Oracle® Database

Oracle Clusterware Installation Guide

11*g* Release 1 (11.1) for AIX

B28258-01

October 2007

ORACLE®

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

B28258-01

Primary Author: Douglas Williams

Contributing Authors: Mark Bauer, Namrata Bhakthavatsalam, Jonathan Creighton, Barb Lundhild, Saar Maoz, Sundar Matpadi, 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, Sana Karam, 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

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 Database Oracle Clusterware Installation Guide for AIX 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.

This preface contains the following topics:

- Intended Audience
- Documentation Accessibility
- Related Documents
- Conventions

Intended Audience

Oracle Database Oracle Clusterware Installation Guide for AIX 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

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be

accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at

http://www.oracle.com/accessibility/

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

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

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, 7 days a week. For TTY support, call 800.446.2398. Outside the United States, call +1.407.458.2479.

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 Database Oracle Clusterware and 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*

- *Oracle Database Installation Guide for AIX*

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

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 (32-Bit)*

Oracle Real Application Clusters Management

- *Oracle Database Oracle Clusterware and 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*

- *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

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):

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.

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.

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.

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:

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

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.

Verify the space available for Oracle Clusterware files using one of the following commands, depending on where you intend to place Oracle Clusterware files:

GPFS:

```
df -k
```

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

```
lslv lv_name
```

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

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

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 -k /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`.

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, on 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.

Check Operating System Packages

Refer to the tables listed inChapter 2 "Checking the Software Requirements"for details.

Set Kernel Parameters

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

"Tuning AIX System Environment"

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"

"Environment Requirements for Oracle Database and Oracle ASM Owners"

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. Then use `smit` to create the oracle user, or enter commands similar to the following:

```
# mkgroup oinstall
# mkgroup dba
# mkuser pgrp=oinstall groups=dba,oinstall oracle
```

Ensure that the Oracle Clusterware `oracle` user has the capabilities CAP_NUMA_ ATTACH, CAP_BYPASS_RAC_VMM, and CAP_PROPAGATE.

To check existing capabilities, enter the following command as `root`; in this example, the Oracle Clusterware `oracle` user is `crs`:

```
# lsuser -a capabilities crs
```

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

```
# chuser capabilities=CAP_NUMA_ATTACH CAP_BYPASS_RAC_VMM CAP_PROPAGATE crs
```

Set the password on the oracle account:

```
passwd oracle
```

Repeat this process for each cluster member node.

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
```

Configure Oracle Installation Owner Shell Limits

For information, review the following section in Chapter 2:

"Configuring Software Owner User Environments"

Configure SSH

For information, review the following section in Chapter 2:

"Configuring SSH on All Cluster Nodes"

OUI uses the ssh and scp commands during installation to run remote commands on and copy files to the other cluster nodes. If SSH is not available, then OUI falls back to rsh and rcp commands. If you want to use SSH for increased security during installation, then you must configure SSH so that the ssh and scp commands used during installation do not prompt for a password.

To configure SSH, complete the following tasks:

Check Existing SSH Configuration on the System

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

```
$ ps -ef grep 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.

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

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.

Create Storage

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

For information, review the following sections in Chapter 4:

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

"Configuring Storage for Oracle Clusterware Files on Raw Devices"

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. To create partitions:

1. Identify or configure the required disk devices.

   The disk devices must be shared on all of the cluster nodes.

2. As the `root` user, enter the following command on any node to identify the device names for the disk devices that you want to use:

   ```
   # lspv | grep -i none
   ```

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

   ```
   hdisk17          0009005fb9c23648                  None
   ```

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

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

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

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

   ```
   # lspv | grep -i "0009005fb9c23648"
   ```

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

   ```
   hdisk18          0009005fb9c23648                  None
   ```

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

5. If the device names are the same on all nodes, then enter commands similar to the following on all nodes to change the owner, group, and permissions on the character raw device files for the disk devices:

   - OCR device:

     ```
     # chown root:oinstall /dev/rhdiskn
     # chmod 640 /dev/rhdiskn
     ```

   - Other devices:

     ```
     # chown oracle:dba /dev/rhdiskn
     # chmod 660 /dev/rhdiskn
     ```

6. If the device name associated with the PVID for a disk that you want to use is different on any node, then you must create a new device file for the disk on each of the nodes using a common unused name.

For the new device files, choose an alternative device file name that identifies the purpose of the disk device. The previous table suggests alternative device file names for each file. For database files, replace *dbname* in the alternative device file name with the name that you chose for the database in step 1.

> **Note:** Alternatively, you could choose a name that contains a number that will never be used on any of the nodes, for example `hdisk99`.

To create a new common device file for a disk device on all nodes, perform these steps on each node:

a. Enter the following command to determine the device major and minor numbers that identify the disk device, where *n* is the disk number for the disk device on this node:

```
# ls -alF /dev/*hdiskn
```

The output from this command is similar to the following:

```
brw------- 1 root system    24,8192 Dec 05 2001  /dev/hdiskn
crw------- 1 root system    24,8192 Dec 05 2001  /dev/rhdiskn
```

In this example, the device file `/dev/rhdiskn` represents the character raw device, `24` is the device major number, and `8192` is the device minor number.

b. Enter a command similar to the following to create the new device file, specifying the new device file name and the device major and minor numbers that you identified in the previous step:

> **Note:** In the following example, you must specify the character `c` to create a character raw device file.

```
# mknod /dev/ora_ocr_raw_256m c 24 8192
```

c. Enter commands similar to the following to change the owner, group, and permissions on the character raw device file for the disk:

– OCR:

```
# chown root:oinstall /dev/ora_ocr_raw_256m
# chmod 640 /dev/ora_ocr_raw_256m
```

– Oracle Clusterware voting disk:

```
# chown oracle:dba /dev/ora_vote_raw_256m
# chmod 660 /dev/ora_vote_raw_256m
```

d. Enter a command similar to the following to verify that you have created the new device file successfully:

```
# ls -alF /dev | grep "24,8192"
```

The output should be similar to the following:

```
brw------- 1 root   system  24,8192 Dec 05 2001  /dev/hdiskn
```

```
crw-r----- 1 root   oinstall 24,8192 Dec 05 2001  /dev/ora_ocr_raw_256m
crw------- 1 root   system  24,8192 Dec 05 2001  /dev/rhdiskn
```

7. To enable simultaneous access to a disk device from multiple nodes, you must set the appropriate Object Data Manager (ODM) attribute listed in the following table to the value shown, depending on the disk type:

Disk Type	Attribute	Value
SSA, FAStT, or non-MPIO-capable disks	reserve_lock	no
ESS, EMC, HDS, CLARiiON, or MPIO-capable disks	reserve_policy	no_reserve

To determine whether the attribute has the correct value, enter a command similar to the following on all cluster nodes for each disk device that you want to use:

```
# lsattr -E -l hdiskn
```

If the required attribute is not set to the correct value on any node, then enter a command similar to one of the following on that node:

- SSA and FAStT devices

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

- ESS, EMC, HDS, CLARiiON, and MPIO-capable devices

  ```
  # chdev -l hdiskn  -a reserve_policy=no_reserve
  ```

