Oracle® Clusterware

Installation Guide

11*g* Release 1 (11.1) for Linux

B28263-13

July 2011

ORACLE®

Oracle Clusterware Installation Guide, 11*g* Release 1 (11.1) for Linux

B28263-13

Copyright © 2007, 2011, Oracle and/or its affiliates. All rights reserved.

Primary Author: Douglas Williams

Contributing Authors: Mark Bauer, Namrata Bhakthavatsalam, Jonathan Creighton, Barb Lundhild, Saar Maoz, Markus Michalewicz, Hanlin Qian, Dipak Saggi, Ara Shakian

Contributors: David Austin, Tanya Bagerman, Aimee Cai, Sumanta Chatterjee, Tracy Chen, Larry Clarke, Sudip Datta, Dave Diamond, Richard Frank, Luann Ho, Julie Hu, Priyesh Jaiswal, Rajiv Jayaraman, Sameer Joshi, Roland Knapp, George Kotsovolos, Raj Kumar, Ranjith Kundapur, Seshasai Koduru, Vivekananda Kolla, Ram Kumar, Sergio Leunissen, Karen Li, Rich Long, Allen Lui, Venkat Maddali, Arnab Maity, Ofir Manor, Sundar Matpadi, Louise Morin, Anil Nair, Shoko Nishijima, Matthew McKerley, Philip Newlan, Goran Olsson, Balaji Pagadala, Soma Prasad, Srinivas Poovala, Sandesh Rao, Sudheendra Sampath, Ghassan Salem, Arun Saral, Vishal Saxena, Sanjay Sharma, David Schreiner, Vivian Schupmann, Janelle Simmons, Khethavath P. Singh, Duane Smith, Malai Stalin, Janet Stern, Jason Straub, Eri Suzuki, Madhu Velukur, Nitin Vengurlekar, Sumana Vijayagopal, Ajesh Viswambharan, Rache Wang, Pierre Wagner, Sergiusz Wolicki, Bin Yan, Jun Yang, Sivakumar Yarlagadda, Gary Young, Shi Zhao, Ricky Zhu

Contents

3 Oracle Real Application Clusters Preinstallation Tasks

4 Configuring Oracle Clusterware Storage

5 Configuring Oracle Real Application Clusters Storage

6   Installing Oracle Clusterware

7   Oracle Clusterware Postinstallation Procedures

8   Deinstallation of Oracle Clusterware

A   Troubleshooting the Oracle Clusterware Installation Process

B How to Perform Oracle Clusterware Rolling Upgrades

Index

List of Tables

Preface

Oracle Clusterware Installation Guide for Linux explains how to install and configure Oracle Clusterware, and how to configure a server and storage in preparation for an Oracle Real Application Clusters installation.

Intended Audience

Oracle Clusterware Installation Guide for Linux provides configuration information for network and system administrators, and database installation information for database administrators (DBAs) who install and configure Oracle Clusterware.

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 complete the process of configuring a system in preparation for an Oracle Clusterware installation, and complete all configuration tasks that require operating system root privileges. When configuration and installation of Oracle Clusterware 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 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

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

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

For more information, refer to the following Oracle resources:

Oracle Clusterware and Oracle Real Application Clusters Documentation

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

This installation guide reviews steps required to complete an Oracle Clusterware installation, and to perform preinstallation steps for Oracle RAC. If you intend to install Oracle Database or Oracle RAC, then review those installation guides for additional information.

Installation Guides

- *Oracle Diagnostics Pack Installation Guide*

- *Oracle Database Installation Guide for Linux*

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

Operating System-Specific Administrative Guides

- *Oracle Clusterware Administration and Deployment Guide*

- *Oracle Database Administrator's Reference, 11g Release 1 (11.1) for UNIX Systems*

- *Oracle Database Platform Guide for Microsoft Windows*

Oracle Real Application Clusters Management

- *Oracle Real Application Clusters Administration and Deployment Guide*

- *Oracle Database 2 Day + Real Application Clusters Guide*

- *Oracle Database 2 Day DBA*

- *Getting Started with the Oracle Diagnostics Pack*

Generic Documentation

- *Oracle Database New Features Guide*

- *Oracle Database Net Services Administrator's Guide*

- *Oracle Database Concepts*

- *Oracle Database Reference*

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

http://oraclestore.oracle.com/

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

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

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

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

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

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

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

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	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 Clusterware Installation and Configuration?

This section describes Oracle Database 11*g* release 1 (11.1) features as they pertain to the installation and configuration of Oracle Clusterware and Oracle Real Application Clusters (Oracle RAC) The topics in this section are:

- Changes in Installation Documentation
- Enhancements and New Features for Installation
- Oracle Validated Configuration RPM

Changes in Installation Documentation

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

Oracle Database 2 Day + Real Application Clusters Guide

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

Oracle Clusterware Installation Guide

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

Oracle Real Application Clusters Installation Guide

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

Oracle Database Storage Administrator's Guide

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

Oracle Clusterware Administration and Deployment Guide

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

Oracle Real Application Clusters Administration and Deployment Guide

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

Enhancements and New Features for Installation

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

Oracle HTTP Server Update

To install Oracle HTTP Server, use the "Oracle Fusion Middleware Web Tier Utilities 11*g* (11.1.1.2.0)" media or download.

New SYSASM Privilege and OSASM operating system group for ASM Administration

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

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

OPROCD and Hangcheck Timer Provide Fencing for Cluster Nodes

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

> **Note:** This guidance for the fencing processes has changed. Previous versions of this release guide indicated that oprocd replaced the hangcheck-timer. This was incorrect. Both processes are required.

Oracle Validated Configuration RPM

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

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

https://linux.oracle.com

1

Summary List: Installing Oracle Clusterware

The following is a summary list of installation configuration requirements and commands. This summary is intended to provide an overview of the installation process. It is written for Asianux Server, Oracle Linux or Red Hat Enterprise Linux; some tasks are slightly different for SUSE Linux.

In addition to providing a summary of the Oracle Clusterware installation process, this list also contains configuration information for preparing a system for Automatic Storage Management (ASM) and Oracle Real Application Clusters (Oracle RAC) installation.

1.1 Verify System Requirements

For more information, review the following section in Chapter 2:

"Checking the Hardware Requirements"

Enter the following commands to check available memory:

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

The minimum required RAM is 1 GB, and the minimum required swap space is 1 GB. Oracle recommends that you set swap space to twice the amount of RAM for systems with 2 GB of RAM or less. For systems with 2 GB to 8 GB RAM, use swap space equal to RAM. For systems with over 8 GB RAM, use .75 times the size of RAM.

```
df -h
```

This command checks the available space on file systems. If you use standard redundancy for Oracle Clusterware files, which is 2 Oracle Cluster Registry (OCR) partitions and 3 voting disk partitions, then you should have at least 1 GB of disk space available on separate physical disks reserved for Oracle Clusterware files. Each partition for the Oracle Clusterware files should be 256 MB in size.

The Oracle Clusterware home requires 650 MB of disk space.

```
df -h /tmp
```

Ensure that you have at least 400 MB of disk space in /tmp. If this space is not available, then increase the partition size, or delete unnecessary files in /tmp.

1.2 Check Network Requirements

For more information, review the following section in Chapter 2:

"Checking the Network Requirements"

The following is a list of address requirements that you must configure on a domain name server (DNS), or configure in the /etc/hosts file for each cluster node:

- You must have three network addresses for each node:

  - A public IP address

  - A virtual IP address, which is used by applications for failover in the event of node failure

  - A private IP address, which is used by Oracle Clusterware and Oracle RAC for internode communication

  -

- The virtual IP address has 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 virtual IP address is on the same subnet as your public interface

- The private IP address has the following requirements:

  - It should be on a subnet reserved for private networks, such as 10.0.0.0 or 192.168.0.0

  - It should use dedicated switches or a physically separate, private network, reachable only by the cluster member nodes, preferably using high-speed NICs

  - It must use the same private interfaces for both Oracle Clusterware and Oracle RAC private IP addresses

  - It cannot be registered on the same subnet that is registered to a public IP address

After you obtain the IP addresses from a network administrator, you can use the utility system-config-network to assign the public and private IP addresses to NICs, or you can configure them manually using ifconfig. Do not assign the VIP address.

Ping all IP addresses. The public and private IP addresses should respond to ping commands. The VIP addresses should not respond.

1.3 Check Operating System Packages

Refer to the tables listed inChapter 2 "Identifying Software Requirements"for details, or use a system configuration script such as the Oracle Validated RPM.

1.4 Set Kernel Parameters

For more information, review the following section in Chapter 2:

"Configuring Kernel Parameters"

Using any text editor, create or edit the /etc/sysctl.conf file, and add or edit lines similar to the following:

```
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
fs.aio-max-size =1048576
fs.file-max = 6815744
```

```
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
```

To enter these kernel settings into the running kernel, enter the following command:

```
# sysctl -p
```

1.5 Configure Groups and Users

For more information, review the following sections in Chapter 2:

"Overview of Groups and Users for Oracle Clusterware Installations"

For information about creating Oracle Database homes, review the following sections in Chapter 3:

"Creating Standard Configuration Operating System Groups and Users"

"Creating Custom Configuration Groups and Users for Job Roles"

"Creating Groups and Users for Oracle Clusterware"

For purposes of evaluation, we will assume that you have one Oracle installation owner, and that this oracle installation software owner name is `oracle`. You must create an Oracle installation owner group (`oinstall`) for Oracle Clusterware. If you intend to install Oracle Database, then you must create an OSDBA group (`dba`). Use the `id oracle` command to confirm the correct group and user configuration.

```
/usr/sbin/groupadd oinstall
/usr/sbin/groupadd dba
/usr/sbin/useradd -m -g oinstall -G dba oracle
id oracle
```

Set the password on the oracle account:

```
passwd oracle
```

1.6 Create Directories

For more information, review the following section in Chapter 2:

"Requirements for Creating an Oracle Clusterware Home Directory"

For information about creating Oracle Database homes, review the following sections in Chapter 3:

"Understanding the Oracle Base Directory Path"

"Creating the Oracle Base Directory Path"

For installations with Oracle Clusterware only, Oracle recommends that you let Oracle Universal Installer (OUI) create the Oracle Clusterware and Oracle Central Inventory (`oraInventory`) directories for you. However, as `root`, you must create a path compliant with Oracle Optimal Flexible Architecture (OFA) guidelines, so that 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.

For example:

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

1.7 Configure Oracle Installation Owner Shell Limits

For information, review the following section in Chapter 2:

"Configuring Software Owner User Environments"

As root, add the following lines to /etc/profile to set shell limits for the user oracle and for root, for the BASH, Korn, and C shells.

Add this ulimit change for each Oracle software owner you want to create

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

1.8 Configure SSH

For information, review the following section in Chapter 2:

"Configuring SSH on All Cluster Nodes"

To configure SSH, complete the following tasks:

1.8.1 Check 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 software owner that you want to use for the installation (crs, oracle), use the command ls -al to ensure that the .ssh directory is owned and writable only by the user.

1.8.2 Configure SSH on Cluster Member Nodes

Complete the following tasks on each node. You must configure SSH separately for each Oracle software installation owner that you intend to use for installation.

- Create .ssh, and create either RSA or DSA keys on each node

- Add all keys to a common authorized_keys file

1.8.3 Enable SSH User Equivalency on Cluster Member Nodes

After you have copied the `authorized_keys` file that contains all keys to each node in the cluster, start SSH on the node, and load SSH keys into memory. Note that you must either use this terminal session for installation, or reload SSH keys into memory for the terminal session from which you run the installation.

1.9 Create Storage

The following outlines the procedure for creating OCR and voting disk partitions on block devices, and creating ASM disks on block devices.

For information, review the following sections in Chapter 4:

"Configuring Storage for Oracle Clusterware Files on a Supported Shared File System"

"Configuring Disk Devices for Oracle Clusterware Files"

1.9.1 Create Disk Partitions for ASM Disks, OCR Disks, and Voting Disks

Create partitions as needed. For OCR and voting disks, create 280MB partitions for new installations, or use existing partition sizes for upgrades.

> **Note:** Every server running one or more database instances that use ASM for storage has an ASM instance. In an Oracle RAC environment, there is one ASM instance for each node, and the ASM instances communicate with each other on a peer-to-peer basis.
>
> Only one ASM instance is permitted for each node regardless of the number of database instances on the node.
>
> If you are upgrading an existing installation, then shut down ASM instances before starting installation, unless otherwise instructed in the upgrade procedure for your platform.

The following outlines the procedure for setting permissions for an ASM, OCR or voting disk partition:

1. On each node, create a permissions file in `/etc/udev/permissions.d`, to change the permissions from default root ownership. On Red Hat Enterprise Linux 4, Oracle Linux4, and Asianux Server, this file should be called `49-oracle.permissions`, so that the kernel loads it before `50-udev.permissions`. On Asianux Server 3, Oracle Linux 5, Red Hat Enterprise Linux 5, and SUSE Enterprise Server 10 systems, this file should be called `51-oracle.permissions`, so that the kernel loads it after `50-udev.permissions`.)

    For each OCR partition, the contents of the *xx*-`oracle.permissions` file are as follows:

    devicepartition:root:oinstall:0640

    For each voting disk partition, the contents of the xx-oracle.permissions file are as follows:

    devicepartition:*crs_user*:oinstall:0640

    For each ASM disk partition, the contents of the xx-oracle.permissions file are as follows:

```
devicepartition:oracle_user:OSDBA:0660
```

For example:

```
# OCR disks
sda1:root:oinstall:0640
# Voting disks
sda2:crs:oinstall:0640
# ASM disks
sdd1:oracle:dba:0660
```

2. Run the command `/sbin/partprobe` from each node.

3. Enter a command similar to the following to mark a disk as an ASM disk:

```
# /etc/init.d/oracleasm createdisk DISK1 /dev/sdb1
```

1.10 Verify Oracle Clusterware Requirements with CVU

For information, review the following section in Chapter 6:

"Verifying Oracle Clusterware Requirements with CVU"

Using the following command syntax, log in as the installation owner user (`oracle` or `crs`), and start Cluster Verification Utility (CVU) to check system requirements for installing Oracle Clusterware. In the following syntax example, replace the variable *mountpoint* with the installation media mountpoint, and replace the variable *node_list* with the names of the nodes in your cluster, separated by commas:

```
/mountpoint/runcluvfy.sh stage -pre crsinst -n node_list
```

1.11 Install Oracle Clusterware Software

For information, review the following sections in Chapter 6:

"Preparing to Install Oracle Clusterware with OUI"

"Installing Oracle Clusterware with OUI"

1. Ensure SSH keys are loaded into memory for the terminal session from which you rn the Oracle Universal Installer (OUI).

2. Navigate to the installation media, and start OUI. For example:

```
$ cd /Disk1
./runInstaller
```

3. Select **Install Oracle Clusterware**, and enter the configuration information as prompted.

1.12 Prepare the System for Oracle RAC and ASM

For information, review the following section in Chapter 5:

"Configuring Disks for Automatic Storage Management"

If you intend to install Oracle RAC, as well as Oracle Clusterware, then Oracle recommends that you use ASM for database file management, and install the Linux ASMLIB RPMs to simplify administration. ASMLib 2.0 is delivered as a set of three Linux packages:

- oracleasmlib-2.0 - the ASM libraries

- oracleasm-support-2.0 - utilities needed to administer ASMLib

- oracleasm - a kernel module for the ASM library

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

Complete the following procedures on each node that you intend to make a member of the cluster:

- Determine the Correct Oracleasm Package

- Download and Install the Oracleasm Package

- Configure ASMLib

1.12.1 Determine the Correct Oracleasm Package

Determine which kernel you are using by logging in as root and running the following command:

```
uname -rm
```

For example:

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

1.12.2 Download and Install the Oracleasm Package

After you determine the kernel version for your system, complete the following task:

1. Open a Web browser using the following URL:

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

2. Select the link for your version of Linux.

3. Download the `oracleasmlib` and `oracleasm-support` packages for your version of Linux.

4. Download the `oracleasm` package corresponding to your kernel version.

5. Log in as `root` and install the ASM packages.

1.12.3 Configure ASMLib

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

```
# /etc/init.d/oracleasm configure
```

Provide information as prompted for your system.

1.12.4 Mark ASM Disk Partitions

For OUI to recognize a disk partition as an ASM disk candidate, you must mark the disk by logging in as `root` and marking the disk partitions that you created for ASM using the following command syntax, where *ASM_DISK_NAME* is the name of the ASM disk group, and *device_name* is the name of the disk device that you want to assign to that disk group:

```
/etc/init.d/oracleasm createdisk ASM_DISK_NAME device_name
```

2

Oracle Clusterware Preinstallation Tasks

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

This chapter contains the following topics:

- About Installing the Linux Operating System
- Reviewing Upgrade Best Practices
- Logging In to a Remote Server Using X Terminal
- Overview of Groups and Users for Oracle Clusterware Installations
- Reviewing Hardware Requirements
- Checking the Hardware Requirements
- Checking the Network Requirements
- Identifying Software Requirements
- Checking the Software Requirements
- Configuring Kernel Parameters
- Configuring SSH on All Cluster Nodes
- Configuring Software Owner User Environments
- Requirements for Creating an Oracle Clusterware Home Directory
- Understanding and Using Cluster Verification Utility
- Checking Oracle Clusterware Installation Readiness with CVU

2.1 About Installing the Linux Operating System

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

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

- Completing a Minimal Linux Installation
- Completing a Default Linux Installation
- About the Oracle Unbreakable Enterprise Kernel for Linux

- [About the Oracle Unbreakable Enterprise Kernel for Linux](#)
- [About the Oracle Validated Configuration RPM](#)
- [Installing the Oracle Validated RPM From Unbreakable Linux Network](#)
- [Installing the Oracle Validated RPM From DVD Disks or Images](#)

2.1.1 Completing a Minimal Linux Installation

To complete a minimal Linux installation, select a minimal install option (either a custom installation where you select the Minimal option from Package Group Selection, or where you deselect all packages except for the Base pack). This installation lacks many RPMs required for installation. However, when you install the Oracle Validated RPM for your platform, the RPM downloads the minimum number of packages required to run Oracle Clusterware and Oracle Database.

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

Oracle Linux 4:

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

Oracle Linux 5:

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

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

https://support.oracle.com

Search for "minimal Linux".

SSH is required for Oracle Grid Infrastructure installation. OpenSSH should be included in the Linux distribution minimal installation. To confirm that SSH packages are installed, enter the following command:

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

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

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

2.1.2 Completing a Default Linux Installation

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

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

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

https://support.oracle.com

 Search for "Default RPM."

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

2.1.3 About the Oracle Unbreakable Enterprise Kernel for Linux

The Oracle Unbreakable Enterprise Kernel for Linux is available for x86-64 platforms. It is based on a stable 2.6.32 Linux kernel, and also includes optimizations developed in collaboration with Oracle Database, Oracle middleware, and Oracle hardware engineering teams to ensure stability and optimal performance for the most demanding enterprise workloads.

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

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

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

The Oracle Unbreakable Linux kernel installs directly on top of the Red Hat compatible Oracle Linux 5 or Red Hat Enterprise Linux 5, starting with Update 5, so you are not required to upgrade to a new major release of the operating system to obtain the benefits and features of this new kernel. You can obtain additional information and download the Oracle Unbreakable Enterprise Kernel for Linux at the following URL:

http://public-yum.oracle.com/

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

2.1.4 About the Oracle Validated Configuration RPM

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

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

- Automatically installs any additional packages needed for installing Oracle Grid Infrastructure and Oracle Database

- Creates an `oracle` user, and creates the oraInventory (`oinstall`) and OSDBA (`dba`) groups for that user

- Sets and verifies `sysctl.conf` settings, system startup parameters, user limits, and driver parameters to values based on recommendations from the Oracle Validated Configurations program

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

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

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

https://linux.oracle.com

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

> **Note:** The Oracle Validated RPM sets kernel parameters and resource limits only for the user account `oracle`. To use multiple software account owners, you must perform system configuration for other accounts manually.
>
> In addition, users and groups are created using the next available ID numbers. If server group and user IDs are not identical on the cluster nodes where you run the Oracle Validated RPM, then it is possible that these ids will be different on the nodes, and the different group and user IDs will cause installation to fail.
>
> To avoid this issue, run the command `id user` on all nodes, where `user` is the software owner user account. If group or user IDs are not identical on all the nodes, then change them as necessary to ensure that they are identical.

2.1.5 Installing the Oracle Validated RPM From Unbreakable Linux Network

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

1. Complete a default Oracle Linux workstation installation, or a default Red Hat Enterprise Linux installation.

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

3. Log in to Unbreakable Linux Network at the following URL:

   https://linux.oracle.com

4. Click the **Systems** tab, and in the System Profiles list, select a registered server. The System Details window opens, and displays the subscriptions for the server.

5. From the Available Channels list, select the `_base` and `_patch` channels corresponding to your Oracle Linux distribution. For example, if your distribution is Oracle Linux 5 Update 5 for x86_64, then select the following:

- Oracle Linux 5 Update 5 installation media copy (x86_64)

- Oracle Linux 5 Update 5 Patch (x86_64)

6. Click **Subscribe**.

7. From a terminal session, as `root`, enter the following command:

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

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

```
el5_u5_i386_base
el5_u5_x86_64_patch
```

8. Open a terminal session as `root`, and install the Oracle Validated Configurations Setup RPM with `up2date`, using the following command:

```
# up2date --install oracle-validated
```

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

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

> **Note:** Check the Oracle Validated Configuration RPM log file to review system configuration changes:
>
> ```
> /var/log/oracle-validated/results/orakernel.log
> ```

2.1.6 Installing the Oracle Validated RPM From DVD Disks or Images

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

1. Obtain Oracle Linux disks either by ordering the Oracle Linux Media Pack from Oracle Store, or downloading disk images from the Oracle E-Delivery Web site for Oracle Linux and Oracle VM.

Oracle Store:

https://shop.oracle.com/store/enterpriselinux

E-Delivery Web Site:

http://edelivery.oracle.com/linux

2. Start the Oracle Linux installation.

3. At the first software selection screen, which lists task-specific software options, there is an option at the bottom of the screen to customize now or customize later. Select **Customize now**, and click **Next**.

4. On the Customize selection page, select **Base System** on the list on the left side of the screen, and then select **System Tools** on the right side of the screen. Then click **Optional Packages**.

5. The Packages in System Tools window opens. Select the Oracle Validated RPM package box from the package list, and click **Next**.

6. Complete other screens to finish installing Oracle Linux.

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

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

2.2 Reviewing Upgrade Best Practices

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

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 version of Oracle Database you are running.

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

To find the most recent updates, and to find best practices recommendations about preupgrade, post-upgrade, compatibility, and interoperability, refer to "Oracle Upgrade Companion." "Oracle Upgrade Companion" is available through Note 466181.1 on Oracle*MetaLink*:

https://metalink.oracle.com

2.3 Logging In to a Remote Server 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, enable remote display as `root`.

