The configuration script starts OUI in Configuration Wizard mode. Provide information as needed for configuration. Each page shows the same user interface and performs the same validation checks that OUI normally does. However, instead of running an installation, The configuration wizard mode validates inputs and configures the installation on all cluster nodes.

2. When you complete inputs, OUI shows you the Summary page, listing all inputs you have provided for the cluster. Verify that the summary has the correct

information for your cluster, and click **Install** to start configuration of the local node.

When configuration of the local node is complete, OUI copies the Oracle Grid Infrastructure configuration file to other cluster member nodes.

3. When prompted, run root scripts.
4. When you confirm that all root scripts are run, OUI checks the cluster configuration status, and starts other configuration tools as needed.

Configuring the Software Binaries Using a Response File

When you install or copy Oracle Grid Infrastructure software on any node, you can defer configuration for a later time. This section provides the procedure for completing configuration after the software is installed or copied on nodes, using the configuration wizard utility (`config.sh`).

To configure the Oracle Grid Infrastructure software binaries using a response file:

1. As the Oracle Grid Infrastructure installation owner (`grid`) start OUI in Oracle Grid Infrastructure configuration wizard mode from the Oracle Grid Infrastructure software-only home using the following syntax, where `Grid_home` is the Oracle Grid Infrastructure home, and `filename` is the response file name:

```
Grid_home/crs/config/config.sh [-debug] [-silent -responseFile filename]
```

For example:

```
$ cd /u01/app/grid/crs/config/
$ ./config.sh -responseFile /u01/app/grid/response/response_file.rsp
```

The configuration script starts OUI in Configuration Wizard mode. Each page shows the same user interface and performs the same validation checks that OUI normally does. However, instead of running an installation, The configuration wizard mode validates inputs and configures the installation on all cluster nodes.

2. When you complete inputs, OUI shows you the Summary page, listing all inputs you have provided for the cluster. Verify that the summary has the correct information for your cluster, and click **Install** to start configuration of the local node.

When configuration of the local node is complete, OUI copies the Oracle Grid Infrastructure configuration file to other cluster member nodes.

3. When prompted, run root scripts.
4. When you confirm that all root scripts are run, OUI checks the cluster configuration status, and starts other configuration tools as needed.

Restrictions for Oracle Berkeley DB

The Oracle Berkeley DB embedded database installation that is included with Oracle Grid Infrastructure for a Cluster 11g release 2 (11.2.0.2) is only for use with the Oracle Grid Infrastructure installation products. Refer to terms of the Berkeley DB license at the following URL for details:

<http://www.oracle.com/technetwork/database/berkeleydb/downloads/index.html>

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.

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 is running on node1,node2
```

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

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.

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](#)

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.

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://support.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 E](#) for information about how to stop database processes in preparation for installing patches.

Recommended Postinstallation Tasks

Oracle recommends that you complete the following tasks as needed after installing Oracle Grid Infrastructure:

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

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.

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. Set MAXUPROCS to 2048.
2. Set NCARGS to 128.

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.

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

The Fast Recovery Area is a unified storage location for all Oracle Database files related to recovery. Database administrators can define the `DB_RECOVERY_FILE_DEST` parameter to the path for the Fast Recovery Area to enable on-disk backups, and rapid recovery of data. Enabling rapid backups for recent data can reduce requests to system administrators to retrieve backup tapes for recovery operations.

When you enable Fast Recovery in the `init.ora` file, all RMAN backups, archive logs, control file automatic backups, and database copies are written to the Fast Recovery Area. RMAN automatically manages files in the Fast Recovery Area by deleting obsolete backups and archive files no longer required for recovery.

Oracle recommends that you create a Fast Recovery Area disk group. Oracle Clusterware files and Oracle Database files can be placed on the same disk group, and you can also place Fast recovery files in the same disk group. However, Oracle recommends that you create a separate Fast Recovery disk group to reduce storage device contention.

The Fast Recovery Area is enabled by setting `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*

Creating the Fast Recovery Area Disk Group

To create a Fast recovery file disk group:

1. Navigate to the Grid home bin directory, and start OracleASM 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**.

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) Oracle Grid Infrastructure installations:

- [General Restrictions for Using Older Oracle Database Versions](#)
- [Using ASMCA to Administer Disk Groups for Older 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](#)

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 11g release 2 (11.2).

However, placing Oracle Database homes on Oracle ACFS that are prior to Oracle Database release 11.2 is not supported, because earlier releases are not designed to use Oracle ACFS.

To use Oracle Real Application Clusters release 9.2 with Oracle Clusterware 11g release 2 (11.2), install and run Oracle*9i* Cluster Manager for the release 9.2 database. Oracle*9i* Cluster Manager can coexist with Oracle Clusterware.

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 are upgrading from releases 11.1.0.7, 11.1.0.6, and 10.2.0.4, Oracle recommends that you check for the latest recommended patches for the release you are upgrading from, and install those patches as needed prior to the upgrade.

For more information on recommended patches, refer to "Oracle Upgrade Companion," which is available through Note 785351.1 on My Oracle Support:

<https://support.oracle.com>

You may also refer to Notes 756388.1 and 756671.1 for the current list of recommended patches for each release

Using ASMCA to Administer Disk Groups for Older Database Versions

Use ASM Configuration Assistant (ASMCA) to create and modify disk groups when you install older Oracle databases and Oracle RAC databases on Oracle Grid Infrastructure installations. Starting with 11g release 2, Oracle ASM is installed as part of an Oracle Grid Infrastructure installation, with Oracle Clusterware. You can no longer use Database Configuration Assistant (DBCA) to perform administrative tasks on Oracle ASM.

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

When Oracle Clusterware 11g release 11.2 is installed on a cluster with no previous Oracle software version, it configures Oracle Database 11g release 11.2 and later releases dynamically. However, dynamic configuration does not occur when you install an Oracle Database release 10.x or 11.1 on the cluster. Before installing 10.x or 11.1 Oracle Databases on an Oracle Clusterware 11g release 11.2 cluster you must establish a persistent configuration. Creating a persistent configuration for a node is called pinning a node.

Note: During an upgrade, all cluster member nodes are pinned automatically, and no manual pinning is required for existing databases. This procedure is required only if you install older database versions after installing Oracle Grid Infrastructure release 11.2 software.

To pin a node in preparation for installing or using 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

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
```

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 Oracle Grid Infrastructure home (Grid home). Do not attempt to use the `lsnrctl` commands from Oracle home locations for previous releases, as they cannot be used with the new release.

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 Oracle 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/rdbms/lib/ins_rdbms.mk ipc_rds ioracle
```

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

3. Relock the Grid home and restart the cluster 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](#)
- [Relinking Oracle Grid Infrastructure for a Cluster Binaries](#)
- [Changing the Oracle Grid Infrastructure Home Path](#)
- [Deconfiguring Oracle Clusterware Without Removing Binaries](#)
- [Removing Oracle Clusterware and Oracle ASM](#)

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

Deciding When to Deinstall Oracle Clusterware

Remove installed components in the following situations:

- You have successfully installed Oracle Clusterware, and you want to remove the Oracle Clusterware installation, either in an educational environment, or a test environment.
- You have encountered errors during or after installing or upgrading Oracle Clusterware, and you want to reattempt an installation.
- Your installation or upgrade stopped because of a hardware or operating system failure.
- You are advised by Oracle Support to reinstall Oracle Clusterware.

Migrating From Oracle Restart to Oracle Clusterware

If you have an Oracle Database installation using Oracle Restart (that is, an Oracle Grid Infrastructure installation for a standalone server), and you want to configure that server as a cluster member node, then complete the following tasks:

1. Inspect the Oracle Restart configuration with `srvctl` using the following syntax, where `db_unique_name` is the unique name for the database, and `lsnrname` is the name of the 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. Stop all of the databases, services, and listeners that you discovered in step 1.
4. If present, unmount all Oracle ACFS filesystems.
5. Deconfigure and deinstall the Oracle Grid Infrastructure installation for a standalone server, using the following command:

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

6. Prepare the server for Oracle Clusterware configuration, as described in this document.
7. As the Oracle Grid Infrastructure installation owner, run Oracle Clusterware Configuration Wizard, and save and stage the response file. For example:

```
$ Grid_home/crs/config/config.sh -silent -responseFile $HOME/GI.rsp
```

8. Run `root.sh` for the Oracle Clusterware Configuration Wizard.
9. Mount the Oracle Restart disk group
10. Enter the `volenable` command to enable all Oracle Restart disk group volumes.
11. Mount all Oracle ACFS filesystems manually.
12. Add back Oracle Clusterware services to the Oracle Clusterware home, using the information you wrote down in step 1, including adding back Oracle ACFS resources. For example:

```
/u01/app/grid/product/11.2.0/grid/bin/srvctl add filesystem -d /dev/asm/db1 -g  
ORestartData -v db1 -m /u01/app/grid/product/11.2.0/db1 -u grid
```

13. 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`, and the service `myservice`, enter the following commands:

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

14. Add each service to the database, using the command `srvctl add service`. For example:

```
srvctl add service -d mydb -s myservice
```

Relinking Oracle Grid Infrastructure for a Cluster Binaries

After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to modify the binaries, then use the

following procedure, where *Grid_home* is the Oracle Grid Infrastructure for a Cluster home:

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

As root:

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

As the Oracle Grid Infrastructure for a Cluster owner:

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

As root again:

```
# cd Grid_home/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.