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

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

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

```
/dev/rhdisk10
```

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
```

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.

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.

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 create disk 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:

- Logging In to a Remote System as root Using X Terminal
- Overview of Groups and Users for Oracle Clusterware Installations
- Creating Groups and Users for Oracle Clusterware
- Checking the Hardware Requirements
- Checking the Network Requirements
- Checking the Software Requirements
- Tuning AIX System Environment
- Configuring SSH on All Cluster Nodes
- Configuring Software Owner User Environments
- Running the rootpre.sh Script
- Requirements for Creating an Oracle Clusterware Home Directory
- Understanding and Using Cluster Verification Utility
- Checking Oracle Clusterware Installation Readiness with CVU

Logging In to a Remote System as root Using X Terminal

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

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

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

- If you are installing the software from an X Window System workstation or X terminal, then:

    1. Start a local terminal session, for example, an X terminal (`xterm`).

2. If you are not installing the software on the local system, then enter a command using the following syntax to enable remote hosts to display X applications on the local X server:

```
$ xhost + remote_host
```

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

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

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

```
$ ssh remote_host
```

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

```
$ ssh somehost.example.com
```

4. If you are not logged in as the `root` user, then enter the following command to switch the user to `root`:

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

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

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

1. Start the X server software.

2. Configure the security settings of the X server software to permit remote hosts to display X applications on the local system.

3. Connect to the remote system where you want to install the software and start a terminal session on that system, for example, an X terminal (`xterm`).

4. If you are not logged in as the `root` user on the remote system, then enter the following command to switch user to `root`:

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

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`.

> **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 2-4 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

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

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.

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.

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 (oraInventory), 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`).

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

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

```
# mkgroup id=501 install
```

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

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 `oinstall` group as their primary group, and each user must share the same Oracle Central Inventory, to prevent corruption of the Central Inventory. Refer to "Creating Custom Configuration Groups and Users for Job Roles" on page 3-3.

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`).

Creating a 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. Enter the following command:

   ```
   # smit security
   ```

2. On the **Security & Users** menu, select **Users**.

3. On the **Users** menu, select **Add a User**.

4. Choose the appropriate menu items to create the Oracle Clusterware user (`crs` or `oracle`). In the **Primary GROUP** field, specify the Oracle Inventory group. Make a note of the information you provide in the entry fields, so that you can provide the same value on other nodes.

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

5. Press F10 to exit.

6. Set the password of the Oracle Clusterware `oracle` user. For example:

   ```
   # passwd crs
   ```

7. Ensure that the Oracle Clusterware `oracle` user has the capabilities CAP_NUMA_ATTACH, CAP_BYPASS_RAC_VMM, and CAP_PROPAGATE.

   To check existing capabilities, enter the following command as `root`; in this example, the Oracle Clusterware `oracle` user is `crs`:

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

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

   ```
   # /usr/bin/chuser capabilities=CAP_NUMA_ATTACH CAP_BYPASS_RAC_VMM CAP_PROPAGATE
   crs
   ```

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

Modifying an Existing Software Owner User

If the user account you want to use for Oracle Clusterware exists, but its primary group is not `oinstall`, then use the following procedure to modify the user to add groups:

1. Enter the following command:

   ```
   # smit security
   ```

2. Choose the appropriate menu items to modify the `oracle` user.

3. In the **Primary GROUP** field, specify the Oracle Inventory group, for example `oinstall`.

4. Press F10 to exit.

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

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 /u01/app/crs
# chown crs:oinstall /u01/app
# mkdir /u01/app/oracle
# chown oracle:oinstall /u01/app/oracle
# chmod -R 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.

Checking the Hardware Requirements

Each system must meet the following minimum hardware requirements:

- At least 1 GB of physical RAM

- Swap space equivalent to the multiple of the available RAM, using the following table as a guideline:

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 AIX systems with 1 GB or more of memory, Oracle recommends that you set the paging space to an initial setting of half the size of RAM plus 4 GB, with an upper limit of 32 GB. During installation, to optimize paging, monitor the paging space use in a separate window. Use the command `lsps -a` to increase or decrease the paging space size. The output of lsps -a should indicate paging space use of less than 25 percent on a healthy system. Refer to *Oracle Database Administrator's Reference for AIX* for more information about configuring paging 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

- 7.5 GB of disk space for the Oracle Database files (Oracle base)

- 2 GB of disk space for a preconfigured database that uses file system storage (optional)

> **Note:** The disk space requirement for databases that use Automatic Storage Management or raw device storage is described in Chapter 4.

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

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

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

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

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

2. To determine the available RAM and swap space, enter the following command:

   ```
   # /usr/sbin/lsps -s
   ```

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

   ```
   # /usr/sbin/lsps -a
   ```

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

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

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

   If there is less than 400 MB of disk space available in the /tmp directory, then complete one of the following steps:

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

   - Set the TEMP and TMPDIR environment variables when setting the oracle user's environment (described later).

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

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

GPFS:

   ```
   # df -k
   ```

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

```
# lslv lv_name
```

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

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

The following table shows the approximate disk space requirements for software files for each installation type:

Installation Type	Requirement for Software Files (GB)
Enterprise Edition	4 GB
Standard Edition	4 GB
Custom (maximum)	4 GB

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

```
# getconf HARDWARE_BITMODE
```

> **Note:** The expected output of this command is 64. If you do not see the expected output, then you cannot install the software on this system.

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