Note: If you log in as another user (for example, `oracle`, or `grid`), then repeat this procedure for those users 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:
#
```

2.4 Overview of Groups and Users for Oracle Clusterware Installations

You must create the following group and user to install Oracle Clusterware:

- The Oracle Inventory group (typically, `oinstall`)

You must create this group the first time that you install Oracle software on the system. In Oracle documentation, this group is referred to as `oinstall`. Membership in this group controls access to OCR keys, CRS resources, and files and directories in the Oracle Clusterware home that are shared among all Oracle database owners.

> **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 `crs`) that you use to install Oracle Clusterware. Refer to "Determining If the Oracle Inventory and Oracle Inventory Group Exists" on page 2-9 to identify an existing Oracle Inventory group.

- Oracle clusterware software owner user (typically, `oracle`, if you intend to create a single software owner user for all Oracle software, or `crs`, if you intend to create separate Oracle software owners.)

  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 Clusterware software, and you can also make this user the owner of the binaries of Automatic Storage Management and Oracle Database or Oracle RAC.

  > **See Also:** *Oracle Database Administrator's Reference for UNIX Systems* and *Oracle Database Administrator's Guide* for more information about the OSDBA and OSOPER groups and the `SYSDBA`, `SYSASM` and `SYSOPER` privileges

2.5 Creating Groups and Users for Oracle Clusterware

Log in as `root`, and use the following instructions to locate or create the Oracle Inventory group and a software owner for Oracle Clusterware:

- Understanding the Oracle Inventory Group
- Understanding the Oracle Inventory Directory
- 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 Clusterware User
- Example of Creating the Oracle Clusterware User and OraInventory Path

2.5.1 Understanding the Oracle Inventory Group

You must have a group whose members are given access to write to the Oracle Central Inventory (`oraInventory`). The Central Inventory contains the following:

- A registry of the Oracle home directories (Oracle Clusterware, Oracle Database, and Automatic Storage Management) 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.

2.5.2 Understanding the Oracle Inventory Directory

The first time you install Oracle software on a system, Oracle Universal Installer checks to see if you have created an Optimal Flexible Architecture (OFA) compliant path in the format u[01-09]/app, such as `/u01/app`, and that the user running the installation has permissions to write to that path. If this is true, then Oracle Universal

Installer creates the Oracle Inventory directory in the path /u[01-09]/app/oraInventory. For example:

```
/u01/app/oraInventory
```

If you have set the environment variable $ORACLE_BASE for the user performing the Oracle Clusterware installation, then OUI creates the Oracle Inventory directory in the path $ORACLE_BASE/../oraInventory. For example, if $ORACLE_BASE is set to /opt/oracle/11, then the Oracle Inventory directory is created in the path /opt/oracle/oraInventory.

If you have created neither an OFA-compliant path nor set $ORACLE_BASE, then the Oracle Inventory directory is placed in the home directory of the user that is performing the installation. For example:

```
/home/oracle/oraInventory
```

As this placement can cause permission errors during subsequent installations with multiple Oracle software owners, Oracle recommends that you either create an OFA-compliant installation path, or set an $ORACLE_BASE environment path.

For new installations, Oracle recommends that you allow OUI to create the Central Inventory directory. By default, if you create an Oracle path in compliance with OFA structure, such as /u01/app, that is owned by an Oracle software owner, then the Central Inventory is created in the path u01/app/oraInventory using correct permissions to allow all Oracle installation owners to write to this directory.

2.5.3 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 (typically, 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.

If you have an existing Oracle 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 Inventory 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).

2.5.4 Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist

If the oraInst.loc file does not exist, then create the Oracle Inventory group by entering a command similar to the following:

```
# /usr/sbin/groupadd -g 501 oinstall
```

The preceding command creates the group oinstall, with the group ID number 501.

2.5.5 Creating the Oracle Clusterware User

You must create a software owner for Oracle Clusterware 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, such as crs, with different group membership, to give separate clusterware and database administrative privileges to those groups in a new Oracle Clusterware and Oracle Database installation.

In Oracle documentation, a user created to own only Oracle Clusterware software installations is called the crs user. A user created to own either all Oracle installations, or only Oracle database installations, is called the oracle user.

> **Note:** If you intend to use multiple Oracle software owners for different Oracle Database homes, then Oracle recommends that you create a separate Oracle software owner for Oracle Clusterware, and install Oracle Clusterware using the Oracle Clusterware software owner.
>
> If you want to create separate Oracle software owners (oracle, crs, asm) to create separate users and separate operating system privileges groups for different Oracle software installations, then note that each of these users must have the oraInventory (or oinstall) group as their primary group, and each user must share the same oraInventory, to prevent corruption of the Central Inventory. Refer to "Creating Custom Configuration Groups and Users for Job Roles" on page 3-4.

2.5.5.1 Determining if an Oracle Software Owner User Exists

To determine whether an Oracle software owner user named oracle or crs 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.

If you want to use the existing user, then ensure that the user's primary group is the Oracle Inventory group (oinstall).

2.5.5.2 Creating or Modifying an Oracle Software Owner User for Oracle Clusterware

If the Oracle software owner (`oracle`, `crs`) user does not exist, or if you require a new Oracle software owner user, then create it. The following procedure uses `crs` as the name of the Oracle software owner.

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

   ```
   # /usr/sbin/useradd -u 501 -g oinstall crs
   ```

   In the preceding command:

   - The `-u` option specifies the user ID. Using this command flag is optional, as you can allow the system to provide you with an automatically generated user ID number. However, you must make note of the user ID number of the user you create for Oracle Clusterware, as you require it later during preinstallation.

   - The `-g` option specifies the primary group, which must be the Oracle Inventory group. For example: `oinstall`.

   Use the `usermod` command to change user id numbers and groups. For example:

   ```
   # id oracle
   uid=500(oracle) gid=500(oracle) groups=500(oracle)
   # /usr/sbin/usermod -u 500 -g 501 -G 500,502 oracle
   # id oracle
   uid=500(oracle) gid=501(oinstall) groups=501(oinstall),500(oracle),502(dba)
   ```

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

   ```
   # passwd crs
   ```

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

2.5.6 Example of Creating the Oracle Clusterware User and OraInventory Path

The following is an example of how to create the Oracle Clusterware software owner (in this case, `crs`), and a path compliant with OFA structure with correct permissions for the oraInventory directory. This example also shows how to create separate Oracle Database and Oracle ASM homes with correct ownership and permissions:

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

At the end of this procedure, you will have the following:

- `/u01` owned by root.

- `/u01/app` owned by `crs:oinstall` with 775 permissions. This ownership and permissions enables OUI to create the oraInventory directory, in the path `/u01/app/oraInventory`.

- `/u01/app/crs` owned by `crs:oinstall` with 775 permissions. These permissions are required for installation, and are changed during the installation process.

- `/u01/app/oracle` owned by `oracle:oinstall` with 775 permissions.
- `/u01/app/asm` owned by `asm:oinstall` with 775 permissions.

2.6 Reviewing Hardware Requirements

Each system must meet the following minimum hardware requirements:

2.6.1 Server Architecture

- Select servers with the same chip architecture; running 32-bit and 64-bit Oracle software versions in the same cluster stack is not supported.
- If servers run the same operating system binary, then Oracle Clusterware and Oracle Real Application Clusters support servers with different hardware in the same cluster.

2.6.2 Memory and Disk or Volume Space

- At least 1 GB of physical RAM
- Swap space equivalent to the multiple of the available RAM, as indicated in the following table:

Available RAM	Swap Space Required
Between 1 GB and 2 GB	1.5 times the size of RAM
Between 2 GB and 8 GB	Equal to the size of RAM
More than 8 GB	.75 times the size of RAM

> **Note:** On Linux, the Hugepages feature allocates non-swappable memory with large page tables using memory-mapped files. If you enable Hugepages, then you should deduct the memory allocated to HugePages from the available RAM before calculating swap space.

- 400 MB of disk space in the `/tmp` directory
- 2 GB of disk space for Oracle Clusterware files, in partitions on separate physical disks, assuming standard redundancy (2 Oracle Cluster Registry partitions and 3 voting disks)

    650 MB of disk space for the Oracle Clusterware home

- For Linux x86 platforms, if you intend to install Oracle Database, allocate 4 GB of disk space for the Oracle base.
- For Linux x86_64 platforms, if you intend to install Oracle Database, allocate 4.6 GB of disk space for the Oracle base.
- If you intend to install Oracle Database single instance, allocate between 1.5 and 3.5 GB of disk space for a preconfigured database that uses file system storage. Oracle RAC software only requires 3.2 GB, and Oracle RAC software with sample schemas requires 3.3 GB.

> **Note:** The volume space requirement for databases that use
> Automatic Storage Management or shared disk storage is described
> in Chapter 5.

Additional disk space, either on a file system or in an Automatic Storage
Management disk group, is required for the Fast recovery area if you choose to
configure automated backups.

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

2.6.3 Display Resolution and Color Depth

The monitor or display visual on which you run the installation must be capable of
displaying a color depth of at least 256 colors.

> **See Also:** Setting Display and X11 Forwarding Configuration for
> operating system display configuration

2.7 Checking the Hardware Requirements

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

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

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

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

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

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

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

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

    ```
    # df -k /tmp
    ```

    This command displays disk space in 1 kilobyte blocks. On most systems, you can
    use the df command with the -h flag (df -h) to display output in
    "human-readable" format, such as "24G" and "10M." If there is less than 400 MB of
    disk space available in the /tmp directory (less than 4194304 1-k blocks), then
    complete one of the following steps:

    - Delete unnecessary files from the /tmp directory to make available the disk
      space required.

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

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

    ```
    # free
    ```

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

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

```
# uname -m
```

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

6. To determine

2.8 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 Real Application Clusters (Oracle RAC) installation:

- Network Hardware Requirements
- IP Address Requirements
- Node Time Requirements
- Network Configuration Options
- Configuring the Network Requirements

> **Note:** For the most up-to-date information about supported network protocols and hardware for Oracle RAC installations, refer to the Certify pages on the Oracle*MetaLink* Web site at the following URL:
>
> https://metalink.oracle.com

2.8.1 Network Hardware Requirements

The following is a list of requirements for network configuration:

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

  If you want to use more than one NIC for the public network or for the private network, then Oracle recommends that you use NIC bonding.

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

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

- For the public network, each network adapter must support TCP/IP.

- For the private network, the interconnect must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better required).

> **Note:** UDP is the default interconnect protocol for Oracle RAC, and TCP is the interconnect protocol for Oracle Clusterware. You must use a switch for the interconnect. Oracle recommends that you use a dedicated switch.
>
> Oracle does not support token-rings or crossover cables for the interconnect.

- For the private network, the endpoints of all designated interconnect interfaces must be completely reachable on the network. There should be no node that is not connected to every private network interface. You can test whether an interconnect interface is reachable using a `ping` command.

2.8.2 IP Address Requirements

Before starting the installation, you must have the following IP addresses available for each node:

- An IP address with an associated host name (or network name) registered in the DNS for the public interface. If you do not have an available DNS, then record the host name and IP address in the system hosts file, `/etc/hosts`.

- One virtual IP (VIP) address with an associated host name registered in a DNS. If you do not have an available DNS, then record the host name and VIP address in the system hosts file, `/etc/hosts`. 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

- A private IP address with a host name for each private interface

  Oracle recommends that you use private network IP addresses for these interfaces (for example: 10.*.*.* or 192.168.*.*). You can use DNS servers, or the `/etc/hosts` file, or both to register the private IP address. Note that if you use DNS servers alone, and the public network becomes unreachable due to NIC or cable failure, then the private IP addresses can fail to resolve.

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

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

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

Note that multiple private interfaces provide load balancing but not failover, unless bonded.

For example, if you intend to use the interfaces `eth2` and `eth3` as interconnects, then before installation, you must configure `eth2` and `eth3` with the private interconnect addresses. If the private interconnect addresses are 10.10.1.1 for `eth2` and 10.10.2.1 for `eth3`, then bond eth2 and eth3 to an interface, such as `bond0`, using a separate subnet such as 10.10.222.0. During installation, define the Oracle Clusterware private node names on 10.10.222.0, and then define 10.10.222.0 (and only that one) as a private interconnect. This ensures that Oracle Clusterware and Oracle RAC are using the same network.

After installation, if you modify interconnects on Oracle RAC with the `CLUSTER_ INTERCONNECTS` initialization parameter, then you must change it to a private IP address, on a subnet that is not used with a public IP address, nor marked as a public subnet by `oifcfg`. Oracle does not support changing the interconnect to an interface using a subnet that you have designated as a public subnet.

> **See Also:** *Oracle Clusterware Administration and Deployment Guide* for further information about setting up and using bonded multiple interfaces

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

Before installation, 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. After installation, configure clients to use either the VIP address, or the host name associated with the VIP. If a node fails, then the node's virtual IP address fails over to another node.

For example, with a two node cluster where each node has one public and one private interface, you might have the configuration shown in the following table for your network interfaces, where the hosts file is `/etc/hosts`:

Node	Host Name	Type	IP Address	Registered In
node1	node1	Public	143.46.43.100	DNS (if available, else the hosts file)
node1	node1-vip	Virtual	143.46.43.104	DNS (if available, else the hosts file)
node1	node1-priv	Private	10.0.0.1	Hosts file
node2	node2	Public	143.46.43.101	DNS (if available, else the hosts file)
node2	node2-vip	Virtual	143.46.43.105	DNS (if available, else the hosts file)
node2	node2-priv	Private	10.0.0.2	Hosts file

To enable VIP failover, the configuration shown in the preceding table defines the public and VIP addresses of both nodes on the same subnet, 143.46.43.

> **Note:** All host names must conform to the RFC 952 standard, which permits alphanumeric characters. Host names using underscores ("_") are not allowed.

2.8.3 Node Time Requirements

Before starting the installation, ensure that each member node of the cluster is set as closely as possible to the same date and time. Oracle strongly recommends using the Network Time Protocol feature of most operating systems for this purpose, with all nodes using the same reference Network Time Protocol server.

On Linux the "-x" flag can be added to the ntpd daemon to prevent the clock from going backwards.

To ensure that your cluster nodes are in synch, and to prevent the clock from going backwards, enter the following command:

```
service ntpd restart -z
```

2.8.4 Network 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 I/O. Failing to provide three separate interfaces in this case can cause performance and stability problems under load.

For high capacity clusters with a small number of multiprocessor servers, to ensure high availability, you may want to configure redundant network interfaces to prevent a NIC failure from reducing significantly the overall cluster capacity. If you are using network storage, and want to provide redundant network interfaces, then Oracle recommends that you provide six network interfaces: two for the public network interface, two for the private network interface, and two for the network storage.

2.8.5 Configuring the Network Requirements

To verify that each node meets the requirements, follow these steps:

1. If necessary, install the network adapters for the public and private networks and configure them with either public or private IP addresses.

2. If you are using a domain name server (DNS), then for each node, register the host names and IP addresses for the public network interfaces in the DNS.

3. Even if you are using a DNS, Oracle recommends that you add lines to the `/etc/hosts` file on each node, specifying the private IP addresses and associated private host names. Oracle also recommends that you add public and virtual IP addresses. Configure the /etc/host file so that it is similar to as shown in the following example, with private interface `eth1`, and private hosts `nodeint1` and `nodeint2:`, where `xxx` represents parts of a valid IP address.

```
#eth0 - PUBLIC
xxx.xxx.100.45   node1.example.com     node1
xxx.xxx.100.46   node2.example.com     node2

#eth1 - PRIVATE
10.0.0.1      nodeint1.example.com    nodeint1
10.0.0.2      nodeint2.example.com    nodeint2

#VIPs
xxx.xxx.100.47   pmvip1.example.com    nodevip1
xxx.xxx.100.48   pmvip2.example.com    nodevip2
```

4. To check network configuration, on each node, enter the following commands:

```
# /usr/bin/netstat -in
```

Ensure that each server is properly identified, and that the interface name and IP address for all network adapters that you want to specify as public or private network interfaces are properly configured. In addition, use the ping command to ensure that each node can obtain a response for the public and private IP addresses from each other node in the cluster.

> **Note:** When you install Oracle Clusterware and Oracle RAC, you will require the public, private and virtual IP addresses. Make a note of the addresses you configured in the /etc/hosts file or DNS.

5. To prevent public network failures with Oracle RAC databases using NAS devices or NFS mounts, enter the following command as root to enable the Name Service Cache Daemon (nscd):

```
# /sbin/service  nscd start
```

2.9 Identifying Software Requirements

Depending on the products that you intend to install, verify that the following operating system software is installed on the system. To check these requirements refer to the section "Checking the Software Requirements" on page 2-25, following this section.

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

For information about current supported Linux distributions, refer to Note ID 169706 on My Oracle Support, which is available from the following URL:

https://support.oracle.com

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

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

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

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

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

- Software Requirements List for x86 (32-bit) Linux Platforms
- Software Requirements List for x86-64 Linux Platforms

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

For installations only of Oracle Clusterware, ensure that you have the kernel versions and packages listed in Table 2–1 and Table 2–2.

If you intend to install Oracle Database or Oracle RAC in addition to Oracle Clusterware, then check Table 2–3 to determine if you must install additional packages for the features you plan to use.

> **Note:** For Asianux Server, Oracle Linux, and Red Hat Enterprise Linux, system requirements are identical by kernel version. Specifically:
>
> Asianux 2, Oracle Linux 4, and Red Hat Enterprise Linux 4 requirements are the same.
>
> Asianux Server 3, Oracle Linux 5, and Red Hat Enterprise Linux 5 requirements are the same.

Table 2–1 Linux x86 (32-bit) Operating System Kernel Requirements

Linux Distribution	Requirements
Asianux Distributions	▪ Asianux 2, kernel 2.6.9 or later ▪ Asianux Server 3, kernel 2.6.18 or later
Oracle Linux Distributions	▪ Oracle Linux 4 Update 4, kernel 2.6.9 or later ▪ Oracle Linux 5, kernel 2.6.9 or later
Red Hat Enterprise Linux Distributions	▪ Red Hat Enterprise Linux 4 Update 4, kernel 2.6.9 or later ▪ Red Hat Enterprise Linux 5, kernel 2.6.9 or later
SUSE Enterprise Linux Distributions	▪ SUSE 10, kernel 2.6.16.21 or later

Table 2–2   Linux x86 (32-bit) Oracle Clusterware and Oracle RAC Package Requirements

Item	Requirements
Asianux 2, Oracle Linux 4, and Red Hat Enterprise Linux 4	The following packages (or later versions) must be installed: ``` binutils-2.15.92.0.2-18 compat-libstdc++-33-3.2.3-47.3 elfutils-libelf-0.97-5 elfutils-libelf-devel-0.97-5 gcc-3.4.5-2 gcc-c++-3.4.5-2 glibc-2.3.4-2.19 glibc-common-2.3.4-2.19 glibc-devel-2.3.4-2.19 glibc-headers-2.3.4-2.19 libaio-devel-0.3.105-2 libaio-0.3.105-2 libgcc-3.4.5 libstdc++-3.4.5-2 libstdc++-devel-3.4.5-2 make-3.80-5 sysstat-5.0.5 unixODBC-2.2.11 unixODBC-devel-2.2.11 ```
Asianux Server 3, Oracle Linux 5, and Red Hat Enterprise Linux 5	The following packages (or later versions) must be installed: ``` binutils-2.17.50.0.6-2.el5 compat-libstdc++-33-3.2.3-61 elfutils-libelf-0.125 elfutils-libelf-devel-0.125 glibc-2.5-12 glibc-common-2.5-12 glibc-devel-2.5-12 glibc-headers-2.3.4-2 gcc-4.1.1-52 gcc-c++-4.1.1-52 libaio-0.3.106 libaio-devel-0.3.106 libgcc-4.1.1-52 libstdc++-4.1.1 libstdc++-devel-4.1.1-52.e15 make-3.81-1.1 sysstat-7.0.0 unixODBC-2.2.11 unixODBC-devel-2.2.11 ```

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

Item	Requirements

SUSE 10 Packages	The following packages (or later versions) must be installed:
	
	`binutils-2.16.91.0.5`
	`compat-libstdc++-5.0.7`
	`glibc-2.4-31.63`
	`glibc-devel-2.4-31.2`
	`gcc-4.1.0`
	`ksh-93r-12.9`
	`libaio-0.3.104`
	`libaio-devel-0.3.104`
	`libelf-0.8.5`
	`libgcc-4.1.0`
	`libstdc++-4.1.0`
	`libstdc++-devel-4.1.0`
	`make-3.80`
	`sysstat-6.0.2`
	`unixODBC-2.2.11`
	`unixODBC-devel-2.2.11`

Table 2–3 Linux x86 Oracle Database Features Package Requirements

Item	Requirement
LDAP package	If you did not perform a default Linux installation, you intend to use LDAP, and you want to use the scripts `odisrvreg`, `oidca`, or `schemasync`, then install the Korn shell RPM for your Linux distribution.
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	Intel C++ Compiler 9.1 or later and the version of GNU C and C++ compilers listed previously for the distribution are supported for use with these products. **Note**: Intel C++ Compiler v9.1 can be used only with gcc 3.4.5, gcc 4.0 or gcc 4.1 standard template libraries to build OCCI applications. Oracle XML Developer's Kit is supported with the same compilers as OCCI.
Oracle ODBC Drivers	If you intend to use Open Database Connectivity (ODBC), then you should install the most recent ODBC Driver Manager for Linux. You should obtain the most current ODBC driver from your operating system vendor. You can read about ODBC at the following URL: http://www.unixodbc.org The ODBC RPMs are only needed if you plan on using ODBC. If you do not plan to use ODBC, then you do not need to install the ODBC RPMs for Oracle Clusterware, Oracle ASM, or Oracle RAC.
Oracle JDBC/OCI Drivers	You can use the following optional JDK version with the Oracle JDBC/OCI drivers; however, it is not required for the installation: ■ Sun JDK 1.5.0-06 (JDK 5.0) with the JNDI extension ■ IBM Java 5.0 32-bit (SR1) or later Note: By default, IBM Java 5.0 (32-bit) is installed with this release.

Table 2–3   (Cont.) Linux x86 Oracle Database Features Package Requirements

Item	Requirement
Oracle Real Application Clusters	For a cluster file system, use one of the following options: **OCFS2** ■ Version 1.2.3 or later For information about OCFS2, refer to the following Web site: http://oss.oracle.com/projects/ocfs2/ For OCFS2 certification status, refer to the Certify page on Oracle*MetaLink*.

2.9.2 Software Requirements List for x86-64 Linux Platforms

For installations only of Oracle Clusterware, ensure that you have the kernel versions and packages listed in Table 2–4 and Table 2–5.

If you intend to install Oracle Database or Oracle RAC in addition to Oracle Clusterware, then check Table 2–6 to determine if you must install additional packages for the features you plan to use.

> **Note:**   For Asianux Server, Oracle Linux, and Red Hat Enterprise Linux, system requirements are identical by kernel version. Specifically:
>
> Asianux 2, Oracle Linux 4, and Red Hat Enterprise Linux 4 requirements are the same.
>
> Asianux Server 3, Oracle Linux 5, and Red Hat Enterprise Linux 5 requirements are the same.
>
> Oracle Unbreakable Enterprise Kernel for Linux 5 Update 5 (2.6.32), available for x86-64 systems, contains several additional features and performance enhancements not available either with Oracle Linux or with other supported Linux distributions.

Table 2–4   Linux x86-64 Operating System Kernel Requirements

Linux Distribution	Requirements
Asianux Distributions	■ Asianux 2, kernel 2.6.9 or later ■ Asianux 3, kernel 2.6.18 or later
Oracle Linux Distributions	■ Oracle Linux 4 Update 7, kernel 2.6.9 or later ■ Oracle Linux 5 Update 2, kernel 2.6.18 or later (compatible with Red Hat Enterprise kernel) ■ Oracle Unbreakable Linux 5 Update 5, kernel 2.6.32-100.0.19 or later
Red Hat Enterprise Linux Distributions	■ Red Hat Enterprise Linux 4 Update 4, kernel 2.6.9 or later ■ Red Hat Enterprise Linux 5, kernel 2.6.9 or later
SUSE Enterprise Linux Distributions	■ SUSE 10, kernel 2.6.16.21 or later

Table 2–5   Linux x86-64 Oracle Clusterware and Oracle RAC Package Requirements

Item	Requirements
Oracle Linux 5 Update 5 with Unbreakable Enterprise Kernel (2.6.32-100.0.19 or later)	Subscribe to the Oracle Linux 5 channel on the Unbreakable Linux Network, and then install the Oracle Validated RPM. This installs the Oracle Unbreakable Linux kernel, and all required kernel packages for Oracle Grid Infrastructure and Oracle Database installations.

Asianux 2, Oracle Linux 4, and Red Hat Enterprise Linux 4	The following packages (or later versions) must be installed:
	
	`binutils-2.15.92.0.2`
	`compat-libstdc++-33-3.2.3`
	`compat-libstdc++-33-3.2.3 (32 bit)`
	`elfutils-libelf-0.97`
	`elfutils-libelf-devel-0.97`
	`gcc-3.4.5`
	`gcc-c++-3.4.5`
	`glibc-2.3.4-2.19`
	`glibc-2.3.4-2.19 (32 bit)`
	`glibc-common-2.3.4`
	`glibc-devel-2.3.4`
	`glibc-devel-2.3.4 (32-bit)`
	`libaio-0.3.105`
	`libaio-0.3.105 (32 bit)`
	`libaio-devel-0.3.105`
	`libgcc-3.4.5`
	`libgcc-3.4.5 (32-bit)`
	`libstdc++-3.4.5`
	`libstdc++-3.4.5 (32 bit)`
	`libstdc++-devel 3.4.5`
	`make-3.80`
	`sysstat-5.0.5`

Asianux Server 3, Oracle Linux 5, and Red Hat Enterprise Linux 5	The following packages (or later versions) must be installed:
	
	`binutils-2.17.50.0.6`
	`compat-libstdc++-33-3.2.3`
	`compat-libstdc++-33-3.2.3 (32 bit)`
	`elfutils-libelf-0.125`
	`elfutils-libelf-devel-0.125`
	`gcc-4.1.1`
	`gcc-c++-4.1.1`
	`glibc-2.5-12`
	`glibc-2.5-12 (32 bit)`
	`glibc-common-2.5`
	`glibc-devel-2.5`
	`glibc-devel-2.5-12 (32 bit)`
	`libaio-0.3.106`
	`libaio-0.3.106 (32 bit)`
	`libaio-devel-0.3.106`
	`libgcc-4.1.1`
	`libgcc-4.1.1 (32 bit)`
	`libstdc++-4.1.1`
	`libstdc++-4.1.1 (32 bit)`
	`libstdc++-devel 4.1.1`
	`make-3.81`
	`sysstat-7.0.0`

Table 2–5   (Cont.) Linux x86-64 Oracle Clusterware and Oracle RAC Package

Item	Requirements
SUSE 10	The following packages (or later versions) must be installed:
|  | ```
binutils-2.16.91.0.5
compat-libstdc++-5.0.7-22.2
gcc-4.1.0
gcc-c++-4.1.0
glibc-2.4-31.63
glibc-32bit-2.4-31.2 (32 bit)
glibc-devel-2.4
glibc-devel-32bit-2.4 (32 bit)
libaio-0.3.104
libaio-32bit-0.3.104 (32 bit)
libelf-0.8.5
libgcc-4.1.0
libstdc++-4.1.0
libstdc++-devel-4.1.0
make-3.80
sysstat-6.0.2
``` |