For upgrades from previous releases, if you want to deinstall the prior release Grid home, then you must first unlock the prior release Grid home. Unlock the previous release Grid home by running the command `rootcrs.pl -unlock` from the previous release home. After the script has completed, you can run the deinstall command.

Changing the Oracle Grid Infrastructure Home Path

After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to change the Grid home path, then use the following example as a guide to detach the existing Grid home, and to attach a new Grid home:

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

1. Detach the existing Grid home by running the following command as the Oracle Grid Infrastructure installation owner (`grid`), where `/u01/app/11.2.0/grid` is the existing Grid home location:

```
$ cd /u01/app/11.2.0/grid/oui/bin
$ ./detachhome.sh -silent -local -invPtrLoc /u01/app/11.2.0/grid/oraInst.loc
```

2. As root, move the Grid binaries from the old Grid home location to the new Grid home location. For example, where the old Grid home is `/u01/app/11.2.0/grid` and the new Grid home is `/u01/app/grid`:

```
# mv /u01/app/11.2.0/grid /u01/app/grid
```

3. Attach the new Grid home by running the following command on each cluster member node as the Oracle Grid Infrastructure installation owner (`grid`), where:

- `ORACLE_HOME (/u01/app/grid)` is the new Grid home location
- `ORACLE_HOME_NAME (grid)` is the Grid home name
- `CLUSTER_NODES (node1,node2,node3)` is a comma-delimited list of cluster member nodes
- `LOCAL_NODE (node1)` is the name of the node on which you are running the `attachhome.sh` script

```
$ cd /u01/app/grid/oui/bin
$ ./attachhome.sh -silent -local -invPtrLoc /u01/app/grid/oraInst.loc\
ORACLE_HOME=/u01/app/grid ORACLE_HOME_NAME=grid
"CLUSTER\NODES={node1,node2,node3} Local_Node=node1
```

If the Grid home that you are attaching is a shared Grid home location, then add the flag `-cfs` to the `attachhome.sh` command line.

4. As root again, enter the following command:

```
# cd /u01/app/grid/crs/install
# perl rootcrs.pl -patch
```

5. Repeat steps 1 through 4 on each cluster member node.

You must relink the Oracle Clusterware and Oracle ASM binaries every time you move the Grid home.

Deconfiguring Oracle Clusterware Without Removing Binaries

Running the `rootcrs.pl` command flags `-deconfig -force` enables you to deconfigure Oracle Clusterware on one or more nodes without removing installed binaries. This feature is useful if you encounter an error on one or more cluster nodes during installation when running the `root.sh` command, such as a missing operating system package on one node. By running `rootcrs.pl -deconfig -force` on nodes where you encounter an installation error, you can deconfigure Oracle Clusterware on those nodes, correct the cause of the error, and then run `root.sh` again.

Note: Stop any databases, services, and listeners that may be installed and running before deconfiguring Oracle Clusterware.

Caution: Commands used in this section remove the Oracle Grid infrastructure installation for the entire cluster. If you want to remove the installation from an individual node, then refer to *Oracle Clusterware Administration and Deployment Guide*.

To deconfigure Oracle Clusterware:

1. Log in as the `root` user on a node where you encountered an error.
2. Change directory to `Grid_home/crs/install`. For example:

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


3. Run `rootcrs.pl` with the `-deconfig -force` flags. For example:

```
# perl rootcrs.pl -deconfig -force
```

Repeat on other nodes as required.

4. If you are deconfiguring Oracle Clusterware on all nodes in the cluster, then on the last node, enter the following command:

```
# perl rootcrs.pl -deconfig -force -lastnode
```

The `-lastnode` flag completes deconfiguration of the cluster, including the OCR and voting disks.

Removing Oracle Clusterware and Oracle ASM

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

- [About the Deinstallation Tool](#)
- [Downloading The Deinstall Tool for Use with Failed Installations](#)
- [Deinstall Command Example for Oracle Clusterware and Oracle ASM](#)
- [Deinstallation Parameter File Example for Grid Infrastructure for a Cluster](#)

About the Deinstallation Tool

The Deinstallation Tool (`deinstall`) stops Oracle software, and removes Oracle software and configuration files on the operating system. It 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 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] -h
```

The default method for running the `deinstall` tool is from the `deinstall` directory in the Grid home. For example:

```
$ cd /u01/app/11.2.0/grid/deinstall
$ ./deinstall
```

In addition, you can run the `deinstall` tool from other locations, or with a parameter file, or select other options to run the tool.

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.

If you run `deinstall` from the `$ORACLE_HOME/deinstall` path, then the `-home` flag is not required because the tool knows from which home it is being run. If you use the standalone version of the tool, then `-home` is mandatory.

- `-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 `deinstall` 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 `deinstall` with this flag, it deconfigures and deinstalls the Oracle software on the local node (the node where `deinstall` 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 `deinstall` 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. If you are running the `deinstall` command from the Oracle home that you plan to deinstall, then you do not need to use the `-paramfile` flag.

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

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

 - 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`.

- -h
Use the help option (-h) to obtain additional information about the command option flags.

Downloading The Deinstall Tool for Use with Failed Installations

You can use the Deinstallation tool (`deinstall`) to remove failed or incomplete installations. It is available as a separate download from the Oracle Technology Network (OTN) Web site.

To download the Deinstallation tool:

1. Go to the following URL:

<http://www.oracle.com/technetwork/database/enterprise-edition/downloads/index.html>

2. Under **Oracle Database 11g Release 2**, click **See All** for the respective platform for which you want to download the Deinstallation Tool.

The Deinstallation tool is available for download at the end of this page.

Deinstall Command Example for Oracle Clusterware and Oracle 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, enter the following command:

```
$ cd /u01/app/11.2.0/grid/deinstall/
$ ./deinstall
```

You can run generate a `deinstall` 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.

Deinstallation Parameter File Example 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 Oracle Grid Infrastructure, containing Oracle ASM log files, Oracle Clusterware logs, and other administrative files) is `/u01/app/11.2.0/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/11.2.0/grid/
VIP1_MASK=255.255.252.0
```

```

VOTING_DISKS=/u02/storage/grid/vdsk
SCAN_PORT=1522
silent=true
ASM_UPGRADE=false
ORA_CRS_HOME=/u01/app/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
HELPJAR_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

```

Note: Do not use quote marks with variables except in the following cases:

- Around addresses in CRS_NODEVIPS:


```

CRS_
NODEVIPS='n1-vip/255.255.252.0/eth0,n2-vip/255.255.252.0/eth0'

```
 - Around interface names in NETWORKS:


```

NETWORKS="eth0"/192.0.2.1\:public,"eth1"/10.0.0.1\:cluster_
interconnect VIP1_IP=192.0.2.2

```
-
-

Troubleshooting the Oracle Clusterware Installation Process

This appendix provides troubleshooting information for installing Oracle Clusterware.

See Also: The Oracle Database 11g Oracle RAC documentation set included with the installation media in the Documentation directory:

- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Database Oracle Clusterware and 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](#)
- [Troubleshooting Issues on AIX](#)
- [Performing Cluster Diagnostics During Oracle Clusterware Installations](#)
- [Interconnect Configuration Issues](#)
- [Storage Configuration Issues](#)

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](#)
- [Failed to connect to server, Connection refused by server, or Can't open display](#)
- [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](#)
- [Time stamp is in the future](#)
- [Timed out waiting for the CRS stack to start](#)

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 error file should show the following error: "java.io.IOException: /home/oracle/OraHome//bin/kfod: not found" (My Oracle Support bulletin 276454.1)

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 Oracle 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.

Action: In a local terminal window, log in as the user that started the X Window session, and enter the following command:

```
$ xhost fully_qualified_remote_host_name
```

For example:

```
$ xhost somehost.example.com
```

Then, enter the following commands, where `workstation_name` is the host name or IP address of your workstation.