```
# bootinfo -K
```

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

Checking the Network Requirements

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Real Application Clusters installation.

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

Network Hardware Requirements

Each node in the cluster must meet the following requirements:

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

■ 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 en0 as the public interface, but on node2 have en1 as the public interface. Public interface names must be the same, so you must configure en0 as public on *both* nodes. You should configure the private interfaces on the same network adapters as well. If en1 is the private interface for node 1, then en1 should be the private interface for node 2.

■ For increased reliability, configure redundant public and private network adapters for each node.

■ 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 recommended).

> **Note:** UDP is the default interconnect protocol for Oracle RAC, and TCP is the interconnect protocol for Oracle Clusterware.
>
> Token-Ring is not supported 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. You can test whether an interconnect interface is reachable using a ping command.

IP Address Requirements

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

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

■ One virtual IP (VIP) address with an associated network name registered in DNS. If you do not have an available DNS, then record the network 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 network name are currently unused

– The VIP is on the same subnet as your public interface

Before installation, check that the default gateway can be accessed by a ping command. During installation, OUI uses the ping command to ensure that the VIP is reachable. 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 network name associated with the VIP. If a node fails, then the node's virtual IP address fails over to another node.

■ 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.*.*). Use the /etc/hosts file on each node to associate private network names with private IP addresses.

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	Interface Name	Type	IP Address	Registered In
rac1	rac1	Public	143.46.43.100	DNS (if available, else the hosts file)
rac1	rac1-vip	Virtual	143.46.43.104	DNS (if available, else the hosts file)
rac1	rac1-priv	Private	10.0.0.1	Hosts file
rac2	rac2	Public	143.46.43.101	DNS (if available, else the hosts file)
rac2	rac2-vip	Virtual	143.46.43.105	DNS (if available, else the hosts file)
rac2	rac2-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. When a node or interconnect fails, then the associated VIP is relocated to the surviving instance, enabling fast notification of the failure to the clients connecting through that VIP. If the application and client are configured with transparent application failover options, then the client is reconnected to the surviving instance.

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.

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. Register the host names and IP addresses for the public network interfaces in DNS.

3. For each node, register one virtual host name and IP address in DNS.

4. For each private interface on every node, add a line similar to the following to the `/etc/hosts` file on all nodes, specifying the private IP address and associated private host name:

   ```
   10.0.0.1     rac1-priv1
   ```

5. To identify the interface name and associated IP address for every network adapter, enter the following command:

   ```
   # netstat -i
   ```

   From the output, identify the interface name and IP address for all network adapters that you want to specify as public or private network interfaces.

   > **Note:** When you install Oracle Clusterware and RAC, you will require this information.

Checking the Software Requirements

Depending on the products that you intend to install, verify that the following software is installed on the system. The procedure following the table describes how to verify whether these requirements are addressed.

> **Note:** Oracle Universal Installer performs checks on your system to verify that it meets the listed requirements. To ensure that these checks pass, verify the requirements before you start Oracle Universal Installer.

Table 2–1 AIX Operating System Kernel Requirements

Item	Requirement
Operating system	AIX 5L version 5.3 Maintenance Level 06 or later (64-bit)
Operating system filesets	The following operating system filesets are required:

	`bos.adt.base`
	`bos.adt.lib`
	`bos.adt.libm`
	`bos.perf.libperfstat`
	`bos.perf.perfstat`
	`bos.perf.proctools`
	`rsct.basic.rte`
	`rsct.compat.clients.rte`
	`xlC.aix50.rte 8.0.0.5`
	`xlC.rte 8.0.0.5`

	You must have the xlC C/C++ runtime filesets for installation, but you do not require the C/C++ compilers. You do not require a license for the xlC runtime filesets
	To obtain the xlC runtime filesets, Oracle Database users who do not install the IBM XL C/C++ Enterprise Edition V8.0 compiler should install the IBM XL C/C++ Enterprise Edition V8.0 for AIX Runtime Environment Component, which contains xLC.aix.rte 8.0.0.5 and xlC.rte 8.0.0.5.
	Download the C/C++ compiler from the following Web site:
	`http://www-1.ibm.com/support/docview.wss?uid=swg24015075`
	Download the C runtime environment file sets, with no license requirement, from the following Web site:
	`http://www-1.ibm.com/support/docview.wss?uid=swg24015077`
Oracle RAC	High Availability Cluster Multi-Processing (HACMP) v5.3
	Note: HACMP is required only if you want to use raw logical volumes for Oracle Clusterware or database file storage. However, it is supported for all installations.
	General Parallel File System (GPFS) v3.1.0.10 or later.
	Note: GPFS is required only if you want to use a cluster file system for Oracle Clusterware or database files.
ADA	OC Systems PowerAda 5.4d
JDK	IBM JDK 1.5 is installed with this release.
Pro*FORTRAN	IBM XL Fortran v. 10.1 for AIX

Table 2–1   (Cont.) AIX Operating System Kernel Requirements

Item	Requirement
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK), GNU Compiler Collection (GCC)	Note: If you do not install the C/C++ compilers, then you require the C/C++ runtime filesets for installation as described in the "Operating system filesets" row in this table. ■ February 2007 XL C/C++ Enterprise Edition V8.0 for AIX PTF 8.0: You can download the PTF from the following link: http://www-1.ibm.com/support/docview.wss?uid=swg24015075 Download the C++ runtime environment file sets, with no license requirement, from the following Web site: http://www-1.ibm.com/support/docview.wss?uid=swg24015077 ■ gcc 3.45
Pro*COBOL	■ AcuCobol 6.1 ■ Micro Focus Server Express 5.0 (Both 32-bit and 64-bit)

Oracle Messaging Gateway	IBM WebSphere MQ V6.0, client and server:
	
	`mqm.Client.Bnd`
	`mqm.Server.Bnd`

To ensure that the system meets these requirements:

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

```
# oslevel -r
```

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

http://www-912.ibm.com/eserver/support/fixes/

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

```
# lslpp -l bos.adt.base bos.adt.lib bos.adt.libm bos.perf.perfstat \
 bos.perf.libperfstat bos.perf.proctools rsct.basic.rte
```

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

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

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

Table 2–2   AIX APAR and Other Operating System Fixes

Installation Type or Product	Requirement
All AIX 5L v5.3 installations	All AIX 5L v. 5.3 installations Authorized Problem Analysis Reports (APARs) for AIX 5L v. 5.3 ML06, and the following AIX fixes:
	IY89080 IY92037 IY94343 IZ01060, or efix for IZ01060 IZ03260, or efix for IZ03260
Oracle JDBC/OCI Drivers AIX 5L v5.3	**Note:** These APARs are required only if you are using the associated JDK version.
	APAR required for JDK 1.4.2 (64-bit):
	■ IY63533: JDK 1.4.2 64-bit SR1 caix64142-20040917
	APARs required for JDK 1.3.1.16 (32-bit):
	■ IY58350: SDK 1.3.1 32-Bit SR7P: CA131IFX-20040721A
	■ IY65305: JAVA142 32-bit PTF: CA142IFX-20041203
RAC AIX 5L v5.3	If you use HACMP, then note the following additional requirements:
	■ AIX: AIX 5.3 TL06 or newer (`bosrte.lvm` must be at least 5.3.0.60)
	■ HACMP: Ensure the following versions are installed: * HACMP v. 5.3 with PTFS (APAR:IY94307) and `cluster.es.clvm` installed * HACMP APAR: IZ01809
	■ RSCT (component of AIX): RSCT (`rsct.basic.rte`) version 2.4.7.3
	APARs required for GPFS v3.1.0.10:
	None.

To ensure that the system meets these requirements:

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

   ```
   # instfix -i -k "IY92312 IY97156 IY76252 IY94335 IY94334"
   ```

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

   ```
   http://www-03.ibm.com/servers/eserver/support/pseries/aixfixes.html
   ```

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

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

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

   ```
   http://www-03.ibm.com/servers/eserver/support/pseries/aixfixes.html
   ```

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

http://www-306.ibm.com/software/

Tuning AIX System Environment

Perform the following system tuning and configuration all cluster nodes.

- Tuning Virtual Memory Manager (VMM)
- Increasing System Block Size Allocation
- Configuring Shell Limits
- Configuring User Process Parameters
- Configuring Network Tuning Parameters

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

Tuning Virtual Memory Manager (VMM)

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

Table 2–3   Recommended Values for Virtual Memory Manager

Parameter	Value
minperm%	3 (default is 20)
maxperm%	90 (default is 80)
maxclient% = 90	90 (default is 80)
lru_file_repage	0 (default is 1)
strict_maxclient	1 (default is 1)
strict_maxperm	0 (default is 0)

For example:

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

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

Increasing System Block Size Allocation

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

For example:

```
chdev -l sys0 -a ncargs='128'
```

Configuring Shell Limits

To improve the performance of the software, you must increase the following shell limits:

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	maxuproc	16384

To increase the shell limits:

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

```
default:
        fsize = -1
        core = -1
        cpu = -1
        data = 512000
        rss = 512000
        stack = 512000
        nofiles = 2000
```

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

```
/etc/lsattr -E -l sys0 -a maxuproc
```

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

```
/etc/chdev -l sys0 -a maxuproc = 16384
```

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

Configuring User Process Parameters

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

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

1. Enter the following command:

```
# smit chgsys
```

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

If necessary, edit the existing value.

3. When you have finished making changes, press F10 to exit.

Configuring Network Tuning Parameters

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

Network Tuning Parameter	Recommended Value
ipqmaxlen	512
rfc1323	1
sb_max	2*655360
tcp_recvspace	65536
tcp_sendspace	65536
udp_recvspace	655360 **Note:** The recommended value of this parameter is 10 times the value of the udp_sendspace parameter. The value must be less than the value of the sb_max parameter.
udp_sendspace	65536 **Note:** This value is suitable for a default database installation. For production databases, the minimum value for this parameter is 4 KB plus the value of the database DB_BLOCK_SIZE initialization parameter multiplied by the value of the DB_MULTIBLOCK_READ_COUNT initialization parameter: (DB_BLOCK_SIZE * DB_MULTIBLOCK_READ_COUNT) + 4 KB

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

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

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

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

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

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

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

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

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

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

      For example:

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

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

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

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

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

- ipqmaxlen parameter:

```
no -r -o ipqmaxlen=512
```

- Other parameter:

```
no -p -o parameter=value
```

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

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

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. If SSH is not available, then OUI falls back to rsh and rcp commands. If you want to use SSH for increased security during installation, then you must configure SSH so that the ssh and scp commands used during installation do not prompt for a password. The SSH configuration procedure in this section describes how to configure SSH using SSH1.

This section contains the following sections:

- Checking Existing SSH Configuration on the System

- Downloading and Installing SSH Filesets

- 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

Checking Existing SSH Configuration on the System

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

```
$ ps -ef | grep 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.