Table 2–6   Linux x86-64 Oracle Database Features Package Requirements

Item	Requirement
LDAP package	If you did not perform a default Linux installation, you intend to use LDAP, and you want to use the scripts odisrvreg, oidca, or schemasync, then install the Korn shell RPM for your Linux distribution.
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	Intel C++ Compiler 9.1 or later and the version of GNU C and C++ compilers listed previously for the distribution are supported for use with these products.
	Note: Intel C++ Compiler v9.1 can be used only with gcc 3.4.5, gcc 4.0 or gcc 4.1 standard template libraries to build OCCI applications.
	Oracle XML Developer's Kit is supported with the same compilers as OCCI.

Table 2–6 (Cont.) Linux x86-64 Oracle Database Features Package Requirements

Item	Requirement
Oracle ODBC Drivers	If you intend to use Open Database Connectivity (ODBC), then you should install the most recent ODBC Driver Manager for Linux.
	You should obtain the most current ODBC driver from your operating system vendor. You can read about ODBC at the following URL:
	http://www.unixodbc.org
	The ODBC RPMs are only needed if you plan on using ODBC. If you do not plan to use ODBC, then you do not need to install the ODBC RPMs for Oracle Clusterware, Oracle ASM, or Oracle RAC.
	To use ODBC, you must also install the following additional 32-bit ODBC RPMs, depending on your operating system:
	Asianux 2, Oracle Linux 4, and Red Hat Enterprise Linux 4:
	`unixODBC-2.2.11 (32 bit) or later` `unixODBC-devel-2.2.11 (32 bit) or later`
	Asianux Server 3, Oracle Linux 5, and Red Hat Enterprise Linux 5:
	`unixODBC-2.2.11 (32 bit) or later` `unixODBC-devel-2.2.11 (32 bit) or later`
	SUSE 10:
	`unixODBC-32bit-2.2.11 (32 bit) or later` `unixODBC-devel-32bit-2.2.11 (32 bit) or later`
Oracle JDBC/OCI Drivers	You can use the following optional JDK version with the Oracle JDBC/OCI drivers; however, it is not required for the installation:
	■ Sun JDK 1.5.0-06 (JDK 5.0) with the JNDI extension
	■ IBM Java 5.0 32-bit (SR1) or later
	Note: By default, IBM Java 5.0 (32-bit) is installed with this release.
Oracle Real Application Clusters	For a cluster file system, use the following option:
	OCFS2
	■ Version 1.2.3 or later
	For information about OCFS2, refer to the following Web site:
	http://oss.oracle.com/projects/ocfs2/
	For OCFS2 certification status, refer to the Certify page on Oracle*MetaLink*.

2.10 Checking the Software Requirements

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

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

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

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

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

```
# uname -r
```

The following is sample output displayed by running this command on a Red Hat Enterprise Linux 4.0 system:

```
2.6.9-55.0.0.0.2.ELsmp
```

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

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

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

```
# rpm -q package_name
```

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

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

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

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

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

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

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

2.11 Configuring Kernel Parameters

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

On all cluster nodes, verify that the kernel parameters shown in the following table are set to values greater than or equal to the recommended value shown. The procedure following the table describes how to verify and set the values.

Parameter	Value	File
semmsl semmns semopm semmni	250 32000 100 128	`/proc/sys/kernel/sem`
shmmax	The minimum of the following (4 GB - 1 byte), or half the size of physical memory (in bytes), whichever is lower.	`/proc/sys/kernel/shmmax`
shmmni	4096	`/proc/sys/kernel/shmmni`
shmall	2097152	`/proc/sys/kernel/shmall`
file-max	6815744	`/proc/sys/fs/file-max`
ip_local_port_range	Minimum: 1024 Maximum: 65000	`/proc/sys/net/ipv4/ip_local_port_range`
rmem_default	262144	`/proc/sys/net/core/rmem_default`
rmem_max	4194304	`/proc/sys/net/core/rmem_max`
wmem_default	262144	`/proc/sys/net/core/wmem_default`
wmem_max	1048576	`/proc/sys/net/core/wmem_max`
aio-max-nr	Maximum: 1048576	`/proc/sys/fs/aio-nr`

Note: If the current value for any parameter is greater than the value listed in this table, then do not change the value of that parameter.

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

1. Enter the commands shown in the following table to view the current values of the kernel parameters:

Note: Make a note of the current values and identify any values that you must change.

Parameter	Command
`semmsl`, `semmns`, `semopm`, and `semmni`	`# /sbin/sysctl -a	grep sem`  This command displays the value of the semaphore parameters in the order listed. For example:  `250 32000 100 128`
`shmall`, `shmmax`, and `shmmni`	`# /sbin/sysctl -a	grep shm`  For example:  `kernel.shmmni = 4096` `kernel.shmall = 2097152` `kernel.shmmax = 2147483648`
`file-max`, `aio-max`	`# /sbin/sysctl -a	grep file-max`  `# /sbin/sysctl -a	grep aio-max`  For example:  `fs.file-max = 6815744` `fs.aio-max-nr=1048576`
`ip_local_port_range`	`# /sbin/sysctl -a	grep ip_local_port_range`  This command displays a range of port numbers. For example:  `net.ipv4.ip_local_port_range = 1024 65000`
`rmem_default`, `rmem_max`, `wmem_default`, and `wmem_max`	`# /sbin/sysctl -a	grep net.core.rmem`  `# /sbin/sysctl -a	grep net.core.wmem`  For example:  `net.core.rmem_default = 262144` `net.core.rmem_max = 4194304` `net.core.wmem_default = 262144` `net.core.wmem_max = 1048576`

2. If the value of any kernel parameter is different from the recommended value, then complete the following process:

   Using any text editor, create or edit the /etc/sysctl.conf file, and add or edit lines similar to the following:

   > **Note:** Include lines only for the kernel parameter values that you want to change. For the semaphore parameters (kernel.sem), you must specify all four values. However, if any of the current system parameter values are greater than the recommended values, then keep using the larger values.

   ```
   kernel.shmall = 2097152
   kernel.shmmax = 2147483648
   kernel.shmmni = 4096
   kernel.sem = 250 32000 100 128
   fs.file-max = 6815744
   fs.aio-max-nr=1048576
   net.ipv4.ip_local_port_range = 1024 65000
   net.core.rmem_default = 262144
   net.core.rmem_max = 4194304
   net.core.wmem_default = 262144
   ```

```
net.core.wmem_max = 1048576
```

By specifying the values in the `/etc/sysctl.conf` file, they persist when you restart the system.

To have these changes take effect immediately so that you do not have to restart the system, enter the following command:

```
# /sbin/sysctl -p
```

3. Enter the command `/sbin/sysctl -a` to confirm that the values are set correctly.

4. On SUSE systems only, enter the following command to cause the system to read the `/etc/sysctl.conf` file when it restarts:

```
# /sbin/chkconfig boot.sysctl on
```

5. On SUSE 10 systems only, use a text editor to change the `/etc/boot.sysctl` parameter `RUN_PARALLEL` flag from `yes` to `no`.

6. On SUSE systems only, you must enter the GID of the `oinstall` group as the value for the parameter `/proc/sys/vm/hugetlb_shm_group`. Doing this grants members of `oinstall` a group permission to create shared memory segments.

For example, where the oinstall group GID is 501:

```
# echo 501 > /proc/sys/vm/hugetlb_shm_group
```

After running this command, use `vi` to add the following text to `/etc/sysctl.conf`, and enable the `boot.sysctl` script to run on system restart:

```
vm.hugetlb_shm_group=501
```

> **Note:** Only one group can be defined as the `vm.hugetlb_shm_ group`.

7. After updating the values of kernel parameters in the `/etc/sysctl.conf` file, either restart the computer, or run the command `sysctl -p` to make the changes in the `/etc/sysctl.conf` file available in the active kernel memory.

8. Repeat steps 1 through 7 (as applicable) on all other nodes in the cluster.

2.11.1 Installing the cvuqdisk Package for Linux

Download and install the operating system package `cvuqdisk`. Without `cvuqdisk`, Cluster Verification Utility is unable to discover shared disks, and you receive the error message "Package cvuqdisk not installed" when you run Cluster Verification Utility. Use the `cvuqdisk` rpm for your hardware (for example, `x86_64`, or `i386`).

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

> **Note:** If you prefer, you can choose to disable Cluster Verification
> Utility shared disk checks by adding the following line to the file
> *CRS_home*/cv/admin/cvuconfig:
>
> CV_RAW_CHECK_ENABLED=FALSE

1. Locate the cvuqdisk RPM package, which is in the directory rpm on the installation media. If you have already installed Oracle Clusterware, then it is located in the directory *CRS_home*/rpm.

2. Copy the cvuqdisk package to each node on the cluster. You should ensure that each node is running the same version of Linux.

3. Log in as root.

4. Using the following command, check to see if you have an existing version of the cvuqdisk package:

   ```
   # rpm -qi cvuqdisk
   ```

   If you have an existing version, then enter the following command to de-install the existing version:

   ```
   # rpm -e cvuqdisk
   ```

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

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

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

   ```
   # rpm -iv cvuqdisk-1.0.1-1.rpm
   ```

2.12 Configuring SSH on All Cluster Nodes

Before you install and use Oracle Clusterware, you must configure secure shell (SSH) on all cluster nodes for the user that you plan to use to install Oracle Clusterware. In the examples that follow, the Oracle software owner listed is the crs user.

If you intend to install Oracle RAC or other Oracle software, then you should configure SSH for each of the other users (oracle, asm or other software owner) that you plan to use to install software, and either do not use a passphrase, or ensure that you and other installation users load SSH keys into memory before running the installation. As you perform this procedure, replace the example with the user name for which you are configuring SSH.

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.

The SSH configuration procedure in this section describes how to configure SSH using SSH1. If SSH is not available, then OUI attempts to use rsh and rcp instead. However, these services are disabled by default on most Linux systems.

This section contains the following:

- Checking Existing SSH Configuration on the System

- Configuring SSH on Cluster Member Nodes
- Enabling SSH User Equivalency on Cluster Member Nodes
- Setting Display and X11 Forwarding Configuration
- Preventing Oracle Clusterware Installation Errors Caused by stty Commands

2.12.1 Checking Existing SSH Configuration on the System

To determine if SSH is running, enter the following command:

```
$ pgrep sshd
```

If SSH is running, then the response to this command is one or more process ID numbers. In the home directory of the software owner that you want to use for the installation (crs, 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.

2.12.2 Configuring SSH on Cluster Member 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, crs, asm), as SSH ignores a private key file if it is accessible by others. When this is done, then start the SSH agent to load keys into memory. In the examples that follow, the RSA 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:

2.12.2.1 Create .SSH, and Create RSA Keys On Each Node

Complete the following steps on each node:

1. Log in as the software owner (in this example, the crs user).

2. To ensure that you are logged in as the Oracle user, and that the user ID matches the expected user ID you have assigned to the Oracle user, enter the commands id and id crs. Ensure that Oracle user group and user and the terminal window process group and user IDs are identical. For example:

```
$ id
uid=502(crs) gid=501(oinstall) groups=501(oinstall),502(crs)
$ id crs
uid=502(crs) gid=501(oinstall) groups=501(oinstall),502(crs)
```

3. If necessary, create the .ssh directory in the crs 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
```

4. Enter the following command:

```
$ /usr/bin/ssh-keygen -t rsa
```

At the prompts:

- Accept the default location for the key file (press **Enter**).

- Enter and confirm a pass phrase unique for this installation user.

This command writes the RSA public key to the ~/.ssh/id_rsa.pub file and the private key to the ~/.ssh/id_rsa 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 RSA key.

2.12.2.2 Add All Keys to a Common authorized_keys File

Complete the following steps:

1. On the local node, change directories to the .ssh directory in the Oracle Clusterware owner's home directory (typically, either crs or oracle).

Then, add the RSA key to the authorized_keys file using the following commands:

```
$ cat id_rsa.pub >> authorized_keys
$ ls
```

In the .ssh directory, you should see the id_rsa.pub keys that you have created, and the file authorized_keys.

2. On the local node, use SCP (Secure Copy) or SFTP (Secure FTP) to copy the authorized_keys file to the oracle user .ssh directory on a remote node. The following example is with SCP, on a node called node2, with the Oracle Clusterware owner crs, where the crs user path is /home/crs:

```
[crs@node1 .ssh]$ scp authorized_keys node2:/home/crs/.ssh/
```

You are prompted to accept an RSA 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 crs 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:

```
[crs@node1 .ssh]$ scp authorized_keys node2:/home/crs/.ssh/
The authenticity of host 'node2 (xxx.xxx.173.152) can't be established.
RSA 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' (RSA) to the list
of known hosts
crs@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, using the pass phrase you created. Then change to the .ssh directory, and using

the `cat` command, add the RSA keys for the second node to the `authorized_keys` file:

```
[crs@node1 .ssh]$ ssh node2
The authenticity of host node2 (xxx.xxx.100.102) can't be established. RSA key
fingerprint is z3:z3:33:z3:z3:33:zz:76:z3:z3:z3.
Are you sure you want to continue connecting? (yes/no)? yes
Enter passphrase for key '/home/oracle/.ssh/id_rsa':
[crs@node2 crs]$ cd .ssh
[crs@node2 ssh]$ cat id_rsa.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.

If you want to confirm that you have all nodes in the `authorized_keys` file, enter the command `more authorized_keys`, and check to see that there is an RSA key for each member node. The file lists the type of key (ssh-rsa), followed by the key, and then followed by the user and server. For example:

```
ssh-rsa AAAABBBB . . . = crs@node1
```

> **Note:** The `crs` user's `/.ssh/authorized_keys` file on every node must contain the contents from all of the `/.ssh/id_rsa.pub` files that you generated on all cluster nodes.

2.12.3 Enabling SSH User Equivalency on Cluster Member 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 Clusterware software owner is named `crs`:

1. On the system where you want to run OUI, log in as the `crs` user.

2. Use the following command syntax, where *hostname1*, *hostname2*, and so on, are the public hostnames (alias and fully qualified domain name) of nodes in the cluster to run SSH from the local node to each node, including from the local node to itself, and from each node to each other node:

```
[crs@nodename]$ ssh hostname1 date
[crs@nodename]$ ssh hostname2 date
     .
     .
     .
```

For example:

```
[crs@node1 crs]$ ssh node1 date
The authenticity of host 'node1 (xxx.xxx.100.101)' can't be established.
RSA 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' (RSA) to the list of
known hosts.
Enter passphrase for key '/home/crs/.ssh/id_rsa':
Mon Dec 4 11:08:13 PST 2006
[crs@node1 crs]$ ssh node1.example.com date
The authenticity of host 'node1.example.com (xxx.xxx.100.101)' can't be
established.
```

```
RSA 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' (RSA) to the
list of known hosts.
Enter passphrase for key '/home/crs/.ssh/id_rsa':
Mon Dec 4 11:08:13 PST 2006
[crs@node1 crs]$ ssh node2 date
Enter passphrase for key '/home/crs/.ssh/id_rsa':
Mon Dec 4 11:08:35 PST 2006
.
.
.
```

At the end of this process, the public hostname for each member node should be registered in the `known_hosts` file for all other cluster member nodes.

If you are using a remote client to connect to the local node, and you see a message similar to "Warning: No xauth data; using fake authentication data for X11 forwarding," then this means that your authorized keys file is configured correctly, but your ssh configuration has X11 forwarding enabled. To correct this issue, proceed to "Setting Display and X11 Forwarding Configuration" on page 2-35.

3. Repeat step 2 on each cluster node member.

4. On each node, enter the following commands to start the SSH agent, and to load the SSH keys into memory:

```
$ exec /usr/bin/ssh-agent $SHELL
$ /usr/bin/ssh-add
```

At the prompt, enter the pass phrase for each key that you generated.

For example:

```
[crs@node1 .ssh]$ exec /usr/bin/ssh-agent $SHELL
[crs@node1 .ssh]$ ssh-add
Enter passphrase for /home/crs/.ssh/id_rsa
Identity added: /home/crs/.ssh/id_rsa (/home/crs/.ssh/id_rsa)
```

These commands start the ssh-agent on the node, and load the RSA keys into memory so that you are not prompted to use pass phrases when issuing SSH commands

If you have configured SSH correctly, then you can now use the `ssh` or `scp` commands without being prompted for a password or a pass phrase. For example:

```
[crs@node1 ~]$ ssh node2 date
Mon Feb 26 23:34:42 UTC 2007
[crs@node1 ~]$ ssh node1 date
Mon Feb 26 23:34:48 UTC 2007
[crs@node1 ~]$ ssh node2
```

If any node prompts for a password or pass phrase, 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.

> **Note:** You must run OUI from this session, or make a note of your
> SSH pass phrase, and remember to repeat step 4 before you start OUI
> from a different terminal session.

2.12.4 Setting Display and X11 Forwarding Configuration

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

  Bourne, Korn, and Bash shells

  ```
  $ export DISPLAY=hostname:0
  ```

  C shell:

  ```
  $ setenv DISPLAY hostname:0
  ```

  For example, if you are using the Bash shell, and if your hostname is `node1`, then
  enter the following command:

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

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

  a. Using any text editor, edit or create the software installation owner's
  `~/.ssh/config` file.

  b. Make sure that the ForwardX11 attribute is set to `no`. For example:

  ```
  Host *
        ForwardX11 no
  ```

2.12.5 Preventing Oracle Clusterware Installation Errors Caused by stty Commands

During an Oracle Clusterware installation, OUI uses SSH to run commands and copy
files to the other nodes. During the installation, hidden files on the system (for
example, `.bashrc` or `.cshrc`) will cause makefile and other installation errors if they
contain `stty` commands.

To avoid this problem, you must modify these files in each Oracle installation owner
user home directory to suppress all output on STDERR, as in the following examples:

- Bourne, Bash, or Korn shell:

  ```
  if [ -t 0 ]; then
     stty intr ^C
  fi
  ```

- C shell:

  ```
  test -t 0
  if ($status == 0) then
     stty intr ^C
  endif
  ```

> **Note:** When SSH is not available, the Installer uses the `rsh` and `rcp` commands instead of `ssh` and `scp`.
>
> If there are hidden files that contain `stty` commands that are loaded by the remote shell, then OUI indicates an error and stops the installation.

2.13 Configuring Software Owner User Environments

You run OUI from the user account that you want to own the Oracle Clusterware installation (`oracle` or `crs`). However, before you start OUI you must configure the environment of the user performing the Oracle Clusterware installation. In addition, create other required Oracle software owners, if needed.

This section contains the following topics:

- Environment Requirements for Oracle Clusterware Software Owner
- Procedure for Configuring Oracle Software Owner Environments
- Setting Shell Limits for the Oracle Software Installation Users

2.13.1 Environment Requirements for Oracle Clusterware Software Owner

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

- Set the installation software owner user (`crs`, `oracle`, `asm`) 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 (`crs`, `oracle`)
- Set the software owner's environment variable DISPLAY environment variables in preparation for the Oracle Clusterware installation

2.13.2 Procedure for Configuring Oracle Software Owner Environments

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

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 `crs` user:

   ```
   $ su - crs
   ```

5. To determine the default shell for the user, enter the following command:

```
$ echo $SHELL
```

6. Open the user's shell startup file in any text editor:

   ■ Bash shell (bash):

     ```
     $ vi .bash_profile
     ```

   ■ Bourne shell (sh) or Korn shell (ksh):

     ```
     $ vi .profile
     ```

   ■ C shell (csh or tcsh):

     ```
     % vi .login
     ```

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

   ```
   umask 022
   ```

8. If the ORACLE_SID, ORACLE_HOME, or ORACLE_BASE environment 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:

    ■ Bash shell:

      ```
      $ . ./.bash_profile
      ```

    ■ Bourne, Bash, or Korn shell:

      ```
      $ . ./.profile
      ```

    ■ C shell:

      ```
      % source ./.login
      ```

11. If you are not installing the software on the local system, then enter a command similar to the following to direct X applications to display on the local system:

    ■ Bourne, Bash, or Korn shell:

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

    ■ C shell:

      ```
      % setenv DISPLAY local_host:0.0
      ```

    In this example, local_host is the host name or IP address of the system that you want to use to display OUI (your workstation or PC).

12. If you determined that the /tmp directory has less than 400 MB of free disk space, then identify a file system with at least 400 MB of free space and set the TEMP and TMPDIR environment variables to specify a temporary directory on this file system:

    > **Note:** You cannot use a shared file system as the location of the temporary file directory (typically /tmp) for Oracle RAC installation. If you place /tmp on a shared file system, then the installation fails.

a. Use the `df -h` command to identify a suitable file system with sufficient free space.

b. If necessary, enter commands similar to the following to create a temporary directory on the file system that you identified, and set the appropriate permissions on the directory:

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

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

* Bourne, Bash, or Korn shell:

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

* C shell:

```
% setenv TEMP /mount_point/tmp
% setenv TMPDIR /mount_point/tmp
```

13. To verify that the environment has been set correctly, enter the following commands:

```
$ umask
$ env | more
```

Verify that the `umask` command displays a value of 22, 022, or 0022 and that the environment variables you set in this section have the correct values.

2.13.3 Setting Shell Limits for the Oracle Software Installation Users

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

Shell Limit	Item in limits.conf	Hard Limit
Maximum number of open file descriptors	nofile	65536
Maximum number of processes available to a single user	nproc	16384

To increase the shell limits:

1. Add the following lines to the `/etc/security/limits.conf` file (the following example shows the `oracle` user as the software owner):

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

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

```
session    required    pam_limits.so
```

3. Depending on your shell environment, make the following changes to the default shell startup file (note that these examples show the user `oracle`):

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

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

- For the C shell (`csh` or `tcsh`), add the following lines to the `/etc/csh.login` file (or the file on SUSE systems)`/etc/csh.login.local`:

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

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

2.14 Requirements for Creating an Oracle Clusterware 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 separate from existing Oracle 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.

- Before installation, it should be owned by the installation owner of Oracle Clusterware (typically `oracle` for a single installation owner for all Oracle software, or `crs` for role-based Oracle installation owners), and set to 750 permissions.

For installations with Oracle Clusterware 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 Clusterware and Oracle Central Inventory (oraInventory) directories for you.

Create an Oracle Clusterware path. For example:

```
# mkdir -p  /u01/app
# chown -R crs:oinstall /u01
```

Alternatively, if you later intend to install Oracle Database software, then create an Oracle base path. OUI automatically creates an OFA-compliant path for Oracle Clusterware derived from the 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:

```
# 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 Clusterware home manually, then do not create the Oracle Clusterware home under Oracle base. Creating an Oracle Clusterware installation in an Oracle base directory will cause succeeding Oracle installations to fail.