Bourne, Bash, or Korn shell:

```
$ DISPLAY=workstation_name: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 the X clock appears, then close the X clock and start Oracle Universal Installer again.

Nodes unavailable for selection from the OUI Node Selection screen

Cause: Oracle Clusterware is either not installed, or the Oracle Clusterware services are not up and running.

Action: Install Oracle Clusterware, or review the status of your Oracle Clusterware. 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.
3. As the `oracle` user, attempt to connect to the node with `ssh` or `rsh`. If you are prompted for a password, then user equivalence is not set up properly.

PROT-8: Failed to import data from specified file to the cluster registry

Cause: Insufficient space in an existing Oracle Cluster Registry device partition, which causes a migration failure while running `rootupgrade.sh`. To confirm, look for the error "utopen:12:Not enough space in the backing store" in the log file `Grid_home /log/hostname/client/ocrconfig_pid.log`.

Action: Identify a storage device that has 280 MB or more available space. Locate the existing raw device name from `/var/opt/oracle/srvConfig.loc`, and copy the contents of this raw device to the new device using the command `dd`.

Time stamp is in the future

Cause: One or more nodes has a different clock time than the local node. If this is the case, then you may see output similar to the following:

```
time stamp 2005-04-04 14:49:49 is 106 s in the future
```

Action: Ensure that all member nodes of the cluster have the same clock time.

PRVE-0033 : Core files are not enabled on node

Cause: The core file setting currently prevents the creation of a core file for process aborts and exceptions.

Action: Enable core file creation as described in [Section , "Enabling Core File Creation."](#)

Timed out waiting for the CRS stack to start

Cause: If a configuration issue prevents the Oracle Grid Infrastructure software from installing successfully on all nodes, then you may see error messages such as "Timed out waiting for the CRS stack to start," or you may notice that Oracle Clusterware-managed resources were not create on some nodes after you exit the installer. You also may notice that resources have a status other than ONLINE.

Action: Deconfigure the Oracle Grid Infrastructure installation without removing binaries, and review log files to determine the cause of the configuration issue. After you have fixed the configuration issue, rerun the scripts used during installation to configure Oracle Clusterware.

See Also: [Section , "Deconfiguring Oracle Clusterware Without Removing Binaries"](#)

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

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 [Section , "Setting Display and X11 Forwarding Configuration"](#) to resolve this issue.

If you see errors similar to the following:

```
stty: standard input: Invalid argument
stty: standard input: Invalid argument
```

These errors are produced if hidden files on the system (for example, `.bashrc` or `.cshrc`) contain `stty` commands. If you see these errors, then refer to [Section , "Preventing Installation Errors Caused by stty Commands"](#) to correct the cause of these errors.

Node Reachability Check or Node Connectivity Check Failed

Cause: One or more nodes in the cluster cannot be reached using TCP/IP protocol, through either the public or private interconnects.

Action: Use the command `/bin/ping address` to check each node address. When you find an address that cannot be reached, check your list of public and private addresses to make sure that you have them correctly configured. If you use third-party vendor clusterware, then 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: [Section , "Creating Groups, Users and Paths for Oracle Grid Infrastructure"](#) for instructions about how to create required groups, and how to configure the installation owner user

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.

Oracle 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 Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide*

Troubleshooting Issues on AIX

The following issues can occur on IBM AIX release 6.1:

Oracle Universal Installer error INS-13001, "Environment does not meet minimum requirements" and CLUVFY Reports "Reference data is not available for verifying prerequisites on this operating system distribution"

Cause: The Verified OS is on supported level:

`/bin/oslevel`

6.1.4.0

```
/bin/oslevel -s
6100-04-01-0944
```

This issue is caused by AIX incorrectly reporting the operating system level. In this example, the value returned by `/bin/oslevel` should be `6.1.0.0`.

Action: Install AIX patch IZ64508 to fix the `oslevel` bug.

Performing Cluster Diagnostics During Oracle Clusterware Installations

If Oracle Universal Installer (OUI) does not display the Node Selection page, then perform clusterware diagnostics by running the `olsnodes -v` command from the binary directory in your Oracle Clusterware home (`CRS_home/bin` on Linux and UNIX-based systems, and `CRS_home\BIN` on Windows-based systems) and analyzing its output. Refer to your clusterware documentation if the detailed output indicates that your clusterware is not running.

In addition, 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.

Interconnect Configuration Issues

If you plan to use multiple network interface cards (NICs) for the interconnect, and you do not configure them during installation or after installation with Redundant Interconnect Usage, then you should use a third party solution to aggregate the interfaces at the operating system level. Otherwise, the failure of a single NIC will affect the availability of the cluster node.

If you install Oracle Grid Infrastructure and Oracle RAC, then they must use the same NIC or aggregated NIC cards for the interconnect.

If you use aggregated 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.

Storage Configuration Issues

The following is a list of issues involving storage configuration:

- [Recovery from Losing a Node Filesystem or Grid Home](#)

Recovery from Losing a Node Filesystem or Grid Home

With Oracle Clusterware release 11.2 and later, if you remove a filesystem by mistake, or encounter another storage configuration issue that results in losing the Oracle Local Registry or otherwise corrupting a node, you can recover the node in one of two ways:

1. Restore the node from an operating system level backup (preferred)
2. Remove the node, and then add the node. With 11.2 and later clusters, profile information for is copied to the node, and the node is restored.

The feature that enables cluster nodes to be removed and added again, so that they can be restored from the remaining nodes in the cluster, is called Grid Plug and Play (GPnP). Grid Plug and Play eliminates per-node configuration data and the need for explicit add and delete nodes steps. This allows a system administrator to take a template system image and run it on a new node with no further configuration. This removes many manual operations, reduces the opportunity for errors, and encourages configurations that can be changed easily. Removal of the per-node configuration makes the nodes easier to replace, because they do not need to contain individually-managed state.

Grid Plug and Play reduces the cost of installing, configuring, and managing database nodes by making their per-node state disposable. It allows nodes to be easily replaced with regenerated state.

Initiate recovery of a node using `addnode` syntax similar to the following, where *lostnode* is the node that you are adding back to the cluster:

If you are using Grid Naming Service (GNS):

```
$ ./addNode.sh -silent "CLUSTER_NEW_NODES=lostnode"
```

If you are not using GNS:

```
$ ./addNode.sh -silent "CLUSTER_NEW_NODES={lostnode}" "CLUSTER_NEW_VIRTUAL_HOSTNAMES=lostnode-vip"
```

Note that you require access to root to be able to run the `root.sh` script on the node you restore, to recreate OCR keys and to perform other configuration tasks. When you see prompts to overwrite your existing information in `/usr/local/bin`, accept the default (n):

```
The file "dbhome" already exists in /usr/local/bin. Overwrite it? (y/n) [n]:
The file "oraenv" already exists in /usr/local/bin. Overwrite it? (y/n) [n]:
The file "coraenv" already exists in /usr/local/bin. Overwrite it? (y/n) [n]:
```

Installing and Configuring Oracle Database Using Response Files

This appendix describes how to install and configure Oracle products using response files. It includes information about the following topics:

- [How Response Files Work](#)
- [Preparing a Response File](#)
- [Running the Installer Using a Response File](#)
- [Running Net Configuration Assistant Using a Response File](#)
- [Postinstallation Configuration Using a Response File](#)

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. Prepare a response file.
2. Run the installer in silent or response file mode.
3. If you completed a software-only installation, then run Net Configuration Assistant and Database Configuration Assistant in silent or response file mode.

These steps are described in the following sections.

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>grid_install.rsp</code>	Silent installation of Oracle Grid Infrastructure installations

Caution: When you modify a response file template and save a file for use, the response file may contain plain text passwords. Ownership of the response file should be given to the Oracle software installation owner only, and permissions on the response file should be changed to 600. Oracle strongly recommends that database administrators or other administrators delete or secure response files when they are not in use.

To copy and modify a response file:

1. Copy the response file from the response file directory to a directory on your system:

```
$ cp /directory_path/response/response_file.rsp local_directory
```

In this example, `directory_path` is the path to the `database` directory on the installation media. If you have copied the software to a hard drive, then you can edit the file in the `response` directory 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: Oracle Universal Installer does not record passwords in the response file.

To record a response file:

1. Complete preinstallation tasks as for a normal installation.
2. Ensure that the Oracle Grid Infrastructure software owner user (typically `grid`) has permissions to create or write to the Oracle home path that you will specify when you run the installer.
3. On each installation screen, specify the required information.
4. When the installer displays the Summary screen, perform the following steps:

- a. Click **Save Response File** and specify a file name and location to save the values for the response file, and click **Save**.
 - b. Click **Finish** to create the response file and continue with the installation.
Click **Save Response File** and 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.
5. 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

Run Oracle Universal Installer at the command line, specifying the response file you created. The Oracle Universal Installer executable, `runInstaller`, provides several options. For help information on the full set of these options, run the `runInstaller` command with the `-help` option. For example:

```
$ directory_path/runInstaller -help
```

The help information appears in a window 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.

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 Oracle Automatic Storage Management Configuration Assistant (ASMCA). Provide the following response file:

```
oracle.assistants.asm|S_ASMPASSWORD=password
```

```
oracle.assistants.asm|S_ASMMONITORPASSWORD=password
```

Example B-2 Password response file for Oracle Real Application Clusters

Oracle Database configuration requires configuring a password for the SYS, SYSTEM, SYSMAN, and DBSNMP passwords for use with Database Configuration Assistant (DBCA). In addition, if you use Oracle ASM storage, then configure the ASMSNMP password. 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
```

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 Server Pools](#)
- [Understanding Out-of-Place Upgrade](#)

Understanding Preinstallation Configuration

This section reviews concepts about Oracle 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, you are prompted to provide an `oraInventory` directory path.

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, 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`, so that the central inventory for all installations is outside of the Oracle base for this particular Oracle installation user.

If you neither enter a 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 do not accept this option, and instead use an OFA-compliant path.