Downloading and Installing SSH Filesets

1. Install the OpenSSL package in the AIX bonus package installation disks, or download OpenSSL from the following URL:

   http://www-306.ibm.com/software/

   You require an IBM ID or Partner World ID to access this site.

2. Install the following file sets from the AIX Base installation media:

   - openssh.base
   - openssh.license
   - openssh.msg.en_US
   - openssh.man.en_US

   You can also obtain filesets from the following URL:

   http://sourceforge.net/projects/openssh-aix

3. Start the sshd daemon by running the following command:

   ```
   /usr/bin/startsrc -s sshd
   ```

   > **Note:** If the AIX server on which OpenSSH is installed also has GSA installed, then the SSH daemon will not start. This is a known problem.
   >
   > To resolve this issue, check to see if the sshd user exists on the system. If not, then create it using the following commands:
   >
   > ```
   > mkgroup sshd
   > mkuser -a pgrp=sshd login=false home=/var/empty gecos="OpenSSH\
   > privilege separation" account_locked=true sshd
   > ```

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:

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 oracle`. 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.

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: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.

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, proceed to
   "Setting Display and X11 Forwarding Configuration" on page 2-23.

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.

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
```

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, `.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.

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.

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 -k command to identify a suitable file system with sufficient free
   space.

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

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

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

   * Bourne, Bash, or Korn shell:

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

   * C shell:

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

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.

Running the rootpre.sh Script

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

Run the rootpre.sh script:

1. Switch user to root:

```
$ su - root
```

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

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

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

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

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

3. Exit from the root account:

```
# exit
```

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

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

Requirements for Creating an Oracle 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)

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 3-3 for information about creating groups, users, and software homes for additional Oracle software installations

Understanding and Using Cluster Verification Utility

Cluster Verification Utility (CVU) is a tool that performs system checks. This guide provides CVU 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

Entering Cluster Verification Utility Commands

CVU 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 CVU 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 CVU command options that you select. For example:

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

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

Using CVU to Determine if Installation Prerequisites are Complete

You can use CVU 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 CVU, notably configuring SSH between all nodes in the cluster, before you can complete a clusterwide status check.

- CVU 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 CVU report, and proceed to the sections of the preinstallation chapter to complete additional steps as needed.

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 CVU command, then CVU shows the correct usage for that command. For example, if you type `runcluvfy.sh stage -pre dbinst`, then CVU shows the correct syntax for the database preinstallation checks that CVU performs with the `dbinst` stage option. The following is a list of context help commands.

- `cluvfy -help` —CVU displays detailed CVU command information.

- `cluvfy comp -list`—CVU displays a list of components that can be checked, and brief descriptions of how each component is checked.

- `cluvfy comp -help`—CVU displays detailed syntax for each of the valid component checks.

- `cluvfy stage -list`—CVU displays a list of valid stages.

- `cluvfy stage -help`—CVU displays detailed syntax for each of the valid stage checks.

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

You can use CVU 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 CVU to check 10. 2 installations, append the command flag `-r 10gR2` to the standard CVU 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 CVU checks for 11*g* release 1 (11.1) requirements.

Verbose Mode and "Unknown" Output

If you run CVU using the -verbose argument, and a CVU command responds with UNKNOWN for a particular node, then this is because the CVU cannot determine whether a check passed or failed. The following is a list of possible causes for an "Unknown" response:

- The node is down

- Executables required by CVU are missing in the /bin directory in the Oracle Clusterware home or Oracle home directory

- The user account starting CVU 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

Checking Oracle Clusterware Installation Readiness with CVU

Use the 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-separated 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 the CVU 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 the CVU 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

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

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 3-3.

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

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.

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

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 you need to create a nobody user, then enter the following command:

   ```
   # smit security
   ```

3. Select **Users**, then **Add a User**.

4. On the **Add a User** form, specify the appropriate options to create an unprivileged nobody user, and make a note of the user ID.

5. Press F10 to exit.

6. Repeat this process on each node in the cluster, ensuring that the user ID for each `nobody` user is identical on each cluster node.

Creating the OSDBA Group

To create the OSDBA group, complete the following procedure:

1. Enter the following command:

   ```
   # smit security
   ```

2. Select **Groups**, then **Add a Group**.

3. On the **Add a Group** form, specify the appropriate options to create an OSDBA group, with a group ID that can be identical on all nodes in the cluster. Make a note of the group ID.

4. Press F10 to exit.

5. Repeat this process on each node in the cluster, ensuring that the group ID for OSDBA is identical on each cluster node.

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 on the basis of 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.

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

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.

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.

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

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 Oracle Software Owner User

- Creating Identical Database Users and Groups on Other Cluster Nodes

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:

1. Enter the following command:

   ```
   # smit security
   ```

2. Select **Groups**, then **Add a Group**.

3. On the **Add a Group** form, specify the appropriate options to create an OSDBA group, with a group ID that can be identical on all nodes in the cluster. Make a note of the group ID.

4. Press F10 to exit.

5. Repeat this process on each node in the cluster, ensuring that the group ID for OSDBA is identical on each cluster node.

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.

1. Enter the following command:

   ```
   # smit security
   ```

2. Select **Groups**, then **Add a Group**.

3. On the **Add a Group** form, specify the appropriate options to create an OSOPER group, with a group ID that can be identical on all nodes in the cluster. Make a note of the group ID.

4. Press F10 to exit.

5. Repeat this process on each node in the cluster, ensuring that the group ID for OSOPER is identical on each cluster node.

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:

1. Enter the following command:

   ```
   # smit security
   ```

2. Select **Groups**, then **Add a Group**.

3. On the **Add a Group** form, specify the appropriate options to create an OSASM group, with a group ID that can be identical on all nodes in the cluster. Make a note of the group ID.

4. Press F10 to exit.

5. Repeat this process on each node in the cluster, ensuring that the group ID for OSASM is identical on each cluster node.

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:

1. Enter the following command:

   ```
   # smit security
   ```

2. Select **Groups**, then **Add a Group**.

3. On the **Add a Group** form, specify the appropriate options to create an ASMDBA group, with a group ID that can be identical on all nodes in the cluster. Make a note of the group ID.

4. Press F10 to exit.

5. Repeat this process on each node in the cluster, ensuring that the group ID for ASMDBA is identical on each cluster node.

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`.

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 3-8.

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

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. Enter the following command:

   ```
   # smit security
   ```

2. Select **Users**, then **Add a User**.

3. On the **Add a User** form, specify the appropriate options to create the `oracle` user. These include the following:

   - User NAME: The oracle user account name. For example: `oracle`.

   - User ID: The user ID of the `oracle` account.

   - PRIMARY Group: The primary group ID for the `oracle` user. For example: `oinstall`.

   - Group SET: The primary and secondary groups for the `oracle` user. For example: `oinstall,dba,oper,hagsuser`

4. Press F10 to exit.

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

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 modify the user group configuration:

1. Enter the following command:

   ```
   # smit security
   ```

2. From the **Security and Users** menu, select **Users**, then **Change/Show Characteristics of a User**.

3. Enter the oracle user name, and press Enter.

4. In the **Change/Show Characteristics of a User** menu, update the entry fields as necessary. When you are finished, press the **Enter** key to save changes and exit.

Creating a Separate ASM Owner

1. Enter the following command:

```
# smit security
```

2. Select **Users**, then **Add a User**.

3. On the **Create User** form, specify the appropriate options to create the `asm` user. These include the following:

   ■ User NAME: The oracle user account name. For example: `asm`.

   ■ User ID: The user ID of the `asm` account.

   ■ PRIMARY Group: The primary group ID for the `asm` user. For example: `oinstall`.

   ■ Group SET: The primary and secondary groups for the `asm` user. For example: `oinstall,asmdba`.

4. Press F10 to exit.

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

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. Use `smit`, or 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.

```
# mkgroup id=501 oinstall
# mkgroup id=502 crs
# mkgroup id=503 dba
# mkgroup id=505 oper
# mkgroup id=504 asm
# mkgroup id=506 asmdba
# mkuser id=502 pgrp=oinstall groups=dba,oracle gecos="RDBMS owner" oracle
passwd db2inst1
```

> **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):

```
# mkuser id='502' pgrp='oinstall' groups='dba','oper' oracle
```

In the preceding command:

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

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

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

> **Note:** You must specify 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 2-18.

6. Complete user environment configuration tasks for each user as described in the section Configuring Software Owner User Environments on page 2-24.

Understanding the Oracle Base Directory Path

This section contains information about preparing an Oracle base directory.

Overview of the Oracle Base directory

During installation, you are prompted to specify an Oracle base location, which is owned by the user performing the installation. You can choose a location with an existing Oracle home, or choose another directory location that does not have the structure for an Oracle base directory. 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.

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.

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. Use smit, or enter shell commands. For example:

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

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
  ```

- 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

4

Configuring Oracle Clusterware Storage

This chapter describes the storage configuration tasks that you must complete before you start Oracle Universal Installer. It includes information about the following tasks:

- Reviewing Storage Options for Oracle Clusterware Files
- Configuring Storage for Oracle Clusterware Files on a Supported Shared File System
- Configuring Storage for Oracle Clusterware Files on Raw Devices

Reviewing Storage Options for Oracle Clusterware Files

This section describes supported options for storing Oracle Clusterware files. It includes the following sections:

- Overview of Storage Options
- Checking for Available Shared Storage with CVU

Overview of Storage Options

Use the information in this overview to help you select your storage option.

> **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/

Overview of Oracle Clusterware Storage Options

There are two ways of storing Oracle Clusterware files:

- **A supported shared file system**: The supported file system includes:
  - **General Parallel File System (GPFS)**: A cluster file system for AIX that provides concurrent file access
  - **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. Note that Direct NFS is not supported for Oracle Clusterware files.

- **Raw partitions**: Raw partitions are raw disks that are accessed either through a logical volume manager (LVM), or through non-LVM file systems.

General Storage Considerations

For all installations, you must choose the storage option that you want to use for Oracle Clusterware files and Oracle Database files. 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 flash recovery area). 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 file in a cluster file system environment. If you do not use a cluster file system, then you must place these files on a shared raw device. Oracle Universal Installer (OUI) automatically initializes the OCR during the Oracle Clusterware installation.

For voting disk file placement, ensure that each voting disk is configured so that it does not share any hardware device or disk, or other single point of failure. An absolute majority of voting disks configured (more than half) must be available and responsive at all times for Oracle Clusterware to operate.

For single-instance Oracle Database installations using Oracle Clusterware for failover, you must use GPFS or shared raw disks if you do not want the failover processing to include dismounting and remounting disks.

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

Table 4–1   Supported Storage Options for Oracle Clusterware Files

Storage Option	OCR and Voting Disks	Oracle Software
Automatic Storage Management	No	No
General Parallel File System (GPFS)	Yes	Yes
• **Note**: Oracle does not recommend the use of GPFS for voting disks if HACMP is used.		
Local storage	No	Yes

Table 4–1   (Cont.) Supported Storage Options for Oracle Clusterware Files

Storage Option	OCR and Voting Disks	Oracle Software
NFS file system ■ **Note:** Requires a certified NAS device. Oracle does not recommend the use of NFS for voting disks if HACMP is used.	Yes	Yes
Raw Logical Volumes Managed by HACMP ■ **Note:** If HACMP is used, then voting disks can only be placed on raw devices or logical volumes.	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.

■ You cannot use Automatic Storage Management to store Oracle Clusterware files, because these files must be accessible before any Automatic Storage Management 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.

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 11g Release 1 (11.1). You cannot reuse the old Oracle9*i* release 9.2 quorum disk for this purpose.

After You Have Selected Disk Storage Options

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

1: Check for available shared storage with CVU

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

2: Configure shared storage for Oracle Clusterware files

■ **To use a file system (local or GPFS) for Oracle Clusterware files**, refer to Configuring Storage for Oracle Clusterware Files on a Supported Shared File System on page 4-4

■ **To use raw devices (partitions) for Oracle Clusterware files**, refer to "Configuring Storage for Oracle Clusterware Files on Raw Devices" on page 4-7

3: Configure storage for Oracle Database files and recovery files

■ **To use a file system for database or recovery file storage**, refer to Configuring Storage for Oracle Clusterware Files on a Supported Shared File System on page 4-4, and ensure that in addition to the volumes you create for Oracle

Clusterware files, you also create additional volumes with sizes sufficient to store database files.

> **See Also:** "Configuring Storage for Oracle Database Files on Shared Storage Devices" in Chapter 5

Checking for Available Shared Storage with CVU

To check for all shared file systems available across all nodes on the cluster with GPFS, use the following command:

```
/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/rhdisk8` and `/dev/rhdisk9`, and your mountpoint is `/dev/dvdrom/`, then enter the following command:

```
/dev/dvdrom/runcluvfy.sh comp ssa -n node1,node2 -s /dev/rhdisk8,/dev/rhdisk9
```

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

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

- Checking NFS Mount Buffer Size Parameters for Oracle Clusterware

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

> **Note:** 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 make this file a shared raw device. Oracle Universal Installer (OUI) automatically initializes the OCR during the Oracle Clusterware installation.

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 on AIX, you must use GPFS.

- If you choose to place your Oracle Cluster Registry (OCR) 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),

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

- If you intend to use a shared file system to store database files, then use 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 Clusterware owner (-, `oracle`) 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 volume
Oracle Clusterware files (OCR and voting disks) with redundancy provided by Oracle software	1	At least 280 MB for each volume
Redundant Oracle Clusterware files with redundancy provided by Oracle software (mirrored OCR and two additional voting disks)	1	At least 280 MB of free space for each OCR location, if the OCR is configured on a file system *or* At least 280 MB available for each OCR location if the OCR is configured on raw devices. *and* At least 280 MB for each voting disk location, with a minimum of three disks.

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).

Checking NFS Mount Buffer Size Parameters for Oracle Clusterware

If you are using NFS, then you must set the values for the NFS buffer size parameters rsize and wsize to at least 16384. Oracle recommends that you use the value 32768. Update the /etc/fstab file on each node with an entry similar to the following:

```
nfs_server:/vol/DATA/oradata  /home/oracle/netapp    nfs\
cio,rw,bg,hard,intr,rsize=32768,wsize=32768,tcp,noac,vers=3,timeo=600
```

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

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 GPFS storage, you must complete this procedure only if you want to place the Oracle Clusterware files on a separate file system to the Oracle base directory.

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

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

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

2. Use the df -k command to determine the free disk space on each mounted file system.

3. From the display, identify the file systems that you want to use:

File Type	File System Requirements
Oracle Clusterware files	Choose a file system with at least 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 has permissions to create directories on the disks where you plan to install Oracle Clusterware, then OUI creates the Oracle Clusterware file directory.

    If the user performing installation does not have write access, then you must create these directories manually using commands similar to the following to

create the recommended subdirectories in each of the mount point directories and set the appropriate owner, group, and permissions on them:

- Oracle Clusterware file directory:

```
# mkdir /mount_point/oracrs
# chown oracle:oinstall /mount_point/oracrs
# chmod 775 /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 subdirectories in each of the mount point directories, and set the appropriate owner, group, and permissions, you have completed GPFS configuration.

Configuring Storage for Oracle Clusterware Files on Raw Devices

The following subsections describe how to configure Oracle Clusterware files on raw partitions.

- Identifying Required Raw Partitions for Clusterware Files
- Configuring Raw Disk Devices for Oracle Clusterware Without HACMP or GPFS
- Configuring HACMP Multinode Disk Heartbeat (MNDHB) for Oracle Clusterware
- Configuring Raw Logical Volumes for Oracle Clusterware
- Creating a Volume Group for Oracle Clusterware
- Configuring Raw Logical Volumes in the New Oracle Clusterware Volume Group
- Importing the Volume Group on the Other Cluster Nodes
- Activating the Volume Group in Concurrent Mode on All Cluster Nodes

Identifying Required Raw Partitions for Clusterware Files

Table 4–3 lists the number and size of the raw partitions that you must configure for Oracle Clusterware files.

> **Note:** Because each file requires exclusive use of a complete disk device, Oracle recommends that, if possible, you use disk devices with sizes that closely match the size requirements of the files that they will store. You cannot use the disks that you choose for these files for any other purpose.

Table 4–3   Raw Partitions Required for Oracle Clusterware Files on AIX

Number	Size for Each Partition (MB)	Purpose
2 (or 1, if you have external redundancy support for this file)	256	Oracle Cluster Registry **Note:** You need to create these raw partitions only once on the cluster. If you create more than one database on the cluster, then they all share the same Oracle Cluster Registry (OCR). You should create two partitions: One for the OCR, and one for a mirrored OCR. If you are upgrading from Oracle9*i* release 2, then you can continue to use the raw device that you used for the SRVM configuration repository instead of creating this new raw device.
3 (or 1, if you have external redundancy support for this file)	256	Oracle Clusterware voting disks **Note:** You need to create these raw partitions only once on the cluster. If you create more than one database on the cluster, then they all share the same Oracle Clusterware voting disk. You should create three partitions: One for the voting disk, and two for additional voting disks.

Configuring Raw Disk Devices for Oracle Clusterware Without HACMP or GPFS

If you are installing Oracle RAC on an AIX cluster without HACMP or GPFS, then you must use shared raw disk devices for the Oracle Clusterware files. You can also use shared raw disk devices for database file storage. However, Oracle recommends that you use Automatic Storage Management to store database files in this situation.

This section describes how to configure the shared raw disk devices for Oracle Clusterware files (Oracle Cluster Registry and Oracle Clusterware voting disk). It also describes how to configure shared raw devices for Oracle ASM and for Database files, if you intend to install Oracle Database, and you need to create new disk devices.

> **Note:**   In the following procedure, you are directed to set physical volume IDs (PVIDs) to confirm that all devices appear under the same name on all nodes. Oracle recommends that you complete the entire procedure, even if you are certain that you do not have PVIDs configured on your system, to prevent the possibility of configuration issues.

To configure shared raw disk devices for Oracle Clusterware files:

1. Identify or configure the required disk devices.

   The disk devices must be shared on all of the cluster nodes.

2. As the `root` user, enter the following command on any node to identify the device names for the disk devices that you want to use:

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

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

   ```
   hdisk17          0009005fb9c23648                    None
   ```

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

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

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

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

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

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

```
hdisk18          0009005fb9c23648                     None
```

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

5. If the device names are the same on all nodes, then enter commands similar to the following on all nodes to change the owner, group, and permissions on the character raw device files for the disk devices:

- OCR device:

  ```
  # chown root:oinstall /dev/rhdiskn
  # chmod 640 /dev/rhdiskn
  ```

- Other devices:

  ```
  # chown oracle:dba /dev/rhdiskn
  # chmod 660 /dev/rhdiskn
  ```

6. If the device name associated with the PVID for a disk that you want to use is different on any node, then you must create a new device file for the disk on each of the nodes using a common unused name.

For the new device files, choose an alternative device file name that identifies the purpose of the disk device. The previous table suggests alternative device file names for each file. For database files, replace *dbname* in the alternative device file name with the name that you chose for the database in step 1.

Note: Alternatively, you could choose a name that contains a number that will never be used on any of the nodes, for example `hdisk99`.

To create a new common device file for a disk device on all nodes, perform these steps on each node:

a. Enter the following command to determine the device major and minor numbers that identify the disk device, where *n* is the disk number for the disk device on this node:

```
# ls -alF /dev/*hdiskn
```

The output from this command is similar to the following:

```
brw------- 1 root system    24,8192 Dec 05 2001  /dev/hdiskn
crw------- 1 root system    24,8192 Dec 05 2001  /dev/rhdiskn
```

In this example, the device file `/dev/rhdiskn` represents the character raw device, `24` is the device major number, and `8192` is the device minor number.

b. Enter a command similar to the following to create the new device file, specifying the new device file name and the device major and minor numbers that you identified in the previous step:

> **Note:** In the following example, you must specify the c flag to create a character raw device file.

```
# mknod /dev/ora_ocr_raw_256m c 24 8192
```

c. Enter commands similar to the following to change the owner, group, and permissions on the character raw device file for the disk:

– OCR:

```
# chown root:oinstall /dev/ora_ocr_raw_256m
# chmod 640 /dev/ora_ocr_raw_256m
```

– Oracle Clusterware voting disk:

```
# chown oracle:dba /dev/ora_vote_raw_256m
# chmod 660 /dev/ora_vote_raw_256m
```

d. Enter a command similar to the following to verify that you have created the new device file successfully:

```
# ls -alF /dev | grep "24,8192"
```

The output should be similar to the following:

```
brw------- 1 root   system   24,8192 Dec 05 2001  /dev/hdiskn
crw-r----- 1 root   oinstall 24,8192 Dec 05 2001  /dev/ora_ocr_raw_256m
crw------- 1 root   system   24,8192 Dec 05 2001  /dev/rhdiskn
```

7. To enable simultaneous access to a disk device from multiple nodes, you must set the appropriate Object Data Manager (ODM) attribute listed in the following table to the value shown, depending on the disk type:

Disk Type	Attribute	Value
SSA, FAStT, or non-MPIO-capable disks	reserve_lock	no
ESS, EMC, HDS, CLARiiON, or MPIO-capable disks	reserve_policy	no_reserve

To determine whether the attribute has the correct value, enter a command similar to the following on all cluster nodes for each disk device that you want to use:

```
# /usr/sbin/lsattr -E -l hdiskn
```