> **See Also:** "Creating Standard Configuration Operating System Groups and Users" on page 3-1, and "Creating Custom Configuration Groups and Users for Job Roles" on page 3-4 for information about creating groups, users, and software homes for additional Oracle software installations

2.15 Understanding and Using Cluster Verification Utility

Cluster Verification Utility (CVU) is a tool that performs system checks. This guide provides Cluster Verification Utility commands to assist you with confirming that your system is properly configured for Oracle Clusterware and Oracle RAC installation.

This section describes the following topics:

- Entering Cluster Verification Utility Commands
- Using CVU to Determine if Installation Prerequisites are Complete
- Using the Cluster Verification Utility Help
- Using Cluster Verification Utility with Oracle Database 10g Release 1 or 2
- Verbose Mode and "Unknown" Output

2.15.1 Entering Cluster Verification Utility Commands

Cluster Verification Utility is provided with two scripts: runcluvfy.sh, which is designed to be used before installation, and cluvfy, which is in the path *CRS_home*/bin. The script runcluvfy.sh contains temporary variable definitions which enable it to be run before installing Oracle Clusterware or Oracle Database. After you

install Oracle Clusterware, use the command `cluvfy` to check prerequisites and perform other system readiness checks.

Before Oracle software is installed, to enter a Cluster Verification Utility command, change directories and start `runcluvfy.sh` using the following syntax:

```
cd /mountpoint
./runcluvfy.sh options
```

In the preceding example, the variable *mountpoint* represents the mountpoint path for the installation media and the variable `options` represents the Cluster Verification Utility command options that you select. For example:

```
$ cd /mnt/dvdrom
$ ./runcluvfy.sh comp nodereach -n node1,node2 -verbose
```

By default, when you enter a Cluster Verification Utility command, it provides a summary of the test. During preinstallation, Oracle recommends that you obtain detailed output by using the `-verbose` argument with the Cluster Verification Utility command. The `-verbose` argument produces detailed output of individual checks. Where applicable, it shows results for each node in a tabular layout.

2.15.2 Using CVU to Determine if Installation Prerequisites are Complete

You can use Cluster Verification Utility to determine which system prerequisites for installation are already completed. Use this option if you are installing Oracle 11*g* release 1 (11.1) on a system with a pre-existing Oracle software installation. In using this option, note the following:

- You must complete the prerequisites for using Cluster Verification Utility, notably configuring SSH between all nodes in the cluster, before you can complete a clusterwide status check.

- Cluster Verification Utility can assist you by finding preinstallation steps that need to be completed, but it cannot perform preinstallation tasks

Use the following syntax to determine what preinstallation steps are completed, and what preinstallation steps must be performed

```
$ ./runcluvfy.sh stage -pre crsinst -n node_list
```

In the preceding syntax example, replace the variable `node_list` with the names of the nodes in your cluster, separated by commas.

For example, for a cluster with mountpoint `/mnt/dvdrom/`, and with nodes `node1`, `node2`, and `node3`, enter the following command:

```
$ cd /mnt/dvdrom/
$ ./runcluvfy.sh stage -pre crsinst -n node1,node2,node3
```

Review the Cluster Verification Utility report, and proceed to the sections of the preinstallation chapter to complete additional steps as needed.

2.15.3 Using the Cluster Verification Utility Help

The `cluvfy` commands have context-sensitive help that shows correct syntax usage based on the command line arguments that you enter.

If you enter an invalid Cluster Verification Utility command, then it shows the correct usage for that command. For example, if you type `runcluvfy.sh stage -pre dbinst`, then Cluster Verification Utility shows the correct syntax for the database

preinstallation checks that Cluster Verification Utility performs with the `dbinst` stage option. The following is a list of context help commands.

- `cluvfy -help`: Cluster Verification Utility displays detailed Cluster Verification Utility command information.

- `cluvfy comp -list`: Cluster Verification Utility displays a list of components that can be checked, and brief descriptions of how each component is checked.

- `cluvfy comp -help`: Cluster Verification Utility displays detailed syntax for each of the valid component checks.

- `cluvfy stage -list`: Cluster Verification Utility displays a list of valid stages.

- `cluvfy stage -help`: Cluster Verification Utility displays detailed syntax for each of the valid stage checks.

2.15.4 Using Cluster Verification Utility with Oracle Database 10g Release 1 or 2

You can use Cluster Verification Utility on the Oracle Database 11*g* release 1 (11.1) media to check system requirements for Oracle Database 10*g* release 1 (10.1) and later installations. To use Cluster Verification Utility to check 10. 2 installations, append the command flag `-r 10gR2` to the standard Cluster Verification Utility system check commands.

For example, to perform a verification check for a Cluster Ready Services 10. 2 installation, on a system where the media mountpoint is `/mnt/dvdrom`, and the cluster nodes are `node1`, `node2`, and `node3`, enter the following command:

```
$ cd /mnt/dvdrom
$ ./runcluvfy.sh stage -pre crsinst -n node1,node2,node3 -r 10gR2
```

> **Note:** If you do not specify a release version to check, then Cluster Verification Utility checks for 11*g* release 1 (11.1) requirements.

2.15.5 Verbose Mode and "Unknown" Output

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

- Executables required by Cluster Verification Utility are missing in the `/bin` directory in the Oracle Clusterware home or Oracle home directory

- The user account starting Cluster Verification Utility does not have privileges to run common operating system executables 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

2.16 Checking Oracle Clusterware Installation Readiness with CVU

Use Cluster Verification Utility (CVU) to check your servers for their readiness to install Oracle Clusterware:

As the installation owner user (`oracle` or `crs`), ensure that you have ssh keys loaded into memory, and enter a command using the following syntax to verify that your cluster is properly configured for Oracle Clusterware installation:

`/`*`mountpoint`*`/runcluvfy.sh stage -pre crsinst -n `*`node_list`* `[-verbose]`

In the preceding syntax example, the variable *`node_list`* is a comma-delimited list of nodes in your cluster. This command checks node reachability, user and group equivalence on each node, node connectivity, and basic system requirements, including kernel versions and packages.

Select the option `-verbose` to receive progress updates as Cluster Verification Utility performs its system checks, and detailed reporting of the test results.

For example, to verify system readiness on a two-node cluster with nodes `node1` and `node2`, with the mountpoint `/mnt/dvdrom`, and with updates and a summary of the verification checks Cluster Verification Utility performs, enter the following command:

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

3

Oracle Real Application Clusters Preinstallation Tasks

This chapter describes the system configuration tasks that are generally completed by the system administrator if you plan to install Oracle Database or Oracle Database with Oracle Real Application Clusters (Oracle RAC). These tasks include creating additional groups and users for the database and for Automatic Storage Management (ASM).

You must complete these tasks before you or a database administrator start Oracle Universal Installer to install Oracle RAC. If you do not plan on installing Oracle Database on this cluster, then you can continue to the next chapter.

This chapter contains the following topics:

- Creating Standard Configuration Operating System Groups and Users
- Creating Custom Configuration Groups and Users for Job Roles
- Understanding the Oracle Base Directory Path
- Creating the Oracle Base Directory Path
- Environment Requirements for Oracle Database and Oracle ASM Owners
- Overview of HugePages

> **Note:** Any task required for Oracle Database is also required for Oracle RAC, unless otherwise stated.

3.1 Creating Standard Configuration Operating System Groups and Users

A standard configuration is a configuration with the default groups and users that Oracle Universal Installer (OUI) displays by default during Oracle database installation, which are not created already for Oracle Clusterware installation.

The following sections describe how to create the required operating system user and groups for Oracle Database or Oracle Database with Oracle RAC and ASM installations:

To allocate separate operating system user privileges to different administrative users, refer to "Creating Custom Configuration Groups and Users for Job Roles" on page 3-4.

- Overview of Groups and Users for Oracle Database Installations
- Creating Standard Operating System Groups and Users

3.1.1 Overview of Groups and Users for Oracle Database Installations

The following operating system groups and user are required if you plan to install Oracle Database:

- The OSDBA group (typically, dba)

    You must create this group the first time you install Oracle Database software on the system. In a standard installation, you are prompted to one group to grant the following privileges to its members:

    - Database Administrator (OSDBA)

    - Database Operator (OSOPER)

    - ASM Administrator (OSASM)

    In addition, members of this group are granted database write access to ASM (OSDBA for ASM).

    The default name for this group is dba.

- An unprivileged user

    Verify that the unprivileged user nobody exists on the system. The nobody user must own the external jobs (extjob) executable after the installation.

3.1.2 Creating Standard Operating System Groups and Users

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

- Verifying That the User nobody Exists

- Creating the OSDBA Group

- Creating Identical Users and Groups on Other Cluster Nodes

3.1.2.1 Verifying That the User nobody Exists

If you intend to install Oracle Database or Oracle RAC, then complete the following procedure to verify that the user nobody exists on the system:

1. To determine if the user exists, enter the following command:

    ```
    # id nobody
    ```

    If this command displays information about the nobody user, then you do not have to create that user.

2. If the nobody user does not exist, then enter a command similar to the following to create it:

    ```
    # /usr/sbin/useradd -u 65001 nobody
    ```

3. Repeat this procedure on all the other nodes in the cluster. Note that the ID number for uid and gid should be the same on all nodes of the cluster.

3.1.2.2 Creating the OSDBA Group

To create the OSDBA group, complete the following procedure:

1. Enter a command similar to the following:

    ```
    # /usr/sbin/groupadd -g 502 dba
    ```

The preceding command creates the group `dba`, with the group ID number 502.

3.1.2.3 Creating Identical Users and Groups on Other Cluster Nodes

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

Oracle software owner users and groups must exist and be identical on all cluster nodes.

Identifying Existing User and Group IDs

To determine the user ID (UID) of the `oracle` user, and the group IDs (GID) of the Oracle Inventory, OSDBA, and OSOPER groups, follow these steps:

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

   ```
   # id oracle
   ```

   The output from this command is similar to the following:

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

2. From the output, identify the user ID (UID) for the user and the group identities (GIDs) for the groups to which it belongs. Ensure that these are identical on each node.

Creating Users and Groups on the Other Cluster Nodes

To create users and groups on the other cluster nodes, repeat the following procedure on each node:

1. Log in to the next cluster node as `root`.

2. Enter commands similar to the following to create groups. Use the `-g` option to specify the correct GID for each group.

   ```
   # /usr/sbin/groupadd -g 501 oinstall
   # /usr/sbin/groupadd -g 502 dba
   ```

   > **Note:** If a group already exists, but has a different group ID, then use the `groupmod` command to modify it if necessary. If you cannot use the same group ID for a particular group on this node, then view the `/etc/group` file on all nodes to identify a group ID that *is* available on every node. You must then specify that ID for the group on all of the nodes.

3. To create the `oracle` user or another required user, enter a command similar to the following (in this example, to create the `oracle` user):

   ```
   # /usr/sbin/useradd -u 501 -g oinstall oracle
   ```

   In the preceding command:

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

- The `-g` option specifies the primary group, which must be the Oracle Inventory group; for example, `oinstall`

> **Note:** If the user already exists, then use the `usermod` command to modify it if necessary. If you cannot use the same user ID for the user on this 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
   ```

3.2 Creating Custom Configuration Groups and Users for Job Roles

A Custom configuration is a configuration with groups and users that divide access privileges granted by membership in separate operating system groups and users.

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

To allocate operating system user privileges to a minimum number of groups and users, refer to Creating Standard Configuration Operating System Groups and Users on page 3-1.

- Overview of Creating Operating System Group and User Options Based on Job Roles

- Creating Database Operating System Groups and Users with Job Role Separation

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

3.2.1 Overview of Creating Operating System Group and User Options Based on Job Roles

This section provides an overview of how to create users and groups to divide access privileges by job roles. Log in as `root` to create these groups and users.

- Users for Oracle Installations with Job Role Separation

- Database Groups for Job Role Installations

- ASM Groups for Job Role Installations

3.2.1.1 Users for Oracle Installations with Job Role Separation

Oracle recommends that you create the following operating system group 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, `crs` for Oracle Clusterware, and `asm` for Oracle ASM.

  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 Clusterware software, and you can also make this user the owner of the binaries of Automatic Storage Management and Oracle Database or Oracle RAC.

  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, 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 more than one database on the system.

  In Oracle documentation, a user created to own the Oracle Clusterware binaries is called the `crs` user.

  If you intend to use Automatic Storage Management (ASM), then Oracle recommends that you create a separate user to own ASM files. In Oracle documentation, that user is referred to as `asm`.

  > **See Also:** *Oracle Database Administrator's Reference for UNIX Systems* and *Oracle Database Administrator's Guide* for more information about the OSDBA, OSASM and OSOPER groups and the `SYSDBA`, `SYSASM` and `SYSOPER` privileges

- **An unprivileged user**

  Verify that the unprivileged user `nobody` exists on the system. The `nobody` user must own the external jobs (extjob) executable after the installation.

3.2.1.2 Database Groups for Job Role 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 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.

  If you want to specify a group name other than the default `dba` group, then you must choose the Custom 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.

  On Automatic Storage Manager (ASM) instances, members of the OSDBA group are given privileges to perform all administrative privileges granted to the

SYSASM privileges, including mounting and dismounting disk groups. This privileges grant is deprecated, and will be removed in a future release.

- 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.

  If you want to use the OSOPER group to create a database administrator group with fewer privileges than the default dba group, then you must choose the Custom 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.

3.2.1.3 ASM Groups for Job Role Installations

SYSASM is a new system privilege that enables the separation of the ASM storage administration privilege from SYSDBA. Members of the database OSDBA group are granted SYSASM privileges, but this privilege is deprecated, and may be removed in a future release.

Use the Custom Installation option to designate separate operating system groups as the operating system authentication groups for privileges on ASM. Before you start OUI, create the following groups and users for ASM

- The Oracle Automatic Storage Management Group (typically asm)

  SYSASM privileges for ASM files provide administrator privileges for storage file equivalent to SYSDBA privileges for the database. In Oracle documentation, the operating system group whose members are granted SYSASM privileges is called the OSASM group, and in command lines, is referred to as asm.

  If you have more than one database on your system, then you must create a separate OSASM group, and a separate ASM user. ASM can support multiple databases.

  Members of the OSASM group can use SQL to connect to an 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. In this release of Oracle Clusterware and Oracle Database, SYSASM privileges and SYSDBA privileges are equivalent, but using SYSDBA privileges to perform ASM management tasks on ASM instances is deprecated. SYSDBA privileges may be limited on ASM instances in a future release.

- The OSDBA group for ASM (typically asmdba)

  Members of the OSDBA group for ASM are granted read and write access to files managed by ASM. The Oracle database software owner (typically oracle) must be a member of this group, and all users with OSDBA membership on databases that you want to have access to the files managed by ASM should be members of the OSDBA group for ASM

3.2.2 Creating Database Operating System 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 for Custom Installations

- [Creating an OSOPER Group](#)

- [Creating the OSASM Group](#)

- [Creating the OSDBA Group for ASM](#)

- [Creating the Oracle Software Owner User](#)

- [Creating a Separate ASM Owner](#)

- [Verifying That the User nobody Exists](#)

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

3.2.2.1 Creating the OSDBA Group for Custom Installations

You must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist, for example, if this is the first installation of Oracle Database software on the system

- An OSDBA group exists, but you want to give a different group of operating system users database administrative privileges for a new Oracle Database installation

If the OSDBA group does not exist or if you require a new OSDBA group, then create it as follows. In the following procedure, use the group name `dba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 502 dba
```

3.2.2.2 Creating an OSOPER Group

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. If you want to use an OSOPER group, then you must create it in the following circumstances:

- If an OSOPER group does not exist; for example, if this is the first installation of Oracle Database software on the system

- If an OSOPER group exists, but you want to give a different group of operating system users database operator privileges in a new Oracle installation

If you require a new OSOPER group, then create it as follows. In the following, use the group name `oper` unless a group with that name already exists.

```
# /usr/sbin/groupadd -g 505 oper
```

3.2.2.3 Creating the OSASM Group

If the OSASM group does not exist or if you require a new OSASM group, then create it as follows. In the following procedure, use the group name `asm` unless a group with that name already exists:

```
# /usr/sbin/groupadd -504 asm
```

3.2.2.4 Creating the OSDBA Group for ASM

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

If the OSDBA group for ASM does not exist or if you require a new OSDBA group for ASM, then create it as follows. In the following procedure, use the group name `asmdba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 506 asmdba
```

3.2.2.5 Creating 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 Clusterware, such as `crs`, and you want to create a separate Oracle software owner for Oracle Database software, such as `dba`.

3.2.2.5.1 Determining if an Oracle Software Owner User Exists To determine whether an Oracle software owner user named `oracle` or `crs` 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. If you want 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:

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

- To modify an existing user, refer to the "Modifying an Existing Oracle Software Owner User" section on page 3-9.

- To create a user, refer to the following section.

3.2.2.5.2 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. In the following procedure, use the user name `oracle` unless a user with that name already exists.

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

   ```
   # /usr/sbin/useradd -u 502 -g oinstall -G dba oracle
   ```

   In the preceding command:

   - The -u option specifies the user ID. Using this command flag is optional, as you can allow the system to provide you with an automatically generated user ID number. However, you must make note of the `oracle` user ID number, as you require it later during preinstallation.

   - The -g option specifies the primary group, which must be the Oracle Inventory group--for example, `oinstall`

- The `-G` option specifies the secondary groups, which must include the OSDBA group, and, if required, the OSOPER group. For example: dba, or dba,oper

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

   ```
   # passwd oracle
   ```

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

```
# /usr/sbin/usermod -g oinstall -G dba[,oper] oracle
```

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

3.2.2.6 Creating a Separate ASM Owner

1. To create `asm`, enter a command similar to the following:

   ```
   # /usr/sbin/useradd -u 504 -g oinstall -G asm asm
   ```

   In the preceding command:

   - The -u option specifies the user ID. Using this command flag is optional, as you can allow the system to provide you with an automatically generated user ID number. However, you must make note of the `asm` ID number, as you require it later during preinstallation.

   - The `-g` option specifies the primary group, which must be the Oracle Inventory group--for example, `oinstall`

   - The `-G` option specifies the secondary groups, which must include the OSASM group. For example: `asm`.

2. Set the password for `asm`:

   ```
   # passwd asm
   ```

3.2.2.7 Verifying That the User nobody Exists

Before installing the software, complete the following procedure to verify that the user nobody exists on the system:

1. To determine if the user exists, enter the following command:

   ```
   # id nobody
   ```

   If this command displays information about the `nobody` user, then you do not have to create that user.

2. If the `nobody` user does not exist, then enter the following command syntax to create it:

   ```
   # /usr/sbin/useradd -u number nobody
   ```

   for example:

   ```
   # /usr/sbin/useradd -u 65555 nobody
   ```

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

3.2.2.8 Creating Identical Database Users and Groups on Other Cluster Nodes

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

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

Identifying Existing User and Group IDs

To determine the user ID (UID) of the crs, `oracle,` or `asm` users, and the group IDs (GID) of the Oracle Inventory, OSDBA, and OSOPER groups, follow these steps:

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

   ```
   # id oracle
   ```

   The output from this command is similar to the following:

   ```
   uid=502(oracle) gid=501(oinstall) groups=502(dba),503(oper),506(asmdba)
   ```

2. From the output, identify the user ID (UID) for the user and the group identities (GIDs) for the groups to which it belongs. Ensure that these ID numbers are identical on each node of the cluster.

Creating Users and Groups on the Other Cluster Nodes

To create users and groups on the other cluster nodes, repeat the following procedure on each node:

1. Log in to the next cluster node as `root`.

2. Enter commands similar to the following to create the `oinstall` and `dba` groups, and if required, the `oper` and `asm` groups. Use the `-g` option to specify the correct GID for each group.

   ```
   # /usr/sbin/groupadd -g 501 oinstall
   # /usr/sbin/groupadd -g 502 crs
   # /usr/sbin/groupadd -g 503 dba
   # /usr/sbin/groupadd -g 505 oper
   # /usr/sbin/groupadd -g 504 asm
   # /usr/sbin/groupadd -g 506 asmdba
   ```

> **Note:** If the group already exists, then use the `groupmod` command to modify it if necessary. If you cannot use the same group ID for a particular group on this node, then view the `/etc/group` file on all nodes to identify a group ID that *is* available on every node. You must then change the group ID on all nodes to the same group ID.

3. To create the `oracle` or asm user, enter a command similar to the following (in this example, to create the `oracle` user):

```
# /usr/sbin/useradd -u 502 -g oinstall -G dba,asmdba,[,oper] oracle
```

In the preceding command:

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

- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example `oinstall`

- The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group. For example: `dba` or `dba,oper,asmdba.`

> **Note:** If the user already exists, then use the `usermod` command to modify it if necessary. If you cannot use the same user ID for the user on every node, then view the `/etc/passwd` file on all nodes to identify a user ID that *is* available on every node. You must then specify that ID for the user on all of the nodes.

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