For new installations, Oracle recommends that you either create an Oracle path in compliance with OFA structure, such as `/u01/app/oraInventory`, that is owned by an Oracle software owner, or you set the Oracle base environment variable to an OFA-compliant value.

If you set an Oracle base variable to a path such as `/u01/app/grid` or `/u01/app/oracle`, then the Oracle Inventory is defaulted to the path `u01/app/oraInventory` using correct permissions to allow all Oracle installation owners to write to this central inventory directory.

By default, the Oracle Inventory directory is not installed under the Oracle base directory for the installation owner. This is because all Oracle software installations

share a common Oracle Inventory, so there is only one Oracle Inventory for all users, whereas there is a separate Oracle base for each user.

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 you identify as private for the cluster interconnect. If you identify multiple interfaces during information for the private network, then Oracle Clusterware configures them with Private Interconnect Usage. Any interface that you identify as private must be on a subnet that connects to every node of the cluster. Oracle Clusterware uses all the interfaces you identify for use as private interfaces.

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

After installation, if you modify interconnects on Oracle RAC with the `CLUSTER_INTERCONNECTS` initialization parameter, then you must change it to a private IP address, on a subnet that is not used with a public IP address. Oracle does not support changing the interconnect to an interface using a subnet that you have designated as a public subnet.

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

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 host names and IP addresses dynamically as nodes are added and removed from the cluster, without requiring additional host configuration in the DNS.

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

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about Grid Naming Service

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

During installation listeners are created. They listen on the SCAN IP addresses provided on nodes for the SCAN IP addresses. Oracle Net Services routes application requests to the least loaded instance providing the service. Because the SCAN addresses resolve to the cluster, rather than to a node address in the cluster, nodes can be added to or removed from the cluster without affecting the SCAN address configuration.

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

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

Clients configured to use IP addresses for Oracle Database releases 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 `REMOTE_LISTENER` parameter must be set to `SCAN:PORT`. Do not set it to a `TNSNAMES` alias with a single address with the SCAN as `HOST=SCAN`.

The SCAN is optional for most deployments. However, clients using Oracle Database 11g release 2 and later policy-managed databases using server pools should access the database using the SCAN. This is because policy-managed databases can run on different servers at different times, so connecting to a particular node virtual IP address for a policy-managed database is not possible.

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 Oracle Automatic Storage Management Cluster File System](#)
- [About Migrating Existing Oracle ASM Instances](#)
- [About Converting Standalone Oracle ASM Installations to Clustered Installations](#)

Understanding Oracle Automatic Storage Management Cluster File System

Oracle Automatic Storage Management has been extended to include a general purpose file system, called Oracle Automatic Storage Management Cluster File System (Oracle ACFS). Oracle ACFS is a new multi-platform, scalable file system, and storage management technology that extends Oracle Automatic Storage Management (Oracle ASM) functionality to support customer files maintained outside of the Oracle Database. Files supported by Oracle ACFS include application binaries and application reports. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

About Migrating Existing Oracle ASM Instances

If you have an Oracle ASM installation from a prior release installed on your server, or in an existing Oracle Clusterware installation, then you can use Oracle Automatic Storage Management Configuration Assistant (ASMCA, located in the path `Grid_home/bin`) to upgrade the existing Oracle ASM instance to Oracle ASM 11g release 2 (11.2), and subsequently configure failure groups, Oracle ASM volumes, and Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

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

During installation, if you chose to use Oracle ASM and ASMCA detects that there is a prior Oracle ASM version installed in another home, then after installing the Oracle ASM 11g release 2 (11.2) binaries, you can start ASMCA to upgrade the existing Oracle ASM instance. You can then choose to configure an Oracle ACFS deployment by creating Oracle ASM volumes and using the upgraded Oracle ASM to create the Oracle ACFS.

On an existing Oracle Clusterware or Oracle RAC installation, if the 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, disk group names on the cluster-enabled Oracle ASM instances must be different from existing standalone disk group names.

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

Understanding Server Pools

The following section provides a short overview of server pools. It contains the following topics:

- [Overview of Server Pools and Policy-based Management](#)
- [How Server Pools Work](#)
- [The Free Server Pool](#)
- [The Generic Server Pool](#)

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about how to configure and administer server pools

Overview of Server Pools and Policy-based Management

With Oracle Clusterware 11g release 2 (11.2) and later, resources managed by Oracle Clusterware are contained in logical groups of servers called server pools. Resources are hosted on a shared infrastructure and are contained within server pools. The resources are restricted with respect to their hardware resource (such as CPU and memory) consumption by policies, behaving as if they were deployed in a single-system environment.

You can choose to manage resources dynamically using server pools to provide policy-based management of resources in the cluster, or you can choose to manage resources using the traditional method of physically assigning resources to run on particular nodes.

The Oracle Grid Infrastructure installation owner has permissions to create and configure server pools, using SRVCTL, Oracle Enterprise Manager Database Control, or Oracle Database Configuration Assistant (DBCA).

Policy-based management:

- Enables dynamic capacity assignment when needed to provide server capacity in accordance with the priorities you set with policies
- Enables allocation of resources by importance, so that applications obtain the required minimum resources, whenever possible, and so that lower priority applications do not take resources from more important applications

- Ensures isolation where necessary, so that you can provide dedicated servers in a cluster for applications and databases

Applications and databases running in server pools do not share resources. Because of this, server pools isolate resources where necessary, but enable dynamic capacity assignments as required. Together with role-separated management, this capability addresses the needs of organizations that have a standardized cluster environments, but allow multiple administrator groups to share the common cluster infrastructure.

How Server Pools Work

Server pools divide the cluster into groups of servers hosting the same or similar resources. They distribute a uniform workload (a set of Oracle Clusterware resources) over several servers in the cluster. For example, you can restrict Oracle databases to run only in a particular server pool. When you enable role-separated management, you can explicitly grant permission to operating system users to change attributes of certain server pools.

Top-level server pools:

- Logically divide the cluster
- Are always exclusive, meaning that one server can only reside in one particular server pool at a certain point in time

Server pools each have three attributes that they are assigned when they are created:

- **MIN_SIZE**: The minimum number of servers the server pool should contain. If the number of servers in a server pool is below the value of this attribute, then Oracle Clusterware automatically moves servers from elsewhere into the server pool until the number of servers reaches the attribute value.
- **MAX_SIZE**: The maximum number of servers the server pool should contain.
- **IMPORTANCE**: A number from 0 to 1000 (0 being least important) that ranks a server pool among all other server pools in a cluster.

When Oracle Clusterware is installed, two server pools are created automatically: Generic and Free. All servers in a new installation are assigned to the Free server pool, initially. Servers move from Free to newly defined server pools automatically. When you upgrade Oracle Clusterware, all nodes are assigned to the Generic server pool, to ensure compatibility with database releases before Oracle Database 11g release 2 (11.2).

The Free Server Pool

The Free server pool contains servers that are not assigned to any other server pools. The attributes of the Free server pool are restricted, as follows:

- **SERVER_NAMES**, **MIN_SIZE**, and **MAX_SIZE** cannot be edited by the user
- **IMPORTANCE** and **ACL** can be edited by the user

The Generic Server Pool

The Generic server pool stores pre-11g release 2 (11.2) databases and administrator-managed databases that have fixed configurations. Additionally, the Generic server pool contains servers that match either of the following:

- Servers that you specified in the **HOSTING_MEMBERS** attribute of all resources of the application resource type

- Servers with names you specified in the `SERVER_NAMES` attribute of the server pools that list the Generic server pool as a parent server pool

The Generic server pool's attributes are restricted, as follows:

- No one can modify configuration attributes of the Generic server pool (all attributes are read-only)
- When you specify a server name in the `HOSTING_MEMBERS` attribute, Oracle Clusterware only allows it if the server is:
 - Online and exists in the Generic server pool
 - Online and exists in the Free server pool, in which case Oracle Clusterware moves the server into the Generic server pool
 - Online and exists in any other server pool and the client is either a cluster administrator or is allowed to use the server pool's servers, in which case, the server is moved into the Generic server pool
 - Offline and the client is a cluster administrator
- When you register a child server pool with the Generic server pool, Oracle Clusterware only allows it if the server names pass the same requirements as previously specified for the resources.

Servers are initially considered for assignment into the Generic server pool at cluster startup time or when a server is added to the cluster, and only after that to other server pools.

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 E, "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.

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.

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_dsa.pub` keys that you have created, and the file `authorized_keys`.

2. On the local node, use SCP (Secure Copy) or SFTP (Secure FTP) to copy the `authorized_keys` file to the `oracle` user `.ssh` directory on a remote node. The following example is with SCP, on a node called `node2`, with the Oracle Grid Infrastructure owner `grid`, where the `grid` user path is `/home/grid`:

```
[grid@node1 .ssh]$ scp authorized_keys node2:/home/grid/.ssh/
```

You are prompted to accept a DSA key. Enter `Yes`, and you see that the node you are copying to is added to the `known_hosts` file.

When prompted, provide the password for the `grid` user, which should be the same on all nodes in the cluster. The `authorized_keys` file is copied to the remote node.