If the required attribute is not set to the correct value on any node, then enter a command similar to one of the following on that node:

■ SSA and FAStT devices

```
# /usr/sbin/chdev -l hdiskn  -a reserve_lock=no
```

■ ESS, EMC, HDS, CLARiiON, and MPIO-capable devices

```
# /usr/sbin/chdev -l hdiskn  -a reserve_policy=no_reserve
```

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

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

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

```
/dev/rhdisk10
```

Configuring HACMP Multinode Disk Heartbeat (MNDHB) for Oracle Clusterware

This section contains the following topics:

- Overview of Requirements for Using HACMP with Oracle Clusterware
- Deploying HACMP and MDNDHB for Oracle Clusterware
- Upgrading an Existing Oracle Clusterware and HACMP Installation

Overview of Requirements for Using HACMP with Oracle Clusterware

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

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

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

Deploying HACMP and MDNDHB for Oracle Clusterware

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

1. Start HACMP.

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

```
# lssrc -s clcomdES
```

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

```
# startsrc -s clcomdES
```

3. Ensure that your versions of HACMP and AIX meet the system requirements listed in Chapter 2, "Checking the Software Requirements".

4. Create HACMP cluster and add the Oracle Clusterware nodes. For example:

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

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

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

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

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

   For example:

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

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

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

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

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

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

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

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

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

For example:

```
# smitty cm_change_show_resources_std_resource_group_menu_dmn.select
- select_resource_group_from_step_6
Resource Group Name shared_storage
Participating Nodes (Default Node Priority) mynode1,mynode2,mynode3
Startup Policy Online On All Available Nodes
Fallover Policy Bring Offline (On Error Node Only)
Fallback Policy Never Fallback
Concurrent Volume Groups [enter_VG_from_step_7]
Use forced varyon of volume groups, if necessary false
Application Servers []
```

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

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

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

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

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

13. Verify and Synchronize HACMP configuration. For example:

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

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

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

Upgrading an Existing Oracle Clusterware and HACMP Installation

Complete the following procedure:

1. Back up all databases, and back up the Oracle Cluster Registry (OCR)

2. Shut down on all nodes all Oracle RAC databases, all node applications, and Oracle Clusterware.

3. Enter the following command to disable Oracle Clusterware from starting when nodes are restarted:

```
# crsctl disable crs
```

4. Shut down HACMP on all nodes.

5. Install HACMP APAR IZ01809, following the directions in the README included with that APAR.

6. Determine if the existing voting disk LVs are already on separate hard disks, and if each of these disks have sufficient space (at least 256 MB for the MNDHB LVs. If this is true, then create a MNDHB LV on each of the hard disks. If this is not true, then create new MNDHB LVs and new voting disk LVs, located on separate hard disks using the following command, responding to the sections in italics with the appropriate information for your system:

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

7. Verify and Synchronize HACMP configuration.

8. Start HACMP on all nodes.

9. If you added new LVs for voting disks in step 5, then replace each of the existing voting disks with the new ones.

10. Enter the following command to re-enable Oracle Clusterware:

```
# crsctl enable CRS
```

11. Start Oracle Clusterware on all nodes, and verify that all resources start correctly.

Configuring Raw Logical Volumes for Oracle Clusterware

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

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

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

- You must use concurrent-capable volume groups for Oracle Clusterware.

- The Oracle Clusterware files require less than 560 MB of disk space, with external redundancy. To make efficient use of the disk space in a volume group, Oracle recommends that you use the same volume group for the logical volumes for both the Oracle Clusterware files and the database files.

- If you are upgrading an existing Oracle9*i* release 2 Oracle RAC installation that uses raw logical volumes, then you can use the existing SRVM configuration repository logical volume for the OCR and create a new logical volume in the same volume group for the Oracle Clusterware voting disk. However, you must remove this volume group from the HACMP concurrent resource group that activates it before you install Oracle Clusterware.

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

    Note: If you are upgrading a database, then you must also create a new logical volume for the SYSAUX tablespace. Refer to the "Configuring Raw Logical Volumes in the New Oracle Clusterware Volume Group" section on page 4-17 for more information about the requirements for the Oracle Clusterware voting disk and SYSAUX logical volumes.

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

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

- All volume groups that you intend to use for Oracle Clusterware must be activated in concurrent mode before you start the installation.

- The procedures in this section describe how to create basic volumes groups and volumes. If you want to configure more complex volumes, (using mirroring, for example), then use this section in conjunction with the HACMP documentation.

Creating a Volume Group for Oracle Clusterware

To create a volume group for the Oracle Clusterware files:

1. If necessary, install the shared disks that you intend to use.

2. To ensure that the disks are available, enter the following command on every node:

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

   The output from this command is similar to the following:

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

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

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

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

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

The output from this command is similar to the following:

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

For each disk, this command shows:

- The disk device name

- Either the 16 character physical volume identifier (PVID) if the disk has one, or `none`

- Either the volume group to which the disk belongs, or `none`

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

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

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

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

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

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

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

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

7. Identify an appropriate major number that is unused on all nodes in the cluster.

8. To create a volume group, enter a command similar to the following, or use SMIT (`smit mkvg`):

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

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

Command Option	SMIT Field	Sample Value and Description
-y *VGname*	**VOLUME GROUP name**	oracle_vg1
		Specify the name for the volume group. The name that you specify could be a generic name, as shown, or it could specify the name of the database that you intend to create.
-y *VGname*	**VOLUME GROUP name**	oracle_vg1
		Specify the name for the volume group. The name that you specify could be a generic name, as shown, or for a database volume group, it could specify the name of the database that you intend to create.
-B	**Create a big VG format Volume Group**	Specify this option to create a big VG format volume group.
		Note: If you are using SMIT, then choose **yes** for this field.
-s *PPsize*	**Physical partition SIZE in megabytes**	32
		Specify the size of the physical partitions for the database. The sample value shown enables you to include a disk up to 32 GB in size (32 MB * 1016).
-V *Majornum*	**Volume Group MAJOR NUMBER**	46
		Specify the device major number for the volume group that you identified in Step 7.
-n	**Activate volume group AUTOMATICALLY at system restart**	Specify this option to prevent the volume group from being activated at system restart.
		Note: If you are using SMIT, then choose **no** for this field.
-C	**Create VG Concurrent Capable**	Specify this option to create a concurrent capable volume group.
		Note: If you are using SMIT, then choose **yes** for this field.
PhysicalVolumes	**PHYSICAL VOLUME names**	hdisk3 hdisk4
		Specify the device names of the disks that you want to add to the volume group.

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

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

Configuring Raw Logical Volumes in the New Oracle Clusterware Volume Group

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

1. Identify the logical volumes that you must create.

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

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

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

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

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

   ```
   # chown oracle:dba /dev/rora_vote_raw_256m
   # chmod 660 /dev/rora_vote_raw_256m
   # chown root:oinstall /dev/rora_ocr_raw_256m
   # chmod 640 /dev/rora_ocr_raw_256m
   ```

Importing the Volume Group on the Other Cluster Nodes

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

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

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

2. Note the PVIDs of the physical devices used by the volume group.

3. To vary off the volume group that you want to use, enter a command similar to the following on the node where you created it:

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

4. On each cluster node, complete the following steps:

   a. Enter the following command to determine the physical volume names associated with the PVIDs you noted previously:

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

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

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

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

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

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

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

```
# chown oracle:dba /dev/rora_vote_raw_256m
# chmod 660 /dev/rora_vote_raw_256m
# chown root:oinstall /dev/rora_ocr_raw_256m
# chmod 640 /dev/rora_ocr_raw_256m
```

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

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

Activating the Volume Group in Concurrent Mode on All Cluster Nodes

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

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

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:

- Reviewing Storage Options for Oracle Database and Recovery Files

- Checking for Available Shared Storage with CVU

- Choosing a Storage Option for Oracle Database Files

- Configuring Storage for Oracle Database Files on a Supported Shared File System

- Configuring Storage for Oracle Database Files on Shared Storage Devices

- Configuring Disks for Automatic Storage Management

- Configuring Database File Storage on Raw Devices

- Desupport of the Database Configuration Assistant Raw Device Mapping File

- Checking the System Setup with CVU

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/

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.