```
# passwd oracle
```

5. Complete SSH configuration for each user as described in the section "Configuring SSH on All Cluster Nodes" on page 2-30.

6. Complete user environment configuration tasks for each user as described in the section Configuring Software Owner User Environments on page 2-36.

3.3 Understanding the Oracle Base Directory Path

This section contains information about preparing an Oracle base directory.

3.3.1 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. However, setting an Oracle base directory may become mandatory in a future release.

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.

3.3.2 Understanding Oracle Base and Oracle Clusterware Directories

Even if you do not use the same software owner to install Oracle Clusterware and Oracle Database, be aware that the `root.sh` script in the clusterware installation changes ownership of the Oracle Clusterware home directory to `root`. For this reason, the Oracle Clusterware home cannot be in the same location as other Oracle software.

3.4 Creating the Oracle Base Directory Path

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

For OUI to recognize the path as an Oracle software path, it must be in the form u0[*1-9*]/app, and it must be writable by any member of the oinstall group.

Oracle recommends that you create an Oracle base path manually. 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:

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

3.5 Environment Requirements for Oracle Database and Oracle ASM Owners

If you create separate Oracle installation owner accounts for the database or ASM, then complete the following tasks for the Oracle Database software owner (oracle) and Oracle ASM software owner (asm).

- If you create an Oracle base path, as described in the preceding section, then set the path to the Oracle base directory as an environment variable for the Oracle database owner. For example:

  ```
  # ORACLE_BASE=/u01/app/oracle; export ORACLE_BASE
  ```

- On Linux systems, ensure that the PATH variable lists $ORACLE_HOME/bin before /usr/X11R6/bin

  This is necessary if you want to run the RMAN (Recovery Manager) binary for Oracle, rather than the rman (RosettaMan) binary in Linux that is located in /usr/X11R6/bin.

- 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 ulimit settings for file descriptors and processes for the installation software owner (asm, oracle)

- Set the software owners' environment variable DISPLAY environment variables in preparation for the ASM or Oracle Database installation

3.6 Overview of HugePages

You can choose to configure HugePages. For some uses, HugePages can provide enhanced performance. However, this feature is an advanced configuration option. It is not a requirement for Oracle RAC.

The following is an overview of HugePages. It does not provide RPM or configuration information. The tasks you must perform for configuration depend on kernel distribution and hardware on your system. If you decide that you want to configure

your cluster nodes to use HugePages, then refer to your distribution documentation and to Oracle Technical Network (OTN) and Oracle*MetaLink* for further information.

3.6.1 What HugePages Provides

HugePages is a feature integrated into the Linux kernel with release 2.6. It is a method to have larger pages where it is useful for working with very large memory. It can be useful for both 32-bit and 64-bit configurations. HugePage sizes vary from 2MB to 256MB, depending on the kernel version and the hardware architecture. For Oracle Databases, using HugePages reduces the operating system maintenance of page states, and increases TLB (Translation Lookaside Buffer) hit ratio.

Without HugePages, the operating system keeps each 4 KB of memory as a page, and when it is allocated to the SGA, then the lifecycle of that page (dirty, free, mapped to a process, and so on) must be kept up to date by the operating system kernel.

With HugePages, the operating system page table (virtual memory to physical memory mapping) is smaller, since each page table entry is pointing to pages from 2 MB to 256 MB in size. Also, the kernel has fewer pages whose lifecyle must be monitored.

For example, if you use HugePages with 64-bit hardware, and you want to map 256 MB of memory, you may need one page table entry (PTE). If you do not use HugePages, and you want to map 256 MB of memory, then you must have 256 MB * 1024 KB/4 KB = 65536 PTEs.

3.6.2 Restrictions for HugePage Configurations

To use HugePages, you must configure Grub to allocate memory for HugePages during system startup. After paging space is reserved, HugePages can be used as needed. However, if the space they require is not reserved in memory during system startup, then a HugePages allocation may fail.

HugePages are not subject to allocation or release after system startup, unless a system administrator changes the hugepages configuration by modifying the number of pages available, or the pool size.

4

Configuring Oracle Clusterware Storage

This chapter describes the storage configuration tasks that you must complete before you start Oracle Universal Installer to install Oracle Clusterware.

This chapter contains the following topics:

- Reviewing Storage Options for Oracle Clusterware Files
- Checking for Available Shared Storage with CVU
- Configuring Storage for Oracle Clusterware Files on a Supported Shared File System
- Configuring Disk Devices for Oracle Clusterware Files

4.1 Reviewing Storage Options for Oracle Clusterware Files

This section describes supported options for storing Oracle Clusterware.

4.1.1 Overview of Oracle Clusterware Storage Options

There are two ways of storing Oracle Clusterware files:

- **A supported shared file system**: Supported file systems include the following:

  - **A supported cluster file system**

    Note: For information about how to download and configure OCFS2, refer to the following URL

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

    See Also: The Certify page on Oracle*Metalink* for supported cluster file systems

  - **Network File System (NFS)**: A file-level protocol that enables access and sharing of files

    See Also: The Certify page on Oracle*Metalink* for supported Network Attached Storage (NAS) devices

- **Block or Raw Devices**: Oracle Clusterware files can be placed on either Block or RAW devices based on shared disk partitions. Oracle recommends using Block devices for easier usage.

4.1.2 General Storage Considerations for Oracle Clusterware

For all installations, you must choose the storage option that you want to use for Oracle Clusterware files and Oracle Database files, or for Oracle Clusterware with Oracle Real Application Clusters databases (Oracle RAC). You do not have to use the same storage option for each file type.

Oracle Clusterware files include voting disks, used to monitor cluster node status, and Oracle Cluster Registry (OCR) which contains configuration information about the cluster. The voting disks and OCR are shared files on a cluster or network file system environment. If you do not use a cluster file system, then you must place these files on shared block devices or shared raw devices. Oracle Universal Installer (OUI) automatically initializes the OCR during the Oracle Clusterware installation.

For voting disk file placement, Oracle recommends that each voting disk is configured so that it does not share any hardware device or disk, or other single point of failure. Any node that does not have available to it an absolute majority of voting disks configured (more than half) will be restarted.

The following table shows the storage options supported for storing Oracle Clusterware files. Oracle Clusterware files include the Oracle Cluster Registry (OCR), a mirrored OCR file (optional), the Oracle Clusterware voting disk, and additional voting disk files (optional).

> **Note:** For information about OCFS2, refer to the following Web site:
>
> http://oss.oracle.com/projects/ocfs2/
>
> For OCFS2 certification status, refer to the Certify page on Oracle*MetaLink*.

Table 4–1 Supported Storage Options for Oracle Clusterware

Storage Option	File Types Supported	
	OCR and Voting Disks	Oracle Software
Automatic Storage Management	No	No
OCFS2	Yes	Yes
Local storage	No	Yes
NFS file system **Note:** Requires a certified NAS device. Direct NFS is not certified for Oracle Clusterware files.	Yes	Yes
Shared disk partitions (block devices or raw devices)	Yes	No

Use the following guidelines when choosing the storage options that you want to use for Oracle Clusterware:

- 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 cannot use ASM to store Oracle Clusterware files, because these files must be accessible before any ASM instance starts.

- 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.

4.1.3 Quorum Disk Location Restriction with Existing 9.2 Clusterware Installations

When upgrading your Oracle9i release 9.2 Oracle RAC environment to Oracle Database 11*g* release 1 (11.1), 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 11*g* release 1 (11.1). You cannot reuse the old Oracle9*i* release 9.2 quorum disk for this purpose.

4.1.4 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: Check for available shared storage with CVU

Refer to Checking for Available Shared Storage with CVU on page 4-3.

2: Configure shared storage for Oracle Clusterware files

- **To use a file system for Oracle Clusterware files**, refer to Configuring Storage for Oracle Clusterware Files on a Supported Shared File System on page 4-4.

- To use block devices for Oracle Clusterware files, refer to Configuring Disk Devices for Oracle Clusterware Files on page 4-8.

4.2 Checking for Available Shared Storage with CVU

To check for all shared file systems available across all nodes on the cluster on a supported shared file system, log in as the installation owner user (`oracle` or `crs`), and use the following syntax:

```
/mountpoint/runcluvfy.sh comp ssa -n node_list
```

If you want to check the shared accessibility of a specific shared storage type to specific nodes in your cluster, then use the following command syntax:

```
/mountpoint/runcluvfy.sh comp ssa -n node_list -s storageID_list
```

In the preceding syntax examples, the variable `mountpoint` is the mountpoint path of the installation media, the variable `node_list` is the list of nodes you want to check, separated by commas, and the variable `storageID_list` is the paths for the storage devices that you want to check.

For example, if you want to check the shared accessibility from `node1` and `node2` of storage devices `/dev/sdb` and `/dev/sdc`, and your mountpoint is `/dev/dvdrom/`, then enter the following command:

```
$ /mnt/dvdrom/runcluvfy.sh comp ssa -n node1,node2 -s /dev/sdb,/dev/sdc
```

If you do not specify storage device IDs in the command, then the command searches for all available storage devices connected to the nodes on the list.

4.3 Configuring Storage for Oracle Clusterware Files on a Supported Shared File System

Oracle Universal Installer (OUI) 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 File System for Oracle Clusterware Files

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

- Checking NFS Mount and Buffer Size Parameters for Clusterware

- Creating Required Directories for Oracle Clusterware Files on Shared File Systems

> **Note:** The OCR is a file that contains the configuration information and status of the cluster. Oracle Universal Installer (OUI) automatically initializes the OCR during the Oracle Clusterware installation. Database Configuration Assistant uses the OCR for storing the configurations for the cluster databases that it creates.
>
> The OCR is a shared file in a cluster file system environment. If you do not use a cluster file system, then you must place this file on a shared storage device.

4.3.1 Requirements for Using a File System for Oracle Clusterware Files

To use a file system for Oracle Clusterware files, the file system must comply with the following requirements:

- To use a cluster file system, it must be a supported cluster file system, as listed in the section "Deciding to Use a Cluster File System for Oracle Clusterware Files" on page 4-5.

- To use an NFS file system, it must be on a certified NAS device. Log in to Oracle*MetaLink* at the following URL, and click the Certify tab to find a list of certified NAS devices.

  https://metalink.oracle.com/

- If you choose to place your Oracle Cluster Registry (OCR) files on a shared file system, then Oracle recommends that one of the following is true:

  - The disks used for the file system are on a highly available storage device, (for example, a RAID device that implements file redundancy).

  - At least two file systems are mounted, and use the features of Oracle Clusterware 11*g* release 1 (11.1) to provide redundancy for the OCR.

- The user account with which you perform the installation (`oracle` or `crs`) must have write permissions to create the files in the path that you specify.

> **Note:** If you are upgrading from Oracle9*i* release 2, then you can continue to use the raw device or shared file that you used for the SRVM configuration repository instead of creating a new file for the OCR.
>
> If you are upgrading Oracle Clusterware, and your existing cluster uses 100 MB OCR and 20 MB voting disk partitions, then you can continue to use those partition sizes.

Use Table 4–2 to determine the partition size for shared file systems.

Table 4–2   Shared File System Volume Size Requirements

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

In Table 4–2, 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 1.3 GB of storage available over a minimum of three volumes (two separate volume locations for the OCR and OCR mirror, and one voting disk on each volume).

> **Note:** When you create partitions with `fdisk` by specifying a device size, such as `+256M`, 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 configuration software checks to ensure that devices contain a minimum of 256MB of available disk space. Therefore, Oracle recommends using at least 280MB for the device size. You can check partition sizes by using the command syntax `fdisk -s `*`partition`*. For example:
>
> ```
> [root@node1]$ fdisk -s /dev/sdb1
> 281106
> ```

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

For Linux x86 (32-bit) and x86-64 (64-bit) platforms, Oracle provides a cluster file system, OCFS2. You can have a shared Oracle home on OCFS2.

If you have an existing Oracle installation, then use the following command to determine if OCFS2 is installed:

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

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

OCFS2:

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

Note: For OCFS2 certification status, refer to the Certify page on Oracle*MetaLink*:

https://metalink.oracle.com

4.3.3 Checking NFS Mount and Buffer Size Parameters for Clusterware

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 are:

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

Update the /etc/fstab file on each node with an entry containing the NFS mount options for your platform. For example, if your platform is x86-64, then update the /etc/fstab files with an entry similar to the following:

```
nfs_server:/vol/CWfiles  /u01/oracle/cwfiles nfs \
rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600,noac
```

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 provide an entry similar to the following for a binaries mount point:

```
nfs_server:/vol/crshome /u01/oracle/crs nfs -yes
rw,bg,hard,nointr,rsize=32768,wsize=32768,tcp,vers=3,timeo=600,actime=0
```

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

https://metalink.oracle.com

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

4.3.4 Creating Required Directories for Oracle Clusterware Files on Shared File Systems

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

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

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

1. If necessary, configure the shared file systems 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. 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 that you want to use. Choose a file system with a minimum of 560 MB of free disk space (one OCR and one voting disk, with external redundancy).

   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 (typically, `crs` or `oracle`) has permissions to create directories on the disks where you plan to install Oracle Clusterware and Oracle Database, then OUI creates the Oracle Clusterware file directory, and 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 the Oracle Clusterware home (or CRS home). For example, where the user is `oracle`, and the CRS home is `oracrs`:

   ```
   # mkdir /mount_point/oracrs
   # chown oracle:oinstall /mount_point/oracrs
   # chmod 750 /mount_point/oracrs
   ```

   > **Note:** After installation, directories in the installation path for the Oracle Cluster Registry (OCR) files should be owned by `root`, and not writable by any account other than `root`.

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

4.4 Configuring Disk Devices for Oracle Clusterware Files

On Linux systems, `O_DIRECT` enables direct read and writes to block devices, avoiding kernel overhead. With Oracle Clusterware release 10.2 and later, Oracle Clusterware files are configured by default to use direct input/output.

With the 2. 6 kernel or later for Red Hat Enterprise Linux, Oracle Linux, and SUSE Enterprise Server, you must create a permissions file to maintain permissions on Oracle Cluster Registry (OCR) and voting disk partitions. If you do not create this permissions file, then permissions on disk devices revert to their default values, `root:disk`, and Oracle Clusterware will fail to start.

On Asianux 2, Red Hat Enterprise Linux 4, and Oracle Linux 4, you must create a permissions file number that is lower than 50.

On Asianux Server 3, Red Hat Enterprise Linux 5, Oracle Linux 5, or SUSE Enterprise Linux 10, you must create a permissions file number that is higher at 50.

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

1. **Create a permissions file** in `/etc/udev/permissions.d`, to change the permissions from default root ownership to root and members of the oinstall group, called `49-oracle.permissions`, or `51-oracle.permissions`, depending on your Linux distribution. In each case, the contents of the *xx*-oracle.permissions file are as follows:

   devicepartition:root:oinstall:0640

   For example, to set permissions for an OCR partition on block device `/dev/sda1`, create the following entry:

   ```
   sda1:root:oinstall:0640
   ```

   Use the section "Example of Creating a Udev Permissions File for Oracle Clusterware" on page 4-8 for a step-by-step example of how to perform this task.

2. **Configure the block devices** from the local node with the required partition space for Oracle Clusterware files. Use the section "Example of Configuring Block Device Storage for Oracle Clusterware" on page 4-9 to help you configure block devices, if you are unfamiliar with creating partitions.

3. **Load updated block device partition tables** on all member nodes of the cluster, using `/sbin/partprobe devicename`. You must do this as `root`.

4. **Change the ownership of OCR partitions to the installation owner** on all member nodes of the cluster. In the session where you run the Installer, the OCR partitions must be owned by the installation owner (`crs`, `oracle`) that performs the Oracle Clusterware installation. The installation owner must own the OCR partitions so that the Installer can write to them. During installation, the Installer changes ownership of the OCR partitions back to `root`. With subsequent system restarts, ownership is set correctly by the `oracle.permissions` file.

5. Enter the command `/sbin/udevstart`. This command should assign the permissions set in the `oracle.permissions` file. Check to ensure that your system is configured correctly.

4.4.1 Example of Creating a Udev Permissions File for Oracle Clusterware

The procedure to create a permissions file to grant oinstall group members write privileges to block devices is as follows:

1. Log in as root.

2. Change to the `/etc/udev/permissions.d` directory:

```
# cd /etc/udev/permissions.d
```

3. Start a text editor, such as vi, and enter the partition information where you want to place the OCR and voting disk files, using the syntax *device*[*partitions*]:root:oinstall:0640. Note that Oracle recommends that you place the OCR and the voting disk files on separate physical disks. For example, to grant oinstall members access to SCSI disks to place OCR files on sda1 and sdb2, and to grant the Oracle Clusterware owner (in this example crs) permissions to place voting disks on sdb3, sdc1 and sda2, add the following information to the file:

```
# OCR disks
sda1:root:oinstall:0640
sdb2:root:oinstall:0640
# Voting disks
sda2:crs:oinstall:0640
sdb3:crs:oinstall:0640
sdc1:crs:oinstall:0640
```

4. Save the file:

   On Asianux 2, Oracle Linux 4, and Red Hat Enterprise Linux 4 systems, save the file as `49-oracle.permissions`.

   On Asianux Server 3, Oracle Linux 5, Red Hat Enterprise Linux 5, and SUSE Enterprise Server 10 systems, save the file as `51-oracle.permissions`.

5. Using the following command, assign the permissions in the udev file to the devices:

```
# /sbin/udevstart
```

4.4.2 Example of Configuring Block Device Storage for Oracle Clusterware

The procedure to create partitions for Oracle Clusterware files on block devices is as follows:

1. log in as root

2. Enter the fdisk command to format a specific storage disk (for example, `/sbin/fdisk /dev/sdb`)

3. Create a new partition, and make the partition 280 MB in size for both OCR and voting disk partitions.

4. Use the command syntax `/sbin/partprobe` *diskpath* on each node in the cluster to update the kernel partition table for the shared storage device on each node.

The following is an example of how to use fdisk to create one partition on a shared storage block disk device for an OCR file:

```
[crs@localnode /] $ su
Password:
[root@localnode /] # /sbin/fdisk /dev/sdb
The number of cylinders for this disk is set to 1024.
Command (m for help): n
Command action
  e    extended
  P    primary partition (1-4)
p
Partition number (1-4): 1
```

```
First cylinder (1-1024, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-1024, default 1)
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-1024, default 1024):+280m

Command (m for help):w

The partition table has been altered!
Calling ioctl () to re-read partition table.
Synching disks.
[root@localnode /]#
exit
[crs@localnode /] $ ssh remotenode
Last login Wed Feb 21 20:23:01 from localnode
[crs@remotenode ~]$ su
Password:
[root@localnode /] # /sbin/partprobe /dev/sdb1
```

> **Note:** Oracle recommends that you create partitions for Oracle
> Clusterware files on physically separate disks.

5

Configuring Oracle Real Application Clusters Storage

This chapter includes storage administration tasks that you should complete if you intend to use Oracle Clusterware with Oracle Real Application Clusters (Oracle RAC).

This chapter contains the following topics:

5.1 Reviewing Storage Options for Oracle Database and Recovery Files

This section describes supported options for storing Oracle Database files, and data files.

> **See Also:** The Oracle Certify site for a list of supported vendors for Network Attached Storage options:
>
> http://www.oracle.com/technology/support/metalink/
>
> Refer also to the Certify site on Oracle*Metalink* for the most current information about certified storage options:
>
> https://metalink.oracle.com/

5.1.1 Overview of Oracle Database and Recovery File Options

There are three ways of storing Oracle Database and recovery files:

- **Automatic Storage Management**: Automatic Storage Management (ASM) is an integrated, high-performance database file system and disk manager for Oracle Database files. It performs striping and mirroring of database files automatically.

> **Note:** For Standard Edition Oracle Database installations using
> Oracle RAC, ASM is the only supported storage option.
>
> Only one 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:

  - 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. If you
    intend to store Oracle Clusterware files on OCFS2, then you must ensure that
    OCFS2 volume sizes are at least 500 MB each.

    > **See Also:** The Certify page on Oracle*Metalink* for supported cluster
    > file systems

  - **NAS Network File System (NFS) listed on Oracle Certify**: 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
    Clusterware.

    > **See Also:** The Certify page on Oracle*Metalink* for supported
    > Network Attached Storage (NAS) devices, and supported cluster file
    > systems

- **Block or Raw Devices**: A partition is required for each database file. If you do not
  use ASM, then for new installations on raw devices, you must use a custom
  installation.

  > **Note:** On Linux, Oracle recommends using block devices for new
  > installations.

5.1.2 General Storage Considerations for Oracle RAC

For all installations, you must choose the storage option that you want to use for
Oracle Database files, or for Oracle Clusterware with Oracle RAC. If you want to
enable automated backups during the installation, then you must also choose the
storage option that you want to use for recovery files (the Fast recovery area). You do
not have to use the same storage option for each file type.

For single-instance Oracle Database installations using Oracle Clusterware for failover,
you must use OCFS2, ASM, block devices (on Linux), or shared raw disks if you do
not want the failover processing to include dismounting and remounting of local file
systems.

The following table shows the storage options supported for storing Oracle Database
files and Oracle Database recovery files. Oracle Database files include data files,
control files, redo log files, the server parameter file, and the password file.

> **Note:** For the most up-to-date information about supported storage options for Oracle RAC installations, refer to the Certify pages on the Oracle*MetaLink* Web site:
>
> https://metalink.oracle.com

Table 5–1 Supported Storage Options for Oracle Database and Recovery Files

Storage Option	File Types Supported	
	Database	Recovery
Automatic Storage Management	Yes	Yes
OCFS2	Yes	Yes
Red Hat Global File System (GFS); for Red Hat Enterprise Linux and Oracle Linux	Yes	Yes
Local storage	No	No
NFS file system	Yes	Yes
Note: Requires a certified NAS device		
Shared raw devices	Yes	No
Shared block devices	Yes	No

Use the following guidelines when choosing the storage options that you want 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 Automatic Storage Management (ASM) as the storage option for database and recovery files.

- For Standard Edition Oracle RAC installations, ASM is the only supported storage option for database or recovery files.

- You cannot use ASM to store Oracle Clusterware files, because these files must be accessible before any ASM instance starts.

- If you intend to use ASM with Oracle RAC, and you are configuring a new ASM instance, then your system must meet the following conditions:

  - All nodes on the cluster have the 11*g* release 1 (11.1) version of Oracle Clusterware installed.

  - Any existing ASM instance on any node in the cluster is shut down.

- If you intend to upgrade an existing Oracle RAC database, or an Oracle RAC database with ASM instances, then you must ensure that your system meets the following conditions:

  - Oracle Universal Installer (OUI) and Database Configuration Assistant (DBCA) are run on the node where the Oracle RAC database or Oracle RAC database with ASM instance is located.

  - The Oracle RAC database or Oracle RAC database with an ASM instance is running on the *same* nodes that you intend to make members of the new

cluster installation. For example, if you have an existing Oracle RAC database running on a three-node cluster, then you must install the upgrade on all three nodes. You cannot upgrade only 2 nodes of the cluster, removing the third instance in the upgrade.

> **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.

5.1.3 After You Have Selected Disk Storage Options

After you have installed and configured Oracle Clusterware storage, and after you have reviewed your disk storage options for Oracle Database files, you must perform the following tasks in the order listed:

1: Check for available shared storage with CVU

Refer to Checking for Available Shared Storage with CVU on page 5-4.

2: Configure storage for Oracle Database files and recovery files

- **To use a shared file system for database or recovery file storage**, refer to Configuring Storage for Oracle Database Files on a Supported Shared File System on page 5-5, 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.

- **To use Automatic Storage Management for database or recovery file storage**, refer to "Configuring Disks for Automatic Storage Management with ASMLIB" on page 5-15

- **To use shared devices for database file storage**, refer to "Configuring Storage for Oracle Database Files on Shared Storage Devices" on page 5-19.

> **Note:** If you choose to configure database files on raw devices, note that you must complete database software installation first, and then configure storage after installation.
>
> You cannot use OUI to configure a database that uses raw devices for storage. In a future release, the option to use raw and block devices for database storage will become unavailable.

5.2 Checking for Available Shared Storage with CVU

To check for all shared file systems available across all nodes on the cluster on a supported shared file system, log in as the installation owner user (`oracle` or `crs`), and use the following syntax:

```
/mountpoint/runcluvfy.sh comp ssa -n node_list
```

If you want to check the shared accessibility of a specific shared storage type to specific nodes in your cluster, then use the following command syntax:

```
/mountpoint/runcluvfy.sh comp ssa -n node_list -s storageID_list
```

In the preceding syntax examples, the variable *mountpoint* is the mountpoint path of the installation media, the variable *node_list* is the list of nodes you want to check, separated by commas, and the variable *storageID_list* is the list of storage device IDs for the storage devices managed by the file system type that you want to check.

For example, if you want to check the shared accessibility from node1 and node2 of storage devices /dev/sdb and /dev/sdc, and your mountpoint is /dev/dvdrom/, then enter the following command:

```
$ /mnt/dvdrom/runcluvfy.sh comp ssa -n node1,node2 -s /dev/sdb,/dev/sdc
```

If you do not specify storage device IDs in the command, then the command searches for all available storage devices connected to the nodes on the list.

5.3 Choosing a Storage Option for Oracle Database Files

Database files consist of the files that make up the database, and the recovery area files. There are five options for storing database files:

- OCFS2, GFS

- Network File System (NFS)

- Automatic Storage Management (ASM)

- Block devices (Database files only--not for the recovery area)

- Raw devices (Database files only--not for the recovery area)

During configuration of Oracle Clusterware, if you selected OCFS2 or NFS, and the volumes that you created are large enough to hold the database files and recovery files, then you have completed required preinstallation steps. You can proceed to Chapter 6, "Installing Oracle Clusterware" on page 6-1.

If you want to place your database files on ASM, then proceed to Configuring Disks for Automatic Storage Management on page 5-11.

If you want to place your database files on raw devices, and manually provide storage management for your database and recovery files, then proceed to "Configuring Storage for Oracle Database Files on Shared Storage Devices" on page 5-19.

> **Note:** Databases can consist of a mixture of ASM files and non-ASM files. Refer to *Oracle Database Administrator's Guide* for additional information about ASM. For OCFS2 certification status, refer to the Certify page on Oracle*MetaLink*.

5.4 Configuring Storage for Oracle Database Files on a Supported Shared File System

Review the following sections to complete storage requirements for Oracle Database files:

- Requirements for Using a File System for Oracle Database Files

- Deciding to Use a Cluster File System for Data Files

- Deciding to Use NFS for Data Files

- Deciding to Use Direct NFS for Datafiles

- Enabling Direct NFS Client Oracle Disk Manager Control of NFS

- Disabling Direct NFS Client Oracle Disk Management Control of NFS
- Checking NFS Mount Buffer Size Parameters for Oracle RAC
- Creating Required Directories for Oracle Database Files on Shared File Systems

5.4.1 Requirements for Using a File System for Oracle Database Files

To use a file system for Oracle Database files, the file system must comply with the following requirements:

- To use a cluster file system, it must be a supported cluster file system, as listed in the section "Deciding to Use a Cluster File System for Data Files" on page 5-6.
- To use an NFS file system, it must be on a certified NAS device.
- If you choose to place your database files on a shared file system, then one of the following must be true:
  - The disks used for the file system are on a highly available storage device, (for example, a RAID device that implements file redundancy).
  - The file systems consist of at least two independent file systems, with the database files on one file system, and the recovery files on a different file system.
- The `oracle` user must have write permissions to create the files in the path that you specify.

Use Table 5–2 to determine the partition size for shared file systems.

Table 5–2   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 5–2, the total required volume size is cumulative. For example, to store all database files on the shared file system, you should have at least 3.4 GB of storage available over a minimum of two volumes.

5.4.2 Deciding to Use a Cluster File System for Data Files

For Linux x86 (32-bit) and x86-64 (64-bit) platforms, Oracle provides a cluster file system, OCFS2. You can have a shared Oracle home on OCFS2.

If you have an existing Oracle installation, then use the following command to determine if OCFS2 is installed:

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

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

OCFS2:

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

> **Note:** For OCFS2 certification status, refer to the Certify page on Oracle*MetaLink*:
>
> https://metalink.oracle.com

5.4.3 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.

5.4.4 Deciding to Use Direct NFS for Datafiles

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

5.4.4.1 About Direct NFS Storage

With Oracle Database 11*g* release 1 (11.1), 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. The mount options used in mounting the file systems are not relevant, as Direct NFS manages settings after installation. Refer to your vendor documentation to complete NFS configuration and mounting.

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

> **Note:** Use NFS servers certified for Oracle RAC. Refer to the following URL for certification information:
>
> https://metalink.oracle.com

5.4.4.2 Using the Oranfstab File with Direct NFS

If you use Direct NFS, then you can choose to use a new file specific for Oracle datafile 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`. The `oranfstab` file is not required to use NFS or Direct NFS.

With Oracle RAC installations, if you want to use Direct NFS, then you must replicate the file `/etc/oranfstab` on all nodes, and keep each `/etc/oranfstab` file synchronized on all nodes.

When the `oranfstab` file is placed in `$ORACLE_HOME/dbs`, the entries in the file are specific to a single database. In this case, all nodes running an Oracle RAC database use the same `$ORACLE_HOME/dbs/oranfstab` file.

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

In all cases, mount points must be mounted by the kernel NFS system, even when they are being served using Direct NFS.

5.4.4.3 Mounting NFS Storage Devices with Direct NFS

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

Direct NFS searches for mount entries in the following order:

1. `$ORACLE_HOME/dbs/oranfstab`

2. `/etc/oranfstab`

3. `/etc/mtab`

Direct NFS uses the first matching entry found.

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

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

If Oracle Database is unable to open an NFS server using Direct NFS, then Oracle Database uses the platform operating system kernel NFS client. In this case, the kernel NFS mount options must be set up as defined in "Checking NFS Mount Buffer Size Parameters for Oracle RAC" on page 5-10. Additionally, an informational message will be logged into the Oracle alert and trace files indicating that Direct NFS could not be established.

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

5.4.4.4 Specifying Network Paths with the Oranfstab File

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

Use the following views for Direct NFS management:

- **v$dnfs_servers**: Shows a table of servers accessed using Direct NFS.

- **v$dnfs_files**: Shows a table of files currently open using Direct NFS.

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

- **v$dnfs_stats**: Shows a table of performance statistics for Direct NFS.

5.4.5 Enabling Direct NFS Client Oracle Disk Manager Control of NFS

Complete the following procedure to enable Direct NFS:

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

   - **Server**: The NFS server name.

   - **Path**: Up to four network paths to the NFS server, specified either by IP address, or by name, as displayed using the `ifconfig` command.

   - **Export**: The exported path from the NFS server.

   - **Mount**: The local mount point for the NFS server.

   Note: On Linux and UNIX platforms, the location of the `oranfstab` file is `$ORACLE_HOME/dbs`.

   The following is an example of an `oranfstab` file with two NFS server entries:

   ```
   server:  MyDataServer1
   path:  132.34.35.12
   path:  132.34.35.13
   export: /vol/oradata1 mount: /mnt/oradata1

   server: MyDataServer2
   path:  NfsPath1
   path:  NfsPath2
   path:  NfsPath3
   path:  NfsPath4
   export: /vol/oradata2 mount: /mnt/oradata2
   export: /vol/oradata3 mount: /mnt/oradata3
   export: /vol/oradata4 mount: /mnt/oradata4
   export: /vol/oradata5 mount: /mnt/oradata5
   ```

2. Oracle Database uses an ODM library, `libnfsodm11.so`, to enable Direct NFS. To replace the standard ODM library, `$ORACLE_HOME/lib/libodm11.so`, with the ODM NFS library, `libnfsodm11.so`, complete the following steps:

   a. Change directory to `$ORACLE_HOME/lib`.

   b. Enter the following commands:

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

5.4.6 Disabling Direct NFS Client Oracle Disk Management Control of NFS

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

- Remove the `oranfstab` file.

- Restore the stub `libodm11.so` file by reversing the process you completed in step 2b, "Enabling Direct NFS Client Oracle Disk Manager Control of NFS"

- 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.

5.4.7 Checking NFS Mount Buffer Size Parameters for Oracle RAC

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

If you are using Direct NFS, note that Direct NFS will not work if the backend NFS server does not support a write size (`wtmax`) of 32768 or greater. In that case, NFS mounts fall back to a standard kernel NFS mount path.

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

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

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

5.4.8 Creating Required 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 a RAC database).

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. Ensure that the file systems are configured to mount automatically when a node restarts.

2. Use the `df -h` command to determine the free disk space on each mounted file system.

3. From the display, identify the file systems that you want to use:

File Type	File System Requirements
Database files	Choose either: - 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 (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 the `oracle` user the owner of these directories, this permits them to be read by multiple Oracle homes, including those with different OSDBA groups.

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

5.5 Configuring Disks for Automatic Storage Management

This section describes how to configure disks for use with Automatic Storage Management. Before you configure the disks, you must determine the number of disks and the amount of free disk space that you require. The following sections describe how to identify the requirements and configure the disks:

- Identifying Storage Requirements for Automatic Storage Management

- Using an Existing Automatic Storage Management Disk Group

- Configuring Disks for Automatic Storage Management with ASMLIB

Note: For Automatic Storage Management installations:

- Although this section refers to disks, you can also use zero-padded files on a certified NAS storage device in an Automatic Storage Management disk group. Refer to *Oracle Database Installation Guide for Linux* for information about creating and configuring NAS-based files for use in an Automatic Storage Management disk group.

- You can run ASM using ASMLIB, or run ASM using raw devices. Oracle recommends that you use raw devices only with upgrades, and migrate to other storage systems for Oracle Database files.

5.5.1 Identifying Storage Requirements for Automatic Storage Management

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

1. Determine whether you want to use Automatic Storage Management for Oracle Database files, recovery files, or both.

   > **Note:** You do not have to use the same storage mechanism for database files and recovery files. You can use the file system for one file type and Automatic Storage Management for the other.
   >
   > If you choose to enable automated backups and you do not have a shared file system available, then you must choose Automatic Storage Management for recovery file storage.

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

   - If you select an installation method that runs Database Configuration Assistant in interactive mode (for example, by choosing the Advanced database configuration option) then you can decide whether you want to use the same Automatic Storage Management disk group for database files and recovery files, or use different disk groups for each file type.

     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 use the same Automatic Storage Management disk group for database files and recovery files.

2. Choose the Automatic Storage Management redundancy level that you want to use for the Automatic Storage Management disk group.

   The redundancy level that you choose for the Automatic Storage Management disk group determines how Automatic Storage Management mirrors files in the disk group and determines the number of disks and amount of free disk space that you require, as follows:

   - External redundancy

     An external redundancy disk group requires a minimum of one disk device. The effective disk space in an external redundancy disk group is the sum of the disk space in all of its devices.

     Because Automatic Storage Management does not mirror data in an external redundancy disk group, Oracle recommends that you select external redundancy only use only RAID or similar devices that provide their own data protection mechanisms for disk devices.

   - Normal redundancy

     In a normal redundancy disk group, to increase performance and reliability, Automatic Storage Management by default uses two-way mirroring. A normal redundancy disk group requires a minimum of two disk devices (or two

failure groups). The effective disk space in a normal redundancy disk group is *half* the sum of the disk space in all of its devices.

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

■ High redundancy

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

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 the database files and recovery files.

Use the following table to determine the minimum number of disks and the minimum disk space requirements for installing the starter database:

Redundancy Level	Minimum Number of Disks	Database Files	Recovery Files	Both File Types
External	1	1.15 GB	2.3 GB	3.45 GB
Normal	2	2.3 GB	4.6 GB	6.9 GB
High	3	3.45 GB	6.9 GB	10.35 GB

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

15 + (2 * *number_of_disks*) + (126 * *number_of_Automatic_Storage_Management_instances*)

For example, for a four-node Oracle RAC installation, using three disks in a high redundancy disk group, you require an additional 525 MB of disk space:

15 + (2 * 3) + (126 * 4) = 525

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

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

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

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

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

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

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

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

   - All of the devices in an Automatic Storage Management disk group should be the same size and have the same performance characteristics.

   - Do not specify more than one partition on a single physical disk as a disk group device. Automatic Storage Management expects each disk group device to be on a separate physical disk.

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

5.5.2 Using an Existing Automatic Storage Management Disk Group

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

- If you select an installation method that runs Database Configuration Assistant in interactive mode (for example, by choosing the Advanced database configuration option), then you can decide whether you want to create a disk group, or to use an existing one.

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

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

> **Note:** The Automatic Storage Management instance that manages the existing disk group can be running in a different Oracle home directory.

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

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

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

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

   ```
   +ASM2:oracle_home_path
   ```

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

2. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the Automatic Storage Management instance that you want to use.

3. Connect to the Automatic Storage Management instance as the `SYS` user with `SYSDBA` privilege and start the instance if necessary:

   ```
   $ $ORACLE_HOME/bin/sqlplus "SYS/SYS_password as SYSDBA"
   SQL> STARTUP
   ```

4. Enter the following command to view the existing disk groups, their redundancy level, and the amount of free disk space in each one:

   ```
   SQL> SELECT NAME,TYPE,TOTAL_MB,FREE_MB FROM V$ASM_DISKGROUP;
   ```

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.

5.5.3 Configuring Disks for Automatic Storage Management with ASMLIB

A disk that is configured for use with Automatic Storage Management is known as a candidate disk.

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

The Automatic Storage Management library driver (ASMLIB) simplifies the configuration and management of the disk devices by eliminating the need to rebind disk devices used with ASM each time the system is restarted, or to configure persistent path information in the operating system configuration files.

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

- Installing and Configuring the ASM Library Driver Software

- Configuring Disk Devices to Use ASM Library Driver on x86 Systems

- Administering the ASM Library Driver and Disks

> **Note:** To create a database during the installation using the ASM library driver, you must choose an installation method that runs DBCA in interactive mode. For example, you can run DBCA in an interactive mode by choosing the Custom installation type, or the Advanced database configuration option. You must also change the default disk discovery string to `ORCL:*`.

5.5.3.1 Installing and Configuring the ASM Library Driver Software

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

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

To install and configure the ASMLIB driver software manually, follow these steps:

1. Enter the following command to determine the kernel version and architecture of the system:

    ```
    # uname -rm
    ```

2. Download the required ASMLIB packages from the OTN Web site:

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

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

    You must install the following packages, where *version* is the version of the ASMLIB driver, *arch* is the system architecture, and *kernel* is the version of the kernel that you are using:

    ```
    oracleasm-support-version.arch.rpm
    oracleasm-kernel-version.arch.rpm
    oracleasmlib-version.arch.rpm
    ```

3. Switch user to the `root` user:

    ```
    $ su -
    ```

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

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

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

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

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

```
# /etc/init.d/oracleasm configure
```

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

Prompt	Suggested Response
Default user to own the driver interface:	**Standard groups and users configuration**: Specify the Oracle software owner user (typically, `oracle`)
	Job role separation groups and users configuration: Specify the ASM software owner user (typically `asm`).
Default group to own the driver interface:	Standard groups and users configuration: Specify the OSDBA group (typically `dba`).
	Job role separation groups and users configuration: Specify the OSASM group (typically `asm`).
Start Oracle Automatic Storage Management Library driver on boot (y/n):	Enter `y` to start the Oracle Automatic Storage Management library driver when the system starts.
Fix permissions of Oracle ASM disks on boot? (y/n)	Specify the Oracle software owner (the owner of the RDBMS installation).

The script completes the following tasks:

- Creates the `/etc/sysconfig/oracleasm` configuration file

- Creates the `/dev/oracleasm` mount point

- Loads the `oracleasm` kernel module

- Mounts the ASMLIB driver file system

Note: The ASMLIB driver file system is not a regular file system. It is used only by the Automatic Storage Management library to communicate with the Automatic Storage Management driver.

---

7. Repeat this procedure on all nodes in the cluster where you want to install Oracle RAC.

5.5.3.2 Configuring Disk Devices to Use ASM Library Driver on x86 Systems

To configure the disk devices that you want to use in an Automatic Storage Management disk group, follow these steps:

1. If you intend to use IDE, SCSI, or RAID devices in the Automatic Storage Management disk group, then follow these steps:

   a. If necessary, install or configure the shared disk devices that you intend to use for the disk group and restart the system.

   b. To identify the device name for the disks that you want to use, enter the following command:

   ```
   # /sbin/fdisk -l
   ```

   Depending on the type of disk, the device name can vary:

Disk Type	Device Name Format	Description
IDE disk	/dev/hd*x*n	In this example, *x* is a letter that identifies the IDE disk and *n* is the partition number. For example, /dev/hda is the first disk on the first IDE bus.
SCSI disk	/dev/sd*x*n	In this example, *x* is a letter that identifies the SCSI disk and *n* is the partition number. For example, /dev/sda is the first disk on the first SCSI bus.

   To include devices in a disk group, you can specify either whole-drive device names or partition device names.

   > **Note:** Oracle recommends that you create a single whole-disk partition on each disk that you want to use.

   c. Use either `fdisk` or `parted` to create a single whole-disk partition on the disk devices that you want to use.

2. Enter a command similar to the following to mark a disk as an Automatic Storage Management disk:

   ```
   # /etc/init.d/oracleasm createdisk DISK1 /dev/sdb1
   ```

   In this example, `DISK1` is a name that you want to assign to the disk.

   > **Note:** The disk names that you specify can contain uppercase letters, numbers, and the underscore character. They must start with an uppercase letter.
   >
   > If you are using a multi-pathing disk driver with Automatic Storage Management, then make sure that you specify the correct logical device name for the disk.

3. To make the disk available on the other nodes in the cluster, enter the following command as `root` on each node:

   ```
   # /etc/init.d/oracleasm scandisks
   ```

   This command identifies shared disks attached to the node that are marked as Automatic Storage Management disks.

5.5.3.3 Administering the ASM Library Driver and Disks

To administer the Automatic Storage Management library driver and disks, use the
`oracleasm` initialization script with different options, as described in Table 5–3.

Table 5–3   ORACLEASM Script Options

Option	Description
configure	Use the `configure` option to reconfigure the Automatic Storage Management library driver, if necessary: `# /etc/init.d/oracleasm configure`
enable disable	Use the `disable` and `enable` options to change the actions of the Automatic Storage Management library driver when the system starts. The `enable` option causes the Automatic Storage Management library driver to load when the system starts: `# /etc/init.d/oracleasm enable`
start stop restart	Use the `start`, `stop`, and `restart` options to load or unload the Automatic Storage Management library driver without restarting the system: `# /etc/init.d/oracleasm restart`
createdisk	Use the `createdisk` option to mark a disk device for use with the Automatic Storage Management library driver and give it a name: `# /etc/init.d/oracleasm createdisk DISKNAME devicename`
deletedisk	Use the `deletedisk` option to unmark a named disk device: `# /etc/init.d/oracleasm deletedisk DISKNAME` **Caution:** Do not use this command to unmark disks that are being used by an Automatic Storage Management disk group. You must delete the disk from the Automatic Storage Management disk group before you unmark it.
querydisk	Use the `querydisk` option to determine if a disk device or disk name is being used by the Automatic Storage Management library driver: `# /etc/init.d/oracleasm querydisk {DISKNAME │ devicename}`
listdisks	Use the `listdisks` option to list the disk names of marked Automatic Storage Management library driver disks: `# /etc/init.d/oracleasm listdisks`
scandisks	Use the `scandisks` option to enable cluster nodes to identify which shared disks have been marked as Automatic Storage Management library driver disks on another node: `# /etc/init.d/oracleasm scandisks`

When you have completed creating and configuring Automatic Storage Management,
with ASMLIB, proceed to Chapter 6, "Installing Oracle Clusterware" on page 6-1.

5.6 Configuring Storage for Oracle Database Files on Shared Storage Devices

The following subsections describe how to configure storage for Oracle Database files.

■   Planning Your Shared Storage Device Creation Strategy

■   [Configuring Block Devices on Linux for ASM](#)

5.6.1 Planning Your Shared Storage Device Creation Strategy

Before installing the Oracle Database 11*g* release 1 (11.1) software with Oracle RAC, if you do not place files on ASM, then create enough partitions of specific sizes to support your database, and also leave a few spare partitions of the same size for future expansion. For example, if you have space on your shared disk array, then select a limited set of standard partition sizes for your entire database. Partition sizes of 50 MB, 100 MB, 500 MB, and 1 GB are suitable for most databases. Also, create a few very small and a few very large spare partitions that are (for example) 1 MB and perhaps 5 GB or greater in size. Based on your plans for using each partition, determine the placement of these spare partitions by combining different sizes on one disk, or by segmenting each disk into same-sized partitions.

> **Note:** Be aware that each instance has its own redo log files, but all instances in a cluster share the control files and data files. In addition, each instance's online redo log files must be readable by all other instances to enable recovery.
>
> In addition to the minimum required number of partitions, you should configure spare partitions. Doing this enables you to perform emergency file relocations or additions if a tablespace data file becomes full.

5.6.2 Configuring Block Devices on Linux for ASM

Use the following procedure to create block device partitions:

1. Use `fdisk` to create disk partitions on block devices for database files.

    To configure block devices and use ASM to manage data files, create one partition for each disk comprising the whole disk, and go to the procedure in the section ["Configuring Disks for Automatic Storage Management with ASMLIB"](#) on page 5-15.

2. On each node, create or modify a permissions file in `/etc/udev/permissions.d`, to change the permissions of the datafiles from default root ownership. On Asianux 2, Oracle Linux 4, and Red Hat Enterprise Linux 4, this file should be called `49-oracle.permissions`, so that the kernel loads it before `50-udev.permissions`. On Asianux Server 3, Oracle Linux 5, Red Hat Enterprise Linux 5, and SUSE Enterprise Server 10, this file should be called `51-oracle.permissions`, so that the kernel loads it after `50-udev.permissions`.

    For each partition, the contents of the *xx*-`oracle.permissions` file are as follows:

    devicepartition:*oracle_db_install_owner*:*OSDBA*:0660

    > **Note:** If you prefer to use manual undo management, instead of automatic undo management, then, instead of the `UNDOTBS`*n* shared storage devices, you must create a single rollback segment tablespace (RBS) on a shared storage device partition that is at least 500 MB in size.

3. For OUI to recognize a disk partition as an ASM disk candidate, you must mark the disk by logging in as `root` and marking the disk partitions that you created for ASM using the following command syntax, where *ASM_DISK_NAME* is the name of the ASM disk group, and *device_name* is the name of the disk device that you want to assign to that disk group:

```
/etc/init.d/oracleasm createdisk ASM_DISK_NAME device_name
```

5.7 Checking the System Setup with CVU

As the `oracle` user, use the following command syntax to start Cluster Verification Utility (CVU) stage verification to check hardware, operating system, and storage setup:

```
/mountpoint/runcluvfy.sh stage -post hwos -n node_list [-verbose]
```

In the preceding syntax example, replace the variable *node_list* with the names of the nodes in your cluster, separated by commas. For example, to check the hardware and operating system of a two-node cluster with nodes `node1` and `node2`, with the mountpoint `/mnt/dvdrom/` and with the option to limit the output to the test results, enter the following command:

```
$ /mnt/dvdrom/runcluvfy.sh stage -post hwos -n node1,node2
```

Select the option `-verbose` to receive detailed reports of the test results, and progress updates about the system checks performed by Cluster Verification Utility.

6

Installing Oracle Clusterware

This chapter describes the procedures for installing Oracle Clusterware for Linux. If you are installing 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:

- Verifying Oracle Clusterware Requirements with CVU
- Preparing to Install Oracle Clusterware with OUI
- Installing Oracle Clusterware with OUI
- Confirming Oracle Clusterware Function

6.1 Verifying Oracle Clusterware Requirements with CVU

Using the following command syntax, log in as the installation owner user (`oracle` or `crs`), and start Cluster Verification Utility (CVU) to check system requirements for installing Oracle Clusterware:

```
/mountpoint/runcluvfy.sh stage -pre crsinst -n node_list
```

In the preceding syntax example, replace the variable `mountpoint` with the installation media mountpoint, and replace the variable `node_list` with the names of the nodes in your cluster, separated by commas.

For example, for a cluster with mountpoint `/mnt/dvdrom/`, and with nodes `node1`, `node2`, and `node3`, enter the following command:

```
$ /mnt/dvdrom/runcluvfy.sh stage -pre crsinst -n node1,node2,node3
```

The Cluster Verification Utility Oracle Clusterware preinstallation stage check verifies the following:

- **Node Reachability**: All of the specified nodes are reachable from the local node.
- **User Equivalence**: Required user equivalence exists on all of the specified nodes.
- **Node Connectivity**: Connectivity exists between all the specified nodes through the public and private network interconnections, and at least one subnet exists that connects each node and contains public network interfaces that are suitable for use as virtual IPs (VIPs).
- **Administrative Privileges**: The `oracle` user has proper administrative privileges to install Oracle Clusterware on the specified nodes.
- **Shared Storage Accessibility**: If specified, the OCR device and voting disk are shared across all the specified nodes.

- **System Requirements**: All system requirements are met for installing Oracle Clusterware software, including kernel version, kernel parameters, memory, swap directory space, temporary directory space, and required users and groups.

- **Kernel Packages**: All required operating system software packages are installed.

- **Node Applications**: The virtual IP (VIP), Oracle Notification Service (ONS) and Global Service Daemon (GSD) node applications are functioning on each node.

> **Note:** Avoid changing host names after you complete the Oracle Clusterware installation, including adding or deleting domain qualifications. Nodes with changed host names must be deleted from the cluster and added back with the new name.

6.1.1 Interpreting CVU Messages About Oracle Clusterware Setup

If the Cluster Verification Utility report indicates that your system fails to meet the requirements for Oracle Clusterware 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

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 oracle user configuration to ensure that the user configuration is properly completed, and that SSH configuration is properly completed.

> **See Also:** "Creating Identical Users and Groups on Other Cluster Nodes" in Chapter 3 on page 3-3, and "Configuring SSH on All Cluster Nodes" in Chapter 2 on page 2-30 for user equivalency configuration instructions

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 node_name date
```

The output from this command should be the timestamp of the remote node identified by the value that you use for node_name. 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 have not attempted to use SSH to connect to the host node before running, then Cluster Verification Utility indicates a user equivalence error. 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 again 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 `$CV_HOME/cv/admin`, open the file `cvu_config` with a text editor, and add or update the key `ORACLE_SRVM_REMOTESHELL` to indicate the ssh path location on your system. For example:

```
# Locations for ssh and scp commands
ORACLE_SRVM_REMOTESHELL=/usr/local/bin/ssh
ORACLE_SRVM_REMOTECOPY=/usr/local/bin/scp
```

Note the following rules for modifying the `cvu_config` file:

- Key entries have the syntax *name=value*
- Each key entry and the value assigned to the key defines one property only
- Lines beginning with the number sign (#) are comment lines, and are ignored
- Lines that do not follow the syntax *name=value* are ignored

When you have changed the path configuration, run Cluster Verification Utility again. If ssh is in another location than the default, you also need to start OUI with additional arguments to specify a different location for the remote shell and remote copy commands. Enter `runInstaller -help` to obtain information about how to use these arguments.

Note: When you or OUI run `ssh` or `rsh` commands, including any login or other shell scripts they start, you may see errors about invalid arguments or standard input if the scripts generate any output. You should correct the cause of these errors.

To stop the errors, remove all commands from the `oracle` user's login scripts that generate output when you run `ssh` or `rsh` commands.

If you see messages about X11 forwarding, then complete the task "Setting Display and X11 Forwarding Configuration" on page 2-35 to resolve this issue.

If you see errors similar to the following:

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

These errors are produced if hidden files on the system (for example, `.bashrc` or `.cshrc`) contain `stty` commands. If you see these errors, then refer to Chapter 2, "Preventing Oracle Clusterware Installation Errors Caused by stty Commands" on page 2-35 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 `oracle` user is created with the correct group membership. Ensure that you have created the required groups, and create or modify the user account on affected nodes to establish required group membership.

> **See Also:** "Creating Standard Configuration Operating System Groups and Users" in Chapter 3 for instructions about how to create required groups, and how to configure the `oracle` user

6.2 Preparing to Install Oracle Clusterware with OUI

Before you install Oracle Clusterware with Oracle Universal Installer (OUI), 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 to install Oracle Clusterware. Mark the box for each task as you complete it, and write down the information needed, so that you can provide it during installation.

❏ **Shut Down Running Oracle Processes**

If you are installing Oracle Clusterware on a node that already has a single-instance Oracle Database 11*g* release 1 (11.1) installation, then stop the existing ASM instances. After Oracle Clusterware is installed, start up the ASM instances again. When you restart the single-instance Oracle database, the ASM instances use the Cluster Synchronization Services (CSSD) Daemon from Oracle Clusterware instead of the CSSDdaemon for the single-instance Oracle database.

You can upgrade some or all nodes of an existing Cluster Ready Services installation. For example, if you have a six-node cluster, then you can upgrade two nodes each in three upgrading sessions.Base the number of nodes that you upgrade in each session on the load the remaining nodes can handle. This is called a "rolling upgrade."

If a Global Services Daemon (GSD) from Oracle9*i* Release 9.2 or earlier is running, then stop it before installing Oracle Database 11*g* release 1 (11.1) Oracle Clusterware 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 Oracle9*i* 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 Clusterware 11*g* release 1 (11.1) software from detecting the Oracle9*i* release 2 nodelist, and causes failure of the Oracle Clusterware 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 home that is running CSS.

❏ **Prepare for Clusterware Upgrade If You Have Existing Oracle Cluster Ready Services Software**

During an Oracle Clusterware installation, if OUI detects an existing Oracle Database 10*g* release 1 (10.1) Cluster Ready Services (CRS), then you are given the option to perform a rolling upgrade by installing Oracle Database 11*g* release 1 (11.1) Oracle Clusterware on a subset of cluster member nodes.

If you intend to perform a rolling upgrade, then you should shut down the CRS stack on the nodes you intend to upgrade, and unlock the Oracle Clusterware home using the script *mountpoint*/clusterware/upgrade/preupdate.sh, which is available on the 11*g* release 1 (11.1) installation media.

If you intend to perform a standard upgrade, then shut down the CRS stack on all nodes, and unlock the Oracle Clusterware home using the script *mountpoint*/clusterware/upgrade/preupdate.sh.

When you run OUI and select the option to install Oracle Clusterware on a subset of nodes, OUI installs Oracle Database 11*g* release 1 (11.1) Oracle Clusterware software into the existing Oracle Clusterware home on the local and remote node subset. When you run the root script, it starts the Oracle Clusterware 11*g* release 1 (11.1) stack on the subset cluster nodes, but lists it as an inactive version.

When all member nodes of the cluster are running Oracle Clusterware 11*g* release 1 (11.1), then the new clusterware becomes the active version.

If you intend to install Oracle RAC, then you must first complete the upgrade to Oracle Clusterware 11*g* release 1 (11.1) on all cluster member nodes before you install the Oracle Database 11*g* release 1 (11.1) version of Oracle RAC.

❏ **Determine the Oracle Inventory location**

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

If you are installing Oracle software for the first time on your system, and your system does not have an Oracle inventory, then you are asked to provide a path for the Oracle inventory, and you are also asked the name of the Oracle Inventory group (typically, oinstall).

> **See Also:** The preinstallation chapters in Chapter 2 for information about creating the Oracle Inventory, and completing required system configuration

❏ **Obtain root account access**

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

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

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

> **Note:** If the language set for the operating system is not supported by Oracle Universal Installer, then Oracle Universal Installer, by default, 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, private node names, and virtual node names for each node in the cluster**

If you install the clusterware during installation, and are not using third-party vendor clusterware, then you are asked to provide a public node name and a private node name for each node.

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. This node name can be either the permanent or the virtual host name.

In addition, ensure that the following are true:

– Determine 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: underscores (_), 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.

– Determine a private node name or private IP address for each node. The private IP address is an address that is accessible only by the other nodes in this cluster. Oracle Database uses private IP addresses for internode, or instance-to-instance Cache Fusion traffic. Oracle recommends that you provide a name in the format *public_hostname*-priv. For example: `myclstr2-priv`.

– 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 *public_hostname*-vip. For example: `myclstr2-vip`.

> **Note:** The following is a list of additional information about node IP addresses:
>
> ■ For the local node only, OUI automatically fills in public, private, and VIP fields. If your system uses vendor clusterware, then OUI may fill additional fields.
>
> ■ Host names, private 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.
>
> ■ Private IP addresses should not be accessible as public interfaces. Using public interfaces for Cache Fusion can cause performance problems.

❏ **Identify shared storage for Oracle Clusterware files and prepare disk partitions if necessary**

During installation, you are asked to provide paths for two files that must be shared across all nodes of the cluster, either on a shared raw device, or a shared file system file:

– The voting disk is a partition that Oracle Clusterware uses to verify cluster node membership and status.

The voting disk must be owned by the user performing the installation (`oracle` or `crs`), and must have permissions set to 640.

– The Oracle Cluster Registry (OCR) contains cluster and database configuration information for the Oracle RAC database and for Oracle Clusterware, including the node list, and other information about cluster configuration and profiles.

The OCR disk must be owned by the user performing the installation (`crs` or `oracle`. That installation user must have `oinstall` as its primary group. The OCR disk partitions must have permissions set to 640, though permissions files used with system restarts should have ownership set to `root:oinstall`. During installation, OUI changes ownership of the OCR disk partitions to root. Provide at least 280 MB disk space for the OCR partitions.

If your disks do not have external storage redundancy, then Oracle recommends that you provide one additional location for the OCR disk, and two additional locations for the voting disk, for a total of five partitions (two for OCR, and three for voting disks). Creating redundant storage locations protects the OCR and voting disk in the event of a disk failure on the partitions you choose for the OCR and the voting disk.

> **See Also:** Chapter 2, "Oracle Clusterware Preinstallation Tasks"

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

If you run OUI during the time that daily cron jobs are running, 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.

6.3 Installing Oracle Clusterware with OUI

This section provides you with information about how to use Oracle Universal Installer (OUI) to install Oracle Clusterware. It contains the following sections:

- Running OUI to Install Oracle Clusterware

- Installing Oracle Clusterware Using a Cluster Configuration File

- Troubleshooting OUI Error Messages for Oracle Clusterware

6.3.1 Running OUI to Install Oracle Clusterware

Complete the following steps to install Oracle Clusterware 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. Unless you have the same terminal window open that you used to set up SSH, enter the following commands:

```
$ exec /usr/bin/ssh-agent $SHELL
$ /usr/bin/ssh-add
```

2. Start the `runInstaller` command from the `/Disk1` directory on the Oracle Database 11*g* release 1 (11.1) installation media. For example:

```
$ cd /Disk1
./runInstaller
```

3. Provide information or run scripts as root when prompted by OUI. If you need assistance during installation, click **Help**.

> **Note:** You must run `root.sh` scripts one at a time. Do not run `root.sh` scripts simultaneously.

4. After you run `root.sh` on all the nodes, OUI runs the Oracle Notification Server Configuration Assistant, Oracle Private Interconnect Configuration Assistant, and Cluster Verification Utility. These programs run without user intervention.

When you have verified that your Oracle Clusterware 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 11*g* release 1 (11.1) with Oracle RAC, then refer to *Oracle Real Application Clusters Installation Guide for Linux*. If you intend to use Oracle Clusterware by itself, then refer to the single-instance Oracle Database installation guide.

> **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 Clusterware

6.3.2 Installing Oracle Clusterware Using a Cluster Configuration File

During installation of Oracle Clusterware, on the Specify Cluster Configuration page, you are given the option either of providing cluster configuration information manually, or of using a cluster configuration file. A cluster configuration file is a text file that you can create before starting OUI, which provides OUI with information about the cluster name and node names that it requires to configure the cluster.

Oracle suggests that you consider using a cluster configuration file if you intend to perform repeated installations on a test cluster, or if you intend to perform an installation on many nodes.

To create a cluster configuration file:

1. On the installation media, navigate to the directory `Disk1/response`.

2. Using a text editor, open the response file `crs.rsp`, and find the section `CLUSTER_CONFIGURATION_FILE`.

3. Follow the directions in that section for creating a cluster configuration file.

6.3.3 Troubleshooting OUI Error Messages for Oracle Clusterware

The following is a list of some common Oracle Clusterware installation issues, and how to resolve them.

PRKC-1044 Failed to check remote command execution

Cause: SSH keys need to be loaded into memory, or there is a user equivalence error.

Action: Run the following commands to load SSH keys into memory:

```
$ exec /usr/bin/ssh-agent $SHELL
$ /usr/bin/ssh-add
```

Note that you must have the passphrase used to set up SSH. If you are not the person who set up SSH, then obtain the passphrase. Note also that the .ssh folder in the user home that is performing the installation must be set with 600 permissions.

In addition, confirm group membership by entering the id command, and entering ID username. For example:

```
$ id
$ id oracle
```

Incorrect permissions on partitions used for OCR or Voting Disks

Cause: The user account performing the installation (oracle or crs) does not have permission to write to these partitions

Action: Make the partitions writable by the user performing installation. For example, use the command chown *user* to make the selected partitions writable by the user (`oracle` or `crs`) performing the installation. During installation, these permissions are changed to root ownership.

6.4 Confirming Oracle Clusterware Function

After installation, log in as root, and use the following command syntax to confirm that your Oracle Clusterware installation is installed and running correctly:

```
CRS_home/bin/crs_stat -t -v
```

For example:

```
[root@node1 /]:/u01/app/crs/bin/crs_stat -t -v
Name        a Type        R/RA   F/FT    Target    State     Host
crs....ac3.gsd application 0/5   0/0     Online    Online    node1
crs....ac3.ons application 0/5   0/0     Online    Online    node1
crs....ac3.vip application 0/5   0/0     Online    Online    node1
crs....ac3.gsd application 0/5   0/0     Online    Online    node2
crs....ac3.ons application 0/5   0/0     Online    Online    node2
crs....ac3.vip application 0/5   0/0     Online    Online    node2
```

You can also use the command crsctl check crs for a less detailed system check. for example:

```
[root@node1 bin] $ ./crsctl check crs
Cluster Synchronization Services appears healthy
Cluster Ready Services appears healthy
Event Manager appears healthy
```

> **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.

7

Oracle Clusterware Postinstallation Procedures

This chapter describes how to complete the postinstallation tasks after you have installed the Oracle Clusterware software.

This chapter contains the following topics:

- Required Postinstallation Tasks
- Recommended Postinstallation Tasks

7.1 Required Postinstallation Tasks

You must perform the following tasks after completing your installation:

- Back Up the Voting Disk After Installation
- Configure Input/Output Fencing
- Download and Install Patch Updates

7.1.1 Back Up the Voting Disk After Installation

After your Oracle Clusterware installation is complete and after you are sure that your system is functioning properly, make a backup of the contents of the voting disk. Use the dd utility. For example:

```
# dd if=/dev/sda1 of=/dev/myvdisk1.bak
```

Also, make a backup copy of the voting disk contents after you complete any node additions or node deletions, and after running any deinstallation procedures.

7.1.2 Configure Input/Output Fencing

Review the following sections to configure input/output fencing:

- About Input/Output Fencing
- Configuring the Hangcheck-timer Module
- Configuring Oracle Clusterware Process Monitor Daemon

7.1.2.1 About Input/Output Fencing

Input/output fencing (IO Fencing) is a required mechanism to ensure that a node that is evicted from the cluster is prevented from writing to the cluster shared storage. The mechanisms that provide fencing are the hangcheck-timer kernel module, and the

Oracle Clusterware Process Monitor Daemon (`oprocd`). The `oprocd` process is installed during Oracle Clusterware installation. Both mechanisms are required.

The hangcheck-timer and `oprocd` are independent mechanisms. The `oprocd` process provides additional hang check capability, and it can catch hang conditions that the hangcheck-timer misses.

7.1.2.2 Configuring the Hangcheck-timer Module

Load the hangcheck-timer kernel module as root, using `insmod` or `modprobe`. The examples in this section show `modprobe` commands.

Check the following hangcheck timer settings:

- `hangcheck_tick` parameter: This parameter defines how often, in seconds, the `hangcheck-timer` checks the node for hangs. The default value is 60 seconds. Oracle recommends that you change the value of `hangcheck_tick` to 1.

- `hangcheck_margin` parameter: This parameter defines how long the timer waits, in seconds, for a response from the kernel. The default value is 180 seconds. Oracle recommends that you change the value of `hangcheck_margin` to 10.

- The `hangcheck_reboot` parameter determines if the hangcheck-timer restarts the node if the kernel fails to respond within the sum of the `hangcheck_tick` and `hangcheck_margin` parameter values. If the value of hangcheck_reboot is equal to or greater than 1, then the `hangcheck-timer` module restarts the system when a hang is detected. If the hangcheck_reboot parameter is set to zero, then the hangcheck-timer module does not restart the node when a hang is detected.

> **Note:** On Linux 2.6 kernels, by default `hangcheck_reboot` is set to 0. The value for `hangcheck_reboot` must always be set to 1 to restart the system if a hang is detected.

These settings assume that the CSS misscount value is set to 30 or 60 seconds, which are the default for release 11*g* and release 10*g* respectively. The value for CSS misscount should always be greater than the sum of `hangcheck_tick` and `hangcheck_margin`.

For optimal cluster performance, test applications with the hangcheck parameter values that Oracle recommends. If you find that the cluster produces false node evictions with these values, then increase the `hangcheck_margin` parameter value, with the help of Oracle Support.

Use the following procedure to configure the hangcheck timer:

1. Log in as root.

2. Check to see if settings for the hangcheck timer are listed in the module configuration file. On Red Hat and on Oracle Linux, that file is `/etc/modprobe.conf`. On SUSE, it is `/etc/modprobe.conf.local`. For example:

   ```
   # more /etc/modprobe.conf |grep hang
   ```

   You should see something similar to the following:

   ```
   options hangcheck-timer hangcheck_tick=1 hangcheck_margin=10 hangcheck_reboot=1
   ```

3. If the hangcheck configuration does not exist, or if it exists but the values are set to different values than those recommended, then enter a command similar to the

following to load them into the configuration file. The following example is for Red Hat and Oracle Linux:

```
# echo "options hangcheck-timer hangcheck_tick=1 hangcheck_margin=10 \
hangcheck_reboot=1" >>/etc/modprobe.conf
```

4. If necessary, enter the following command to remove the existing hangcheck-timer values:

```
# /sbin/modprobe -r hangcheck-timer
```

5. Enter the following command to load the new hangcheck-timer values:

```
# /sbin/modprobe -v hangcheck-timer
```

6. To confirm that the hangcheck module is loaded, enter the following command:

```
# /sbin/lsmod | grep hang
```

The output should be similar to the following:

```
 hangcheck_timer       3289  0
```

7. To ensure that the module is loaded every time the system restarts, verify that the local system startup file contains the command `/sbin/modprobe -v hangcheck-timer`, or add it if necessary:

    -   Red Hat:

        Add the command to the `/etc/rc.d/rc.local` file.

    -   SUSE:

        Add the command to the `/etc/init.d/boot.local` file.

8. Repeat this process on each node that you intend to make a member of the cluster.

7.1.2.3 Configuring Oracle Clusterware Process Monitor Daemon

The Oracle Clusterware Process Monitor Daemon (`oprocd`) process is part of the Oracle Clusterware software installation. It is started automatically by Oracle Clusterware to detect system hangs. When it detects a system hang, it restarts the hung node.

Oracle has found wide variations in scheduling latencies observed across operating systems and versions of operating systems. Because of these scheduling latencies, the default values for `oprocd` can be overly sensitive, particularly under heavy system load, resulting in unnecessary `oprocd`-initiated restarts (false restarts).

Oracle recommends that you address scheduling latencies with your operating system vendor to reduce or eliminate them as much as possible, as they can cause other problems.

To overcome these scheduling latencies, Oracle recommends that you set the Oracle Clusterware parameter `diagwait` to the value `13`. This setting increases the time for failed nodes to flush final trace files, which helps to debug the cause of a node failure. You must shut down the cluster to change the `diagwait` setting.

If you require more aggressive failover times to meet more stringent service level requirements, then you should open a service request with Oracle Support to receive advice about how to tune for lower failover settings.

> **Note:** Changing the `diagwait` parameter requires a clusterwide shutdown. Oracle recommends that you change the `diagwait` setting either immediately after the initial installation, or during a scheduled outage.

To change the `diagwait` setting:

1. Log in as root, and run the following command on all nodes, where *CRS_home* is the home directory of the Oracle Clusterware installation:

   ```
   # CRS_home/bin/crsctl stop crs
   ```

2. Enter the following command, where *CRS_home* is the Oracle Clusterware home:

   ```
   # CRS_home/bin/oprocd stop
   ```

   Repeat this command on all nodes.

3. From one node of the cluster, change the value of the `diagwait` parameter to 13 seconds by issuing the following command as `root`:

   ```
   # CRS_home/bin/crsctl set css diagwait 13 -force
   ```

4. Restart the Oracle Clusterware by running the following command on all nodes:

   ```
   # CRS_home/bin/crsctl start crs
   ```

5. Run the following command to ensure that Oracle Clusterware is functioning properly:

   ```
   # CRS_home/bin/crsctl check crs
   ```

7.1.3 Download and Install Patch Updates

Refer to the Oracle*MetaLink* Web site for required patch updates for your installation. To download required patch updates:

1. Use a Web browser to view the Oracle*MetaLink* Web site:

   https://metalink.oracle.com

2. Log in to Oracle*MetaLink*.

   > **Note:** If you are not an Oracle*MetaLink* registered user, then click **Register for MetaLink** and register.

3. On the main Oracle*MetaLink* 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 number of the patch that you want to download.

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 10*g* to uncompress the Oracle patch updates that you downloaded from Oracle*MetaLink*. The unzip utility is located in the `$ORACLE_HOME/bin` directory.

13. Refer to Appendix B on page B-1 for information about how to stop database processes in preparation for installing patches.

7.2 Recommended Postinstallation Tasks

Oracle recommends that you complete the following tasks after installing Oracle Clusterware.

7.2.1 Back Up the root.sh Script

Oracle recommends that you back up the `root.sh` script after you complete an installation. If you install other products in the same Oracle home directory, then the Oracle Universal Installer (OUI) 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.

7.2.2 Run CVU Postinstallation Check

After installing Oracle Clusterware, check the status of your Oracle Clusterware installation with the command `cluvfy stage -post crsinst`, using the following syntax:

```
cluvfy stage -post crsinst -n node_list [-verbose]
```

8

Deinstallation of Oracle Clusterware

This chapter describes how to remove Oracle Clusterware.