Your output should be similar to the following, where `xxx` represents parts of a valid IP address:

```
[grid@node1 .ssh]$ scp authorized_keys node2:/home/grid/.ssh/
The authenticity of host 'node2 (xxx.xxx.173.152)' can't be established.
DSA key fingerprint is 7e:60:60:ae:40:40:d1:a6:f7:4e:zz:me:a7:48:ae:f6:7e.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'node1,xxx.xxx.173.152' (dsa) to the list
of known hosts
grid@node2's password:
authorized_keys      100%    828    7.5MB/s    00:00
```

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 host names (alias and fully qualified domain name) of nodes in the cluster to run SSH from the local node to each node, including from the local node to itself, and from each node to each other node:

```
[grid@nodename]$ ssh hostname1 date
[grid@nodename]$ ssh hostname2 date
.
.
.
```

For example:

```
[grid@node1 grid]$ ssh node1 date
The authenticity of host 'node1 (xxx.xxx.100.101)' can't be established.
DSA key fingerprint is 7z:60:60:zz:48:48:z1:a0:f7:4e.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'node1,xxx.xxx.100.101' (DSA) to the list of
known hosts.
Mon Dec 4 11:08:13 PST 2006
[grid@node1 grid]$ ssh node1.example.com date
The authenticity of host 'node1.example.com (xxx.xxx.100.101)' can't be
established.
DSA key fingerprint is 7z:60:60:zz:48:48:z1:a0:f7:4e.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'node1.example.com,xxx.xxx.100.101' (DSA) to the
list of known hosts.
Mon Dec 4 11:08:13 PST 2006
[grid@node1 grid]$ ssh node2 date
Mon Dec 4 11:08:35 PST 2006
.
.
.
```

At the end of this process, the public host name for each member node should be registered in the `known_hosts` file for all other cluster nodes.

If you are using a remote client to connect to the local node, and you see a message similar to "Warning: No xauth data; using fake authentication data for X11 forwarding," then this means that your authorized keys file is configured correctly, but your SSH configuration has X11 forwarding enabled. To correct this issue, proceed to [Section , "Setting Display and X11 Forwarding Configuration."](#)

3. Repeat step 2 on each cluster node member.

If you have configured SSH correctly, then you can now use the `ssh` or `scp` commands without being prompted for a password. For example:

```
[grid@node1 ~]$ ssh node2 date
Mon Feb 26 23:34:42 UTC 2009
[grid@node1 ~]$ ssh node1 date
Mon Feb 26 23:34:48 UTC 2009
```

If any node prompts for a password, then verify that the `~/.ssh/authorized_keys` file on that node contains the correct public keys, and that you have created an Oracle software owner with identical group membership and IDs.

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](#)
- [Unset Oracle Environment Variables](#)
- [About Oracle ASM and Oracle Grid Infrastructure Installation and Upgrade](#)
- [Restrictions for Clusterware and Oracle ASM Upgrades](#)
- [Preparing to Upgrade an Existing Oracle Clusterware Installation](#)
- [Performing Rolling Upgrades From an Earlier Release](#)
- [Updating DB Control and Grid Control Target Parameters](#)
- [Unlocking the Existing Oracle Clusterware Installation](#)
- [Downgrading Oracle Clusterware After an Upgrade](#)
- [Checking Cluster Health Monitor Repository Size After Upgrading](#)

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.

Unset Oracle Environment Variables

Unset Oracle environment variables.

If you have set `ORA_CRS_HOME` as an environment variable, following instructions from Oracle Support, then unset it before starting an installation or upgrade. You should never use `ORA_CRS_HOME` as an environment variable except under explicit direction from Oracle Support.

Check to ensure that installation owner login shell profiles (for example, `.profile` or `.cshrc`) do not have `ORA_CRS_HOME` set.

If you have had an existing installation on your system, and you are using the same user account to install this installation, then unset the following environment variables: `ORA_CRS_HOME`; `ORACLE_HOME`; `ORA_NLS10`; `TNS_ADMIN`; and any other environment variable set for the Oracle installation user that is connected with Oracle software homes.

About Oracle ASM and Oracle Grid Infrastructure Installation and Upgrade

In past releases, Oracle Automatic Storage Management (Oracle ASM) was installed as part of the Oracle Database installation. With Oracle Database 11g release 2 (11.2), Oracle ASM is installed when you install the Oracle Grid Infrastructure components and shares an Oracle home with Oracle Clusterware when installed in a cluster such as with Oracle RAC or with Oracle Restart on a standalone server.

If you have an existing Oracle ASM instance, you can either upgrade it at the time that you install Oracle Grid Infrastructure, or you can upgrade it after the installation, using Oracle ASM Configuration Assistant (ASMCA). However, be aware that a number of Oracle ASM features are disabled until you upgrade Oracle ASM, and Oracle Clusterware management of Oracle ASM does not function correctly until Oracle ASM is upgraded, because Oracle Clusterware only manages Oracle ASM when it is running in the Oracle Grid Infrastructure home. For this reason, Oracle recommends that if you do not upgrade Oracle ASM at the same time as you upgrade Oracle Clusterware, then you should upgrade Oracle ASM immediately afterward.

You can perform out-of-place upgrades to an Oracle ASM instance using Oracle ASM Configuration Assistant (ASMCA). In addition to running ASMCA using the graphic user interface, you can run ASMCA in non-interactive (silent) mode.

In prior releases, you could use Database Upgrade Assistant (DBUA) to upgrade either an Oracle Database, or Oracle ASM. That is no longer the case. You can only use DBUA to upgrade an Oracle Database instance. Use Oracle ASM Configuration Assistant (ASMCA) to upgrade Oracle ASM.

See Also: *Oracle Database Upgrade Guide* and *Oracle Database Storage Administrator's Guide* for additional information about upgrading existing Oracle ASM installations

Restrictions for Clusterware and Oracle ASM Upgrades

Be aware of the following restrictions and changes for upgrades to Oracle Grid Infrastructure installations, which consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM):

- 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, 11.1.0.6, or 11.2.
- If you have Oracle ACFS file systems on Oracle Grid Infrastructure 11g release 2 (11.2.0.1), you upgrade Oracle Grid Infrastructure to 11g release 2 (11.2.0.2), and you take advantage of Redundant Interconnect Usage and add one or more additional private interfaces to the private network, then you must restart the Oracle ASM instance on each upgraded cluster member node.
- 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 11gR2 Upgrade Companion" Note 785351.1 on My Oracle Support:

<https://support.oracle.com>

- To upgrade existing 11.2.0.1 Oracle Grid Infrastructure installations to Oracle Grid Infrastructure 11.2.0.2, you must first do at least one of the following:
 - Patch the release 11.2.0.1 Oracle Grid Infrastructure home with the 9413827 and 9706490 patches.
 - Install Oracle Grid Infrastructure Patch Set 1 (GI PSU1) or Oracle Grid Infrastructure Patch Set 2 (GI PSU2).
- 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 Oracle Grid Infrastructure for a Cluster home for Oracle Clusterware and Oracle ASM 11g release 2 (11.2).
- With Oracle Clusterware 11g release 1 and later releases, the same user that owned the Oracle Clusterware 10g software must perform the Oracle Clusterware 11g upgrade. Before Oracle Database 11g, either all Oracle software installations were owned by the Oracle user, typically `oracle`, or Oracle Database software was owned by `oracle`, and Oracle Clusterware software was owned by a separate user, typically `crs`.
- Oracle ASM and Oracle Clusterware both run in the Oracle Grid Infrastructure home.
- During a major version upgrade to 11g release 2 (11.2), the software in the 11g release 2 (11.2) Oracle 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 Oracle Grid Infrastructure upgrade, use the `srvctl` from the existing database homes.
- During Oracle Clusterware installation, if there is a single instance 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 single instance (non-clustered) Oracle ASM installation is on a remote node, which is a node other than the local node (the node on which the Oracle Grid Infrastructure installation is being performed), then it will remain a single instance Oracle ASM installation. However, during installation, if you select to place the Oracle Cluster Registry (OCR) and voting disk files on Oracle ASM, then a clustered Oracle ASM installation is created on all nodes in the cluster, and the single instance Oracle ASM installation on the remote node will become nonfunctional.

See Also: *Oracle Database Upgrade Guide*

Preparing to Upgrade an Existing Oracle Clusterware Installation

If you have an existing Oracle Clusterware installation, then you upgrade your existing cluster by performing an out-of-place upgrade. You cannot perform an in-place upgrade.

Verifying 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.

See Also: *Oracle Database Upgrade Guide*

Checks to Complete Before 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
```
 3. If you have AIX HACMP vendor clusterware on an Oracle Clusterware release 11.1 (11.1.0.7) cluster, then be prepared to do the following:
 - a. Install the new Clusterware version in a software-only installation.
 - b. Bring down the release 11.1.0.7 Oracle Clusterware stack.
 - c. Run the Oracle Grid Infrastructure release 11.2 rootpre.sh script
 - d. Run the Oracle Grid Infrastructure 11.2 rootupgrade.sh script.

Performing Rolling Upgrades From an Earlier Release

Use the following procedures to upgrade Oracle Clusterware or Oracle Automatic Storage Management:

- [Performing a Rolling Upgrade of Oracle Clusterware](#)
- [Performing a Rolling Upgrade of Oracle Automatic Storage Management](#)

Note: When you upgrade to Oracle Clusterware 11g release 2 (11.2), Oracle 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.

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 `rootupgrade.sh` script on each node, that node's instances are shut down and then started up again by the 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. If you are upgrading from 11.2.0.1 to 11.2.0.2, then all nodes are selected by default. You cannot select or de-select the nodes.

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.
4. If you have an Oracle Clusterware 11g release 11.1 (11.1.0.7) installation running with an AIX HACMP cluster, then bring down the release 11.1.0.7 Oracle Clusterware stack on each node that you want to upgrade in this installation session. If you do not have an HACMP cluster, then proceed to step 2.

5. If you have AIX HACMP vendor clusterware on an Oracle Clusterware release 11.1 (11.1.0.7) cluster, then run the Oracle Grid Infrastructure release 11.2 `rootpre.sh` script on each node of the cluster that you want to upgrade.
6. 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.

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.

7. After running the `rootupgrade.sh` script on the last node in the cluster, if you left the check box with ASMCA marked, as is the default, then ASM Configuration Assistant runs automatically, and the Oracle Clusterware upgrade is complete. If you unclicked the box on the interview stage of the upgrade, then ASMCA is not run automatically.

If an earlier version of Oracle 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.

Performing a Rolling Upgrade of Oracle 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 Oracle 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.

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`

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

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 Oracle Grid Infrastructure for a cluster home, or Grid home), the path for the CRS_HOME parameter in some parameter files must be changed. If you do not change the parameter, then you encounter errors such as "cluster target broken on 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.

Unlocking the Existing Oracle Clusterware Installation

After upgrade from previous releases, if you want to deinstall the previous release Oracle Grid Infrastructure Grid home, then you must first change the permission and ownership of the previous release Grid home. Log in as root, and change the permission and ownership of the previous release Grid home using the following command syntax, where *oldGH* is the previous release Grid home, *swowner* is the Oracle Grid Infrastructure installation owner, and *oldGHParent* is the parent directory of the previous release Grid home:

```
#chmod -R 755 oldGH
#chown -R swowner oldGH
#chown swowner oldGHParent
```

For example:

```
#chmod -R 755 /u01/app/11.1.0.6/grid
#chown -R grid /u01/app/11.1.0.6/grid
#chown grid /u01/app/11.1.0.6
```

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/11.2.0/grid/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/11.2.0/grid/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.

Checking Cluster Health Monitor Repository Size After Upgrading

If you are upgrading from a prior release using IPD/OS to Oracle Grid Infrastructure release 2 (11.2.0.2 and later), then you should review the Cluster Health Monitor repository size (the CHM repository). Oracle recommends that you review your CHM repository needs, and enlarge the repository size if you want to maintain a larger CHM repository.

Note: Your previous IPD/OS repository is deleted when you install Oracle Grid Infrastructure, and you run the `root.sh` script on each node.

By default, the CHM repository size is one Gigabyte (1 GB), regardless of the size of the cluster.

To enlarge the CHM repository, use the following command syntax, where *RETENTION_TIME* is the size of CHM repository in number of seconds:

```
oculumon manage -repos resize RETENTION_TIME
```

The value for *RETENTION_TIME* must be more than 3600 (one hour) and less than 259200 (three days). If you enlarge the CHM repository size, then you must ensure that there is local space available for the repository size you select on each node of the cluster. If there is not sufficient space available, then you can move the repository to shared storage.

For example, to set the repository size to four hours:

```
$ ocumon manage -repos resize 14400
```

Index

A

ACFS. *See* Oracle ACFS.
activating volume groups, 3-24
AIX
 tuning virtual memory manager, 2-37
and asmcmd errors, 2-6
APAR
 download location, 2-33
APAR download location, 2-33
ASM
 and multiple databases, 2-12
 and rolling upgrade, E-6
 characteristics of failure groups, 3-30, 3-36
 creating the OSDBA for ASM group, 2-13
 disk groups, 3-28
 failure groups, 3-28
 examples, 3-30, 3-36
 identifying, 3-30, 3-36
 files not supported on GPFS, 3-5
 number of instances on each node, 1-6, 3-2
 OSASM or ASM administrator, 2-11
 OSDBA for ASM group, 2-12
 recommendations for disk groups, 3-28
 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, 3-29
 space required for preconfigured database, 3-29
 storage option for data files, 3-3
 storing Oracle Clusterware files on, 3-5
ASM group
 creating, 2-13
ASMCA
 Used to create disk groups for older Oracle
 Database releases on Oracle ASM, 5-5
ASMSNMP, B-8
Automatic Storage Management Cluster File System.
 See Oracle ACFS.
Automatic Storage Management. *See* ASM.

B

Bash shell
 default user startup file, 2-43
binaries

 relinking, 5-6
Bourne shell
 default user startup file, 2-43

C

C compiler
 requirement, 2-31
 See also Pro*C/C++
C shell
 default user startup file, 2-43
central inventory, 2-10
 about, C-1
central inventory. *See Also* OINSTALL group, and
 Oracle Inventory group
cfgmgr command, 3-19, 3-21
changing host names, 4-3
chdev command, 3-19, 3-22
checking disk availability for raw disks, 3-19, 3-21
checking maintenance level, 2-33
checking version, 2-33
chmod command, 3-27
chown command, 3-27
clients
 connecting to SCANs, C-4
cloning
 cloning a Grid home to other nodes, 4-8
cluster configuration file, 4-7
cluster file system
 storage option for data files, 3-3
cluster name
 requirements for, 4-4
cluster nodes
 activating volume groups, 3-24
 importing raw disk disk group, 3-23
 private node names, 4-4
 public node names, 4-4
 specifying uids and gids, 2-15
 virtual node names, 4-4
Cluster Time Synchronization Service, 2-40
Cluster Verification Utility
 fixup scripts, 2-2
 user equivalency troubleshooting, A-4
clusterware diagnostics, A-7
commands
 asmca, 3-37, 4-6, 5-4, E-6

- asmcmd, 2-6
- chmod, 3-27
- chown, 3-27
- crsctl, 4-10, 5-5, E-3, E-7
- dd, xix
- df, 1-2, 1-3
- groupadd, 2-16
- id, 2-15
- mkdir, 3-27
- passwd, 2-16
- ping, 2-23
- rootcrs.pl, 5-6
 - and deconfigure option, 6-4
- rootupgrade.sh, E-3
- sqlplus, 2-6
- srvctl, E-3
- umask, 2-42, 2-43
- unset, E-4
- useradd, 2-14, 2-16
- xhost, 2-3
- xterm, 2-4
- configuring new disks, 3-19, 3-21
- creating a volume group, 3-19, 3-21
- creating raw logical volumes, 3-18
- creating volume groups, 3-19, 3-22
- cron jobs, 4-5
- crs_install.rsp file, B-3
- CSD
 - download location for WebSphere MQ, 2-34
- ctsdd, 2-40
- custom database
 - failure groups for ASM, 3-30, 3-36
 - requirements when using ASM, 3-29
- Custom installation type
 - reasons for choosing, 2-11

D

- data files
 - creating separate directories for, 3-24, 3-26
 - setting permissions on data file directories, 3-27
 - storage options, 3-3
- data loss
 - minimizing with ASM, 3-30, 3-36
- database files
 - supported storage options, 3-5
- databases
 - ASM requirements, 3-29
- dba group
 - raw device group, 3-18, 3-24
- dba group. *See* OSDBA group
- DBCA
 - no longer used for Oracle ASM diskgroup administration, 5-5
- dbca.rsp file, B-3
- deconfiguring Oracle Clusterware, 6-4
- default file mode creation mask
 - setting, 2-42, 2-43
- deinstall, 6-1
- deinstallation, 6-1

- desupported
 - raw disks, xvii
- device numbers
 - identifying major numbers, 3-19, 3-22
- df command, 2-44
- diagnostics, A-7
- Direct NFS
 - disabling, 3-27
 - enabling, 3-12, 3-13
 - for data files, 3-8
 - minimum write size value for, 3-9
- directory
 - creating separate data file directories, 3-24, 3-26
 - permission for data file directories, 3-27
- disk group
 - ASM, 3-28
 - recommendations for Oracle ASM disk groups, 3-28
- disk groups
 - recommendations for, 3-28
- disk space
 - checking, 2-21
 - requirements for preconfigured database in ASM, 3-29
- disks
 - checking availability for raw disks, 3-19, 3-21
 - configuring new disks, 3-19, 3-21
 - identifying LVM disks, 3-19, 3-21
- DISPLAY environment variable
 - setting, 2-43

E

- emulator
 - installing from X emulator, 2-4
- enterprise.rsp file, B-3
- environment
 - configuring for oracle user, 2-42
- environment variables
 - DISPLAY, 2-43
 - ORACLE_BASE, C-2
 - ORACLE_HOME, 2-6, E-4
 - ORACLE_SID, E-4
 - removing from shell startup file, 2-43
 - SHELL, 2-43
 - TEMP and TMPDIR, 2-20, 2-44
- errors
 - X11 forwarding, 2-44, D-4
- /etc/security/limits.so file, 2-37
- Exadata
 - relinking binaries example for, 5-6
- examples
 - ASM failure groups, 3-30, 3-36

F

- failure group
 - ASM, 3-28
 - characteristics of ASM failure group, 3-30, 3-36
 - examples of ASM failure groups, 3-30, 3-36

- features, new, xv
- file mode creation mask
 - setting, 2-42, 2-43
- file system
 - storage option for data files, 3-3
- file systems, 3-5
- files
 - .bash_profile, 2-43
 - dbca.rsp, B-3
 - editing shell startup file, 2-43
 - enterprise.rsp, B-3
 - /etc/security/limits.so, 2-37
 - .login, 2-43
 - oraInst.loc, 2-5
 - .profile, 2-43
 - response files, B-3
- filesets, 2-29
 - checking, 2-33
- fixup script, 2-2
 - about, 1-1

G

- Gateway
 - See Oracle Messaging Gateway
- General Parallel File System. See GPFS
- gid
 - identifying existing, 2-15
 - specifying, 2-16
 - specifying on other nodes, 2-15
- globalization
 - support for, 4-3
- GNS
 - about, 2-24
- GPFS, 3-2
 - database file storage on, 3-5
- grid home
 - and Oracle base restriction, 2-7
 - default path for, 2-46
 - disk space for, 2-19
 - unlocking, 5-6
- grid infrastructure owner (grid), 2-10
- grid naming service. See GNS
- grid user, 2-10
- group IDs
 - identifying existing, 2-15
 - specifying, 2-16
 - specifying on other nodes, 2-15
- groups
 - checking for existing OINSTALL group, 2-5
 - creating identical groups on other nodes, 2-15
 - creating the ASM group, 2-13
 - creating the OSDBA for ASM group, 2-13
 - creating the OSDBA group, 2-13
 - OINSTALL, 2-5, 2-6
 - OSASM (asmadmin), 2-11
 - OSDBA (dba), 2-11
 - OSDBA for ASM (asmdba), 2-12
 - OSDBA group (dba), 2-11
 - OSOPER (oper), 2-11

- OSOPER for ASM, 2-12
- OSOPER group (oper), 2-11
 - required for installation owner user, 2-10
 - specifying when creating users, 2-16
 - using NIS, 2-10, 2-15

H

- HACMP, 3-13
 - database storage on, 3-5
 - deploying, 3-14
 - upgrading, 3-16
 - verifying free disk space on, 2-21
- hardware requirements, 2-19
- high availability IP addresses, 2-22
- host names
 - changing, 4-3
 - legal hostnames, 4-4

I

- IBM WebSphere MQ
 - requirement, 2-32
- id command, 2-15
- identifying disks for LVM, 3-19, 3-21
- identifying LVM disks, 3-19, 3-21
- importing raw disk disk groups, 3-23
- importvg command, 3-24
- initializing disks for LVM, 3-19, 3-22
- INS-32026 error, 2-7
- installation
 - and cron jobs, 4-5
 - and globalization, 4-3
 - cloning a Grid infrastructure installation to other nodes, 4-8
 - response files, B-3
 - preparing, B-3, B-4
 - templates, B-3
 - silent mode, B-5
 - using cluster configuration file, 4-7
- installation types
 - and ASM, 3-29
- instfix command, 2-33
- interfaces
 - requirements for private interconnect, C-4
- intermittent hangs
 - and socket files, 4-10, 4-11

J

- JDK requirements, 2-29
- job role separation users, 2-10

K

- Korn shell
 - default user startup file, 2-43

L

- legal hostnames, 4-4

- limits.so file, 2-37
- log file
 - how to access during installation, 4-6
- .login file, 2-43
- lsdev command, 3-19, 3-21
- lslpp command, 2-33
- lspv command, 3-19, 3-21, 3-23
- LVM
 - creating a volume group, 3-19, 3-21
 - creating raw logical volumes, 3-18
 - creating volume groups, 3-19, 3-22
 - identifying available disks, 3-19, 3-21
 - identifying major device numbers, 3-19, 3-22
 - identifying volume group devices, 3-19, 3-21
 - initializing disks, 3-19, 3-22
 - recommendations for ASM, 3-28

M

- maintenance level
 - checking, 2-33
- major device numbers
 - identifying, 3-19, 3-22
- mask
 - setting default file mode creation mask, 2-42, 2-43
- memory requirements, 2-19
- Messaging Gateway
 - See* Oracle Messaging Gateway
- mixed binaries, 2-29
- mkdir command, 3-27
- mkvg command, 3-19, 3-22
- mode
 - setting default file mode creation mask, 2-42, 2-43
- multiple databases
 - and ASM, 2-12
- multiple oracle homes, 2-6, 3-27
- My Oracle Support, 5-1

N

- Net Configuration Assistant (NetCA)
 - response files, B-6
 - running at command prompt, B-6
- netca, 4-6
- netca.rsp file, B-3
- Network Information Services
 - See* NIS
- new features, xv
- NFS, 3-5, 3-10
 - and data files, 3-10
 - and Oracle Clusterware files, 3-6
 - buffer size parameters for, 3-10
 - Direct NFS, 3-8
 - for data files, 3-10
 - rsize, 3-10
- NIS
 - alternative to local users and groups, 2-10
- nofile
 - shell limit, 2-37
- noninteractive mode. *See* response file mode

- nproc
 - shell limit, 2-37
- NTP protocol
 - and slewing, 2-41

O

- OCR. *See* Oracle Cluster Registry
- OINSTALL group
 - about, C-1
 - and oraInst.loc, 2-5
 - checking for existing, 2-5
 - creating on other nodes, 2-15
- OINSTALL group. *See Also* Oracle Inventory group
- olsnodes command, A-7
- oper group. *See* OSOPER group
- operating system
 - checking version, 2-33
 - different on cluster members, 2-29
 - limitation for Oracle ACFS, C-6
- operating system requirements, 2-29
- optimal flexible architecture
 - and oraInventory directory, C-2
- Oracle ACFS
 - about, C-6
- Oracle base
 - grid homes not permitted under, 2-47
- Oracle base directory
 - about, C-3
 - grid home must not be in an Oracle Database
 - Oracle base, 2-7
 - minimum disk size for, 2-20
- Oracle Berkeley DB