This chapter contains the following topics:

- Deciding When to Deinstall Oracle Clusterware
- Relocating Single-instance ASM to a Single-Instance Database Home
- Removing Oracle Clusterware

> **See Also:** Product-specific documentation for requirements and restrictions, if you want to remove an individual product

8.1 Deciding When to Deinstall Oracle Clusterware

Remove installed components in the following situations:

- You have encountered errors during or after installing or upgrading Oracle Clusterware, and you want to re-attempt 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.

- You have successfully installed Oracle Clusterware, and you want to remove the Clusterware installation, either in an educational environment, or a test environment.

- You have successfully installed Oracle Clusterware, but you want to downgrade to a previous release.

8.2 Relocating Single-instance ASM to a Single-Instance Database Home

If you have a single-instance Oracle Database on Oracle Clusterware, and you want to remove Oracle Clusterware, then use the following syntax to add the local CSS configuration to the ASM home:

ASM_home/bin/localconfig add

For example:

```
$ cd /u01/app/asm/bin/
$ ./localconfig add
```

8.3 Removing Oracle Clusterware

The scripts `rootdelete.sh` and `rootdeinstall.sh` remove Oracle Clusterware from your system. After running these scripts, run Oracle Universal Installer to remove the Oracle Clusterware home. The following sections describe the scripts, and later, provide exact procedure to the removal of the Oracle Clusterware software.

8.3.1 About the rootdelete.sh Script

The `rootdelete.sh` script should be run from the Oracle Clusterware home on each node. It stops the Oracle Clusterware stack, removes `inittab` entries, and deletes some of the Oracle Clusterware files. It can also be used to downgrade the Oracle Cluster Registry from the existing release to a previous release. The script uses the following syntax:

```
# rootdelete.sh options
```

Options:

- **paramfile**: Use a parameter file containing configuration information for the `rootdelete.sh` command. Provide the path and name of the parameter file. For example: `-paramfile /usr/oracle/cwdeletepar`.

- **local | remote**: Use **local** if you are running `rootdelete.sh` on the local node, and use **remote** if you are running the script on one of the other nodes. The local node is the one from which you run OUI (in other words, the last surviving node), and on which you run `rootdeinstall.sh`.

- **nosharedvar | sharedvar**: Use **nosharedvar** if the directory path for `ocr.loc` (in `/etc/oracle` or `/var/opt/oracle`) is not on a shared file system. Use **sharedvar** if the directory path for `ocr.loc` is in a shared location. The default is **nosharedvar**.

- **sharedhome | nosharedhome**: Use **sharedhome** if the Oracle Clusterware home is shared across the nodes. Otherwise, use **nosharedhome**. The default is **sharedhome**.

- **downgrade**: Use this option if the Oracle Clusterware is downgraded to a previous Oracle Clusterware version. The **-downgrade** option takes the following flags:

  - **-version**: Use this option to specify the version to which you want to downgrade. The default is 10.2.

  - **-force**: Use this option to force cleanup of root configuration

For example, to run the `rootdelete.sh` script from an Oracle Clusterware home in the path `/u01/app/crs`, where you are running the script on a remote node, and the ocr.loc file is in `/etc/oracle` on each node, enter the following command:

```
# cd /u01/app/crs/install/
# ./rootdelete.sh remote nosharedvar
```

8.3.2 Example of the rootdelete.sh Parameter File

You can create a parameter file for `rootdelete.sh` to repeat deinstallation steps. You may want to do this if you intend to perform repeated reinstallations, as in a test environment. The following is an example of a parameter file for `rootdelete.sh`; terms that change relative to system configuration are indicated with italics:

```
CLUSTER_NODES=mynode1,mynode2
INVENTORY_LOCATION=u01/app/oracle/oraInventory
```

```
CRS_HOME=true
ORA_CRS_HOME=/u01/app/crs
ORACLE_OWNER=oracle
DBA_GROUP=oinstall
```

8.3.3 About the rootdeinstall.sh Script

The `rootdeinstall.sh` script should be run on the local node only, after rootdelete.sh has been run on all nodes of the cluster. Use this command either to remove the Oracle Clusterware OCR file, or to downgrade your existing installation. The `rootdeinstall.sh` script has the following command options:

- **paramfile**: A parameter file containing configuration information for the rootdelete.sh command

- **downgrade**: Use this option if the database is downgraded to a previous Oracle Clusterware version. Use the **-version** flag to specify the version to which you want to downgrade. The default is 10.2.

8.3.4 Removing Oracle Clusterware

Complete the following procedure to remove Oracle Clusterware:

1. Log in as the `oracle` user, and shut down any existing Oracle Database instances on each node, with normal or immediate priority. For example:

   ```
   $ Oracle_home/bin/srvctl stop database -d db_name
   $ Oracle_home/bin/srvctl stop asm -n node
   $ Oracle_home/bin/srvctl stop nodeapps -n node
   ```

2. Use Database Configuration Assistant and NETCA to remove listeners, Automatic Storage Management instances, and databases from the system. This removes the Oracle Clusterware resources associated with the listeners, Automatic Storage Management instances, and databases on the cluster.

3. On each remote node, log in as the `root` user, change directory to the Oracle Clusterware home, and run the `rootdelete` script with the options `remote nosharedvar nosharedhome`. For example:

   ```
   [root@node2 /] # cd /u01/app/crs/install
   [root@node2 /install] # ./rootdelete.sh remote nosharedvar nosharedhome
   ```

4. On the local node (the last cluster member node), log in as the `root` user, change directory to the Oracle Clusterware home, and run the `rootdelete` script with the options `local nosharedvar nosharedhome`. For example:

   ```
   [root@node1 /] # cd /u01/app/crs/install
   [root@node1 /install] # ./rootdelete.sh local nosharedvar nosharedhome
   ```

5. On the local node, run the script `rootdeinstall`. For example:

   ```
   [root@node1 install]# ./rootdeinstall.sh
   ```

6. Log in as the `oracle` user, and run Oracle Universal Installer to remove the Oracle Clusterware home. For example

   ```
   $ cd /u01/app/crs/oui/bin
   $ ./runInstaller -deinstall -removeallfiles
   ```

A

Troubleshooting the Oracle Clusterware Installation Process

This appendix provides troubleshooting information for installing Oracle Clusterware.

> **See Also:** The Oracle Database 11*g* Oracle RAC documentation set included with the installation media in the Documentation directory:
>
> - *Oracle Clusterware Administration and Deployment Guide*
> - *Oracle Real Application Clusters Administration and Deployment Guide*

This appendix contains the following topics:

- Install OS Watcher and RACDDT
- General Installation Issues
- Missing Operating System Packages On Linux
- Performing Cluster Diagnostics During Oracle Clusterware Installations
- Interconnect Errors

A.1 Install OS Watcher and RACDDT

To address troubleshooting issues, Oracle recommends that you install OS Watcher, and if you intend to install an Oracle RAC database, RACDDT. You must have access to Oracle*MetaLink* to download OS Watcher and RACDDT.

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

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

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

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

https://metalink.oracle.com

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

A.2 General Installation Issues

The following is a list of examples of types of errors that can occur during installation. It contains the following issues:

- An error occurred while trying to get the disks
- Failed to connect to server, Connection refused by server, or Can't open display
- MEMORY_TARGET not supported on this system
- Nodes unavailable for selection from the OUI Node Selection screen
- Node nodename is unreachable
- PROT-8: Failed to import data from specified file to the cluster registry
- Time stamp is in the future
- YPBINDPROC_DOMAIN: Domain not bound

An error occurred while trying to get the disks

Cause: There is an entry in /etc/oratab pointing to a non-existent Oracle home. The OUI error file should show the following error: "java.io.IOException: /home/oracle/OraHome//bin/kfod: not found" (Oracle*Metalink* bulletin 276454.1)

Action: Remove the entry in /etc/oratab pointing to a non-existing Oracle home.

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.

MEMORY_TARGET not supported on this system

Cause: On Linux systems, insufficient `/dev/shm` size for PGA and SGA.

If you are installing on a Linux system, note that **Memory Size (SGA and PGA)**, which sets the initialization parameter `MEMORY_TARGET` or `MEMORY_MAX_TARGET`, cannot be greater than the shared memory file system (`/dev/shm`) on your operating system.

Action: Increase the `/dev/shm` mountpoint size. For example:

```
# mount -t tmpfs shmfs -o size=4g /dev/shm
```

Also, to make this change persistent across system restarts, add an entry in `/etc/fstab` similar to the following:

```
shmfs /dev/shm tmpfs size=4g 0
```

Nodes unavailable for selection from the OUI Node Selection screen

Cause: Oracle 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. Review the section "Configuring SSH on All Cluster Nodes" on page 2-30.

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

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

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.

YPBINDPROC_DOMAIN: Domain not bound

Cause: This error can occur during postinstallation testing when a node public network interconnect is pulled out, and the VIP does not fail over. Instead, the

node hangs, and users are unable to log in to the system. This error occurs when the Oracle home, listener.ora, Oracle log files, or any action scripts are located on an NAS device or NFS mount, and the name service cache daemon `nscd` has not been activated.

Action: Enter the following command on all nodes in the cluster to start the nscd service:

```
/sbin/service  nscd start
```

A.3 Missing Operating System Packages On Linux

You have missing operating system packages on your system if you receive error messages such as the following during Oracle Clusterware, Oracle RAC, or Oracle Database installation:

```
libstdc++.so.5: cannot open shared object file: No such file or directory
libXp.so.6: cannot open shared object file: No such file or directory
```

Typically, errors such as these occur if you have not fully checked required operating system packages during preinstallation, and failed to confirm that all required packages were installed. Run Cluster Verification Utility (CVU), either from the shiphome mount point (`runcluvfy.sh`), or from an installation directory (*CRS_home*`/bin`). Cluster Verification Utility reports which required packages are missing.

If you have a Linux support network configured, such as the Red Hat network or Oracle Unbreakable Linux support, then you can also use the `up2date` command to determine the name of the package. For example:

```
# up2date --whatprovides libstdc++.so.5
compat-libstdc++-33.3.2.3-47.3
```

A.4 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.

> **Note:** If you encounter unexplained installation errors during or
> after a period when cron jobs are run, then your cron job may have
> deleted temporary files before the installation is finished. Oracle
> recommends that you complete installation before daily cron jobs are
> run, or disable daily cron jobs that perform cleanup until after the
> installation is completed.

A.5 Interconnect Errors

If you use more than one NIC for the interconnect, then you must use NIC bonding, or
the interconnect will fail.

If you install Oracle Clusterware and Oracle RAC, then they must use the same NIC or
bonded NIC cards for the interconnect.

If you use bonded NIC cards, then they must be on the same subnet.

B

How to Perform Oracle Clusterware Rolling Upgrades

This appendix describes how to perform Oracle Clusterware rolling upgrades. Because you must stop database processes on the nodes you intend to upgrade when you perform an Oracle Clusterware upgrade, it includes information about how to stop processes in Oracle Real Application Clusters (Oracle RAC) databases.

The instructions in this section specify a single node, and assume that you are upgrading one node at a time. To upgrade a subset of nodes together, you can specify a list of nodes (the subset), where the example commands specify a single node. For example, instead of `-n node`, specify `-n node1,node2,node3`.

This appendix contains the following topics:

> **Note:** You can use the procedures in this chapter to perform rolling upgrades of Oracle Clusterware from any Oracle Clusterware 10*g* or Oracle Clusterware 11*g* installation to the latest patchset update. For example, you can use these procedures to prepare to upgrade from Oracle Clusterware 10.2.0.1 to 10.2.0.3.

- Back Up the Oracle Software Before Upgrades
- Restrictions for Clusterware Upgrades to Oracle Clusterware 11g
- Verify System Readiness for Patchset and Release Upgrades
- Installing a Patch Set On a Subset of Nodes
- Installing an Upgrade On a Subset of Nodes

B.1 Back Up the Oracle Software Before Upgrades

Before you make any changes to the Oracle software, whether you intend to upgrade or patch part of the database or clusterware, or all of your cluster installation, Oracle recommends that you create a backup of the Oracle software.

B.2 Restrictions for Clusterware Upgrades to Oracle Clusterware 11*g*

To upgrade existing Oracle Clusterware and Cluster Ready Services installations to Oracle Clusterware 11*g*, you must first upgrade the existing installations to a minimum patch level. The minimum patch level is listed in the following table:

Table B–1   Minimum Oracle Clusterware Patch Levels Required for Rolling Upgrades to 11g

Oracle Clusterware Release	Minimum Patch Level Required
10*g* Release 2	10.2.0.3, or 10.2.0.2 with CRS bundle # 2 (Patch 526865)
10*g* Release 1	10.1.0.3

To upgrade your Oracle Clusterware installation to the minimum patch level using a rolling upgrade, follow the directions in the Patch Readme file.

> **Note:**   You can use the procedures in this chapter to prepare to perform rolling upgrades of Oracle Clusterware from any Oracle Clusterware 10g release 10.2 or Oracle Clusterware 11g installation to the latest patchset update. For example, you can use these procedures to prepare to upgrade from Oracle Clusterware 10.2.0.1 to 10.2.0.3.

> **See Also:**   *Oracle Database Upgrade Guide* for additional information about upgrades, and check the following site on the Oracle Technology Network for relevant information about rolling upgrades:
>
> http://www.oracle.com/technology/deploy/availability
>
> My Oracle Support also has an Upgrade Companion for each release that provide additional upgrade information. It is available at the following URL:
>
> https://metalink.oracle.com/

B.3 Verify System Readiness for Patchset and Release Upgrades

If you are completing a patchset update of your database or clusterware, then after you download the patch software, and before you start to patch or upgrade your database, review the Patch Set Release Notes that accompany the patch to determine if your system meets the system requirements for the operating system and the hardware platform.

Use the following Cluster Verification Utility (CVU) command to assist you with system checks in preparation for starting a database patch or upgrade, where *node* is the node, or comma-delimited subset of nodes, that you want to check, and *inventory_group* is the name of the Oracle Inventory group:

cluvfy stage -pre crsinst -n *node* -orainv *inventory_group*

For example, to perform a system check on nodes node1, node2 and node3, and where the Oracle Inventory group is oinstall, enter the following command:

```
$ cluvfy stage -pre crsinst -n node1,node2,node3 -orainv oinstall
```

> **Note:**   Before you start an upgrade, Oracle recommends that you download the latest Cluster Verification Utility version from Oracle Technology Network at the following URL:
>
> http://otn.oracle.com

B.4 Installing a Patch Set On a Subset of Nodes

To patch Oracle Clusterware, review the instructions in the Patch Set README document for additional instructions specific to the patchset.

Before you shut down any processes that are monitored by Enterprise Manager Grid Control, set a blackout in Grid Control for the processes that you intend to shut down. This is necessary so that the availability records for these processes indicate that the shutdown was planned downtime, rather than an unplanned system outage.

To patch a subset of nodes, complete the following steps:

> **Note:**   You must perform these steps in the order listed.

1. Change directory to the Oracle Clusterware home. As `root`, run the `preupdate.sh` script on the local node, and on all other nodes in the subset that you intend to upgrade. Use the following command syntax, where *clusterware_home* is the path to the existing Oracle Clusterware home, and *installation_owner* is the Oracle Clusterware installation owner:

    ./preupdate.sh -crshome *clusterware_home* -crsuser *installation_owner*

    For example:

    ```
    # cd $ORACLE_HOME/install
    # ./preupdate.sh -crshome /opt/crs -crsuser oracle
    ```

    The script output should be similar to the following:

    ```
    Shutting down Oracle Cluster Ready Services (CRS):
    Stopping resources. This could take several minutes.
    Successfully stopped CRS resources.
    Stopping CSSD.
    Shutting down CSS daemon.
    Shutdown request successfully issued.
    Shutdown has begun. The daemons should exit soon.
    Checking to see if Oracle CRS stack is down...
    Oracle CRS stack is down now.
    ```

2. Confirm that you are logged in as the Oracle Clusterware installation owner, and start Oracle Universal Installer to install the software.

    For example:

    ```
    $ whoami
    crs
    $ cd /cdrom/clusterware/
    ./runInstaller
    ```

    Provide information as prompted by the Installer.

    > **Note:**   You cannot change the owner of the Oracle Clusterware home.

3. During an Oracle Clusterware installation, if Oracle Universal Installer detects an existing Oracle Clusterware 10*g* release 1 or release 2 installation, then you are given the option to perform a rolling upgrade by installing the patch on a subset of cluster member nodes.

You can patch the entire cluster, and then run the `root.sh` patch script in a rolling fashion to make the patch update active.

4. After you select the nodes you want to upgrade, the Installer installs the patch software in the existing Oracle Clusterware home on the local and the remote subset of nodes you have selected.

   OUI prompts you to run the appropriate root script for the patchset. The script starts the Oracle Clusterware stack on the upgraded subset of nodes. However, it lists it as an inactive version.

5. After you upgrade the initial node or subset of nodes, repeat steps 1 through 4 for each remaining node or subset of nodes until all nodes in the cluster have been patched, and the new version is the active version.

   When all member nodes of the cluster are running the new Oracle Clusterware release, then the new clusterware becomes the active version. Otherwise, the older Oracle Clusterware release is still used.

   To list the version of Oracle Clusterware that is installed on a node, enter the following command, where *CRShome* is the Oracle Clusterware home, and *nodename* is the name of the node:

   ```
   # CRShome/bin/crsctl query crs softwareversion [nodename]
   ```

   To list the Oracle Clusterware software version that is running on a node, enter the following command, where *CRShome* is the Oracle Clusterware home:

   ```
   # CRShome/bin/crsctl query crs activeversion
   ```

   If you intend to install or upgrade Oracle RAC, then you must first complete the upgrade to Oracle Clusterware 11*g* release 1 (11.1) on all cluster member nodes before you install the Oracle Database 11*g* release 1 (11.1) version of Oracle RAC.

6. Check with Oracle Support to confirm you have installed any recommended patch sets, bundle patches or critical patches.

   To check for the latest recommended patches for Oracle Clusterware and Oracle Real Application Clusters, log on to the following site:

   https://metalink2.oracle.com

   Click **Patches & Updates**, and click Oracle Database from the Recommended Patches list. Provide information as prompted.

B.5 Installing an Upgrade On a Subset of Nodes

To upgrade an Oracle Clusterware release, you must shut down all Oracle Database instances on the subset of nodes you want to upgrade before modifying the Oracle software. Review the instructions in the release upgrade README document for additional instructions specific to the upgrade.

Before you shut down any processes that are monitored by Enterprise Manager Grid Control, set a blackout in Grid Control for the processes that you intend to shut down. This is necessary so that the availability records for these processes indicate that the shutdown was planned downtime, rather than an unplanned system outage.

To shut down Oracle processes and upgrade a subset of nodes, complete the following steps:

> **Note:** You must perform these steps in the order listed.

1. Shut down any processes that may be accessing a database on each node you intend to upgrade. For example, shut down Oracle Enterprise Manager Database Control.

2. Shut down all Oracle RAC instances on the nodes you intend to upgrade. To shut down an Oracle RAC instance for a database, enter the following command *db_name* is the name of the database and *inst_name* is the database instance:

   ```
   $ Oracle_home/bin/srvctl stop instance -d db_name -i inst_name
   ```

3. Shut down all ASM instances on all nodes you intend to upgrade. To shut down an ASM instance, enter the following command, where *ASM_home* is the ASM home location, and *node* is the name of the node where the ASM instance is running:

   ```
   $ ASM_home/bin/srvctl stop asm -n node
   ```

   > **Note:** If you shut down ASM instances, then you must first shut down all database instances on the nodes you intend to upgrade that use ASM, even if these databases run from different Oracle homes.

4. Stop all listeners on the node. To shut down listeners on the node, enter the following command, where *nodename* is the name of the node, and the listener is running from the ASM home:

   ```
   $ ASM_home/bin/srvctl stop listener -n node_name
   ```

5. Stop all node applications on all nodes. To stop node applications running on a node, enter the following command, where *node* is the name of the node where the applications are running

   ```
   $ Oracle_home/bin/srvctl stop nodeapps -n node
   ```

6. Change directory to the Oracle Clusterware home. As `root`, run the `preupdate.sh` script on the local node, and on all other nodes in the subset that you intend to upgrade. Use the following command syntax, where *clusterware_home* is the path to the existing Oracle Clusterware or Cluster Ready Services home, and *installation_owner* is the Oracle Clusterware installation owner:

   ./preupdate.sh -crshome *clusterware_home* -crsuser *installation_owner*

   For example:

   ```
   # cd $ORACLE_HOME/install
   # ./preupdate.sh -crshome /opt/crs -crsuser oracle
   ```

   The script output should be similar to the following:

   ```
   Shutting down Oracle Cluster Ready Services (CRS):
   Stopping resources. This could take several minutes.
   Successfully stopped CRS resources.
   Stopping CSSD.
   Shutting down CSS daemon.
   Shutdown request successfully issued.
   Shutdown has begun. The daemons should exit soon.
   ```

```
Checking to see if Oracle CRS stack is down...
Oracle CRS stack is down now.
```

7. Confirm that you are logged in as the Oracle Clusterware installation owner, and start Oracle Universal Installer to install the software.

   For example:

   ```
   $ whoami
   crs
   $ cd /cdrom/clusterware/
   ./runInstaller
   ```

   Provide information as prompted by the Installer.

   > **Note:** You cannot change the owner of the Oracle Clusterware home.

8. During an Oracle Clusterware installation, if Oracle Universal Installer detects an existing Oracle Clusterware 10*g* release 1 or release 2 installation, then you are given the option to perform a rolling upgrade by installing Oracle Clusterware 11*g* release 1 on a subset of cluster member nodes.

   By default, all the cluster member nodes are checked for upgrade. To perform a rolling upgrade of a subset of nodes, uncheck the cluster member nodes you do not want to upgrade on the Specify Hardware Cluster Installation Mode installation screen.

9. After you select the nodes you want to upgrade, the Installer installs the Oracle Clusterware 11g release 1 software in the existing Oracle Clusterware home on the local and the remote subset of nodes you have selected.

   OUI prompts you to run the appropriate root script for the release or patchset. The script starts the Oracle Clusterware 11g release 1 stack on the upgraded subset of nodes. However, it lists it as an inactive version.

10. After you upgrade the initial node or subset of nodes, repeat steps 1 through 10 for each remaining node or subset of nodes until all nodes in the cluster have been upgraded and the new version is the active version.

    When all member nodes of the cluster are running the new Oracle Clusterware release, then the new clusterware becomes the active version. Otherwise, the older Oracle Clusterware release is still used.

    To list the version of Oracle Clusterware that is installed on a node, enter the following command, where *CRShome* is the Oracle Clusterware home, and *nodename* is the name of the node:

    ```
    # CRShome/bin/crsctl query crs softwareversion [nodename]
    ```

    To list the Oracle Clusterware software version that is running on a node, enter the following command, where *CRShome* is the Oracle Clusterware home:

    ```
    # CRShome/bin/crsctl query crs activeversion
    ```

    If you intend to install or upgrade Oracle RAC, then you must first complete the upgrade to Oracle Clusterware 11*g* release 1 (11.1) on all cluster member nodes before you install the Oracle Database 11*g* release 1 (11.1) version of Oracle RAC.

11. Check with Oracle Support, and apply any recommended patch sets, bundle patches or critical patches.

To check for the latest recommended patches for Oracle Clusterware and Oracle Real Application Clusters, log on to the following site:

https://metalink2.oracle.com

Click **Patches & Updates**, and click Oracle Database from the Recommended Patches list. Provide information as prompted.

Index

wsize, 4-6, 5-10
wsize parameter, 4-6, 5-10

X