 - restrictions for, 4-9
- Oracle Cluster Registry
 - configuration of, 4-5
 - mirroring, 3-7
 - partition sizes, 3-7
 - supported storage options, 3-5
- Oracle Clusterware
 - and file systems, 3-5
 - and upgrading Oracle ASM instances, 1-6, 3-2
 - installing, 4-1
 - rolling upgrade of, 4-2
 - supported storage options for, 3-5
 - upgrading, 3-7
- Oracle Clusterware files
 - ASM disk space requirements, 3-29
- Oracle Clusterware Installation Guide
 - replaced by Oracle Grid Infrastructure Installation Guide, xv, 4-1
- Oracle Database
 - creating data file directories, 3-24, 3-26
 - data file storage options, 3-3
 - privileged groups, 2-11
 - requirements with ASM, 3-29
- Oracle Database Configuration Assistant
 - response file, B-3
- Oracle Disk Manager
 - and Direct NFS, 3-12

- Oracle Enterprise Manager
 - ASMSNMP account, B-8
- Oracle Grid Infrastructure response file, B-3
- oracle home, 2-6
 - ASCII path restriction for, 4-5
 - multiple oracle homes, 2-6, 3-27
- Oracle Inventory group
 - about, C-1
 - checking for existing, 2-5
 - creating, 2-5
 - creating on other nodes, 2-15
 - oraInst.loc file and, 2-5
- Oracle Messaging Gateway
 - requirements, 2-32
- Oracle Net Configuration Assistant
 - response file, B-3
- Oracle patch updates, 5-1
- Oracle Software Owner user
 - configuring environment for, 2-42
 - creating, 2-6, 2-7, 2-14
 - creating on other nodes, 2-15
 - determining default shell, 2-43
 - raw device owner, 3-18, 3-24
 - required group membership, 2-10
- Oracle software owner user
 - description, 2-10
- Oracle Universal Installer
 - response files
 - list of, B-3
- Oracle Upgrade Companion, 2-2
- oracle user
 - configuring environment for, 2-42
 - creating, 2-6, 2-7, 2-14
 - creating on other nodes, 2-15
 - description, 2-10
 - determining default shell, 2-43
 - raw device owner, 3-18, 3-24
 - required group membership, 2-10
- ORACLE_BASE environment variable
 - removing from shell startup file, 2-43
- ORACLE_HOME environment variable
 - removing from shell startup file, 2-43
- ORACLE_SID environment variable
 - removing from shell startup file, 2-43
- oraInst.loc
 - and central inventory, 2-5
 - contents of, 2-5
- oraInst.loc file
 - location, 2-5
 - location of, 2-5
- oraInventory, 2-10
 - about, C-1
 - creating, 2-5
- oraInventory. *See Also* Oracle Inventory group
- OSASM group, 2-11
 - about, 2-11
 - and multiple databases, 2-12
 - and SYSASM, 2-12
 - creating, 2-13

- about, 2-11
- OSDBA group
 - and SYSDBA privilege, 2-11
 - creating, 2-13
 - creating on other nodes, 2-15
 - description, 2-11
 - raw device group, 3-18, 3-24
- OSDBA group for ASM
 - creating, 2-13
- oslevel command, 2-33
- OSOPER for ASM group
 - about, 2-11
 - creating, 2-13
- OSOPER group
 - and SYSOPER privilege, 2-11
 - creating, 2-13
 - creating on other nodes, 2-15
 - description, 2-11

P

- partition
 - using with ASM, 3-28
- passwd command, 2-8, 2-16
- passwords
 - specifying for response files, B-2
 - See also* security
- patch updates
 - download, 5-1
 - install, 5-1
 - My Oracle Support, 5-1
- patches
 - download location, 2-33
- PC X server
 - installing from, 2-4
- permissions
 - for data file directories, 3-27
- physical RAM requirements, 2-19
- policy-managed databases
 - and SCANs, C-5
- postinstallation
 - patch download and install, 5-1
 - root.sh back up, 5-2
- Precompilers
 - requirements, 2-31
- preconfigured database
 - ASM disk space requirements, 3-29
 - requirements when using ASM, 3-29
- privileged groups
 - for Oracle Database, 2-11
- Pro*C/C++
 - requirements, 2-31
- .profile file, 2-43
- PRVF-5436 error, 2-41
- PTF
 - checking, 2-33
 - download location, 2-33
- PTF download location, 2-33

R

RAC
 configuring raw logical volumes, 3-19, 3-21

RAID
 and mirroring Oracle Cluster Registry and voting disk, 3-7
 recommended ASM redundancy level, 3-28

RAM requirements, 2-19

raw disk sizes, 3-18

raw disks
 activating volume group on cluster nodes, 3-24
 and upgrades, 3-4
 checking availability of disks, 3-19, 3-21
 creating raw logical volumes, 3-18
 identifying disks, 3-19, 3-21
 importing on disk group on cluster nodes, 3-23
 initializing disks for LVM, 3-19, 3-22
 required sizes, 3-18
 specifying owner and permissions, 3-18, 3-24
 upgrading existing partitions, 3-7

raw disks desupported, xvii

recovery files
 supported storage options, 3-5

redundancy level
 and space requirements for preconfigured database, 3-29

Redundant Interconnect Usage, 2-22

relinking grid infrastructure home binaries, 5-6

relinking Oracle Grid Infrastructure home binaries, 6-2, 6-3

requirements, 3-29
 hardware, 2-19

response file installation
 preparing, B-3
 response files
 templates, B-3
 silent mode, B-5

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-3
 dbca.rsp, B-3
 enterprise.rsp, B-3
 general procedure, B-2
 Net Configuration Assistant, B-6
 netca.rsp, B-3
 passing values at command line, B-1
 passwords, B-2
 security, B-2
 specifying with Oracle Universal Installer, B-5

response files. *See also* silent mode

rolling upgrade
 ASM, 4-2, E-6
 Oracle Clusterware, 4-2

root user
 logging in as, 2-3

root.sh, 4-6
 back up, 5-2
 running, 4-3

rsize parameter, 3-10

run level, 2-19

S

SCAN listener, C-5

SCANS, 1-4, 2-24
 understanding, C-4
 use of SCANS required for clients of policy-managed databases, C-5

scripts
 root.sh, 4-3

security
 dividing ownership of Oracle software, 2-9
 See also passwords
 setting shell limits, 2-37

shell
 determining default shell for oracle user, 2-43

SHELL environment variable
 checking value of, 2-43

shell limits, 2-37

shell startup file
 editing, 2-43
 removing environment variables, 2-43

silent mode
 about, B-1
 reasons for using, B-2
 See also response files., B-1

silent mode installation, B-5

single client access names. *See* SCAN addresses

smit command, 2-8

software requirements, 2-29
 checking software requirements, 2-33

specifying owner and permissions of raw disks, 3-18, 3-24

ssh
 and X11 Forwarding, 2-44
 automatic configuration from OUI, 2-41
 configuring, D-1
 when used, 2-41

startup file
 for shell, 2-43

storage
 GPFS
 ASM files not supported on, 3-5
 HACMP, 3-5
 supported storage options
 Oracle Clusterware, 3-5

suppressed mode
 reasons for using, B-2

swap space
 requirements, 2-19

SYSASM, 2-11
 and OSASM, 2-12

SYSDBA
 using database SYSDBA on ASM
 deprecated, 2-12

- SYSDBA privilege
 - associated group, 2-11
- SYSOPER privilege
 - associated group, 2-11

T

- TEMP environment variable, 2-20
 - setting, 2-44
- temporary directory, 2-20
- temporary directory. *See* /tmp directory
- temporary disk space
 - checking, 2-20
 - freeing, 2-20
 - requirements, 2-19
- /tmp directory
 - checking space in, 2-20
 - freeing space in, 2-20
- TMPDIR environment variable, 2-20
 - setting, 2-44
- Troubleshooting
 - DBCA does not recognize Oracle ASM disk size and fails to create diskgroups, 5-5
- troubleshooting
 - and deinstalling, 6-1
 - asmcmd errors and oracle home, 2-6
 - automatic SSH configuration from OUI, 2-41
 - deconfiguring Oracle Clusterware to fix causes of root.sh errors, 6-4
 - disk space errors, 4-5
 - DISPLAY errors, 2-44
 - environment path errors, 4-5
 - INS-13001, A-6
 - intermittent hangs, 4-10, 4-11
 - log file, 4-6
 - permissions errors and oraInventory, C-1
 - permissions errors during installation, C-2
 - Reference data is not available for verifying prerequisites, A-6
 - root.sh errors, 6-4
 - run level error, 2-19
 - sqlplus errors and oracle home, 2-6
 - ssh, D-1
 - ssh configuration failure, D-2
 - ssh errors, 2-45
 - stty errors, 2-45
 - unexplained installation errors, 4-5
 - user equivalency, A-4, D-1
 - user equivalency error due to different user or group IDs, 2-8, 2-14
 - user equivalency errors, 2-6
 - voting disk backups with dd command, xix
 - X11 forwarding error, 2-45

U

- uid
 - identifying existing, 2-15
 - specifying, 2-16
 - specifying on other nodes, 2-15

- umask command, 2-42, 2-43
- uninstall, 6-1
- uninstalling, 6-1
- UNIX commands
 - cfgmgr, 3-19, 3-21
 - chdev, 3-19, 3-22
 - importvg, 3-24
 - instfix, 2-33
 - lsdev, 3-19, 3-21
 - lslpp, 2-33
 - lspv, 3-19, 3-21, 3-23
 - mkvg, 3-19, 3-22
 - oslevel, 2-33
 - passwd, 2-8
 - smit, 2-8
 - swap, 2-20
 - swapon, 2-20
 - varyoffvg, 3-23
 - varyonvg, 3-21, 3-23
 - xterm, 2-4
- upgrade
 - of Oracle Clusterware, 4-2
 - restrictions for, E-2
 - unsetting environment variables for, E-4
- upgrades, 2-2
 - and SCANs, C-5
 - of Oracle ASM, E-6
 - using raw disks with, 3-4
- upgrading
 - and existing Oracle ASM instances, 1-6, 3-2
 - and OCR partition sizes, 3-7
 - and voting disk partition sizes, 3-7
 - shared Oracle Clusterware home to local grid homes, 2-47
- user equivalence
 - testing, A-4
- user equivalency errors
 - groups and users, 2-8, 2-14
- user IDs
 - identifying existing, 2-15
 - specifying, 2-16
 - specifying on other nodes, 2-15
- useradd command, 2-14, 2-16
- users
 - creating identical users on other nodes, 2-15
 - creating the oracle user, 2-6, 2-7, 2-14
 - oracle software owner user (oracle), 2-10
 - specifying groups when creating, 2-16
 - using NIS, 2-10, 2-15

V

- varyoffvg command, 3-23
- varyonvg command, 3-21, 3-23
- virtual memory manager
 - tuning, 2-37
- volume group
 - creating, 3-19, 3-21
- volume groups
 - creating, 3-19, 3-22

- voting disks
 - backing up with dd command deprecated, xix
 - configuration of, 4-4
 - mirroring, 3-7
 - partition sizes, 3-7
 - supported storage options, 3-5

W

- WebSphere MQ
 - CSD download location, 2-34
 - requirement, 2-32
- workstation
 - installing from, 2-3
- wsize, 3-10
- wsize parameter, 3-10
- wtmax, 3-9
 - minimum value for Direct NFS, 3-9

X

- X emulator
 - installing from, 2-4
- X window system
 - enabling remote hosts, 2-3, 2-4
- X11 forwarding
 - error, 2-44
- X11 forwarding errors, D-4
- xhost command, 2-3
- xterm command, 2-4