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

  - **General Parallel File System (GPFS)**: Note that if you intend to use GPFS 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 GPFS, then you must ensure that GPFS volume sizes are at least 500 MB each.

  - **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

- **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.

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 flash 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 ASM, 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
General Parallel File System (GPFS) ■ **Note:** You cannot place ASM files on GPFS.	Yes	Yes

Table 5–1 (Cont.) Supported Storage Options for Oracle Database and Recovery Files

Storage Option	File Types Supported	
	Database	**Recovery**
Local storage	No	No
NFS file system	Yes	Yes
■ **Note:** Requires a certified NAS device		
Raw Logical Volumes Managed by HACMP	Yes	Yes

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.

- 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.

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" on page 5-12

- **To use shared devices for database file storage**, refer to "Configuring Storage for Oracle Database Files on Shared Storage Devices" on page 5-10.

> **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.

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/rhdisk6` and `/dev/rhdisk7`, and your mountpoint is `/dev/dvdrom/`, then enter the following command:

```
$ /mnt/dvdrom/runcluvfy.sh comp ssa -n node1,node2 -s /dev/rhdisk6,/dev/rhdisk7
```

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.

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:

- General Parallel File System (GPFS)

- Network File System (NFS)

- Automatic Storage Management (ASM)

- Raw devices managed by HACMP

During configuration of Oracle Clusterware, if you selected GPFS 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 5-12.

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 5-10.

> **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.

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 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

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 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.

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.

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

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.

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.

Mounting NFS Storage Devices with Direct NFS

Direct NFS determines mount point settings to NFS storage devices based on the configurations in `/etc/filesystems`.

Direct NFS searches for mount entries in the following order:

1. `$ORACLE_HOME/dbs/oranfstab`

2. `/etc/oranfstab`

3. `/etc/filesystems`

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 5-9. 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.

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.

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, `libnfsodm10.so`, to enable Direct NFS. To replace the standard ODM library, `$ORACLE_HOME/lib/libodm10.so`, with the ODM NFS library, `libnfsodm10.so`, complete the following steps:

   a. Change directory to `$ORACLE_HOME/lib`.

   b. Enter the following commands:

   ```
   cp libodm10.so libodm10.so_stub
   ```

```
ln -s libnfsodm10.so libodm10.so
```

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 `libodm10.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.

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 at least 16384. Oracle recommends that you use the value 32768.

If you are using Direct NFS, then set the `rsize` and `wsize` values to 32768. Direct NFS will not serve an NFS server with write size values (`wtmax`) less than 32768.

For example, if you decide to use `rsize` and `wsize` buffer settings with the value 32768, then update the `/etc/fstab` file on each node with an entry similar to the following:

```
nfs_server:/vol/DATA/oradata  /home/oracle/netapp    nfs\
cio,rw,bg,hard,intr,rsize=32768,wsize=32768,tcp,noac,vers=3,timeo=600
```

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

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 -k` command to determine the free disk space on each mounted file system.

3. From the display, identify the file systems that you want to use:

File Type	File System Requirements
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 (flash recovery area):

```
# mkdir /mount_point/flash_recovery_area
# chown oracle:oinstall /mount_point/flash_recovery_area
# chmod 775 /mount_point/flash_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 NFS configuration for Oracle Database shared storage.

Configuring Storage for Oracle Database Files on Shared Storage Devices

The following subsections describe how to configure Oracle Clusterware files on raw devices.

■ Planning Your Shared Storage Device Creation Strategy

■ Identifying Required Shared Partitions for Database Files

Planning Your Shared Storage Device Creation Strategy

Before installing the Oracle Database 11*g* release 1 (11.1) software with Oracle RAC, 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.

Identifying Required Shared Partitions for Database Files

> **Note:** For new installations, Oracle recommends that you do not use raw devices for database files.

Table 5–3 lists the number and size of the shared partitions that you must configure for database files.

Table 5–3   Shared Devices or Logical Volumes Required for Database Files on AIX

Number	Partition Size (MB)	Purpose
1	500	SYSTEM tablespace
1	300 + (Number of instances * 250)	SYSAUX tablespace
Number of instances	500	UNDOTBS*n* tablespace (One tablespace for each instance)
1	250	TEMP tablespace
1	160	EXAMPLE tablespace
1	120	USERS tablespace
2 * number of instances	120	Two online redo log files for each instance
2	110	First and second control files
1	5	Server parameter file (SPFILE)
1	5	Password file

> **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.

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 Database File Storage for Automatic Storage Management and Raw Devices

---

Note: 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 AIX 5L Based Systems (64-Bit)* for information about creating and configuring NAS-based files for use in an Automatic Storage Management disk group.

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

Physical disk paths have path names similar to the following:

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

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.

For Oracle RAC installations, if you choose to enable automated backups and you do not have a shared file system available, 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 flash 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 non-interactive mode, 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 disk space that you require, as follows:

   - External redundancy

     An external redundancy disk group requires a minimum of one disk device. The effective disk space in an external redundancy disk group is the sum of the disk space in all of its devices.

     Because Automatic Storage Management does not mirror data in an external redundancy disk group, Oracle recommends that you use only RAID or similar devices that provide their own data protection mechanisms as disk devices in this type of disk group.

   - Normal redundancy

     In a normal redundancy disk group, Automatic Storage Management uses two-way mirroring by default, to increase performance and reliability. 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 use 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:** You need to 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. 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, 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.

For information about completing this task, refer to the "Configuring Database File Storage for Automatic Storage Management and Raw Devices" section on page 5-16.

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 whether an Automatic Storage Management instance is configured on the system:

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

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

   ```
   +ASM2:oracle_home_path:N
   ```

   In this example, `+ASM2` is the system identifier (SID) of the Automatic Storage Management instance, with 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, Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

Configuring Database File Storage for Automatic Storage Management and Raw Devices

To configure disks for use with Automatic Storage Management on AIX, follow these steps:

1. On AIX-based systems, you must apply Program Technical Fix (PTF) U496549 or higher to your system before you use ASM.

2. If necessary, install the shared disks that you intend to use for the Automatic Storage Management disk group and restart the system.

3. To make sure that the disks are available, enter the following command on every node:

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

The output from this command is similar to the following:

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

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

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

5. Enter the following command on any node to identify the device names for the physical disks that you want to use:

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

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

```
hdisk2      0000078752249812    None
```

In this example, `hdisk2` is the device name of the disk and `0000078752249812` is the physical volume ID (PVID). The disks that you want to use might have a PVID, but they must not belong to a volume group.

6. Enter commands similar to the following to clear the PVID from each disk device that you want to use:

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

> **Note:** If the disk device has a PVID, then ASM will fail to mount the diskgroup created on the disk device.

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

```
# /usr/sbin/lspv | grep -i 0000078752249812
```

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

```
hdisk18         0000078752249812        None
```

Depending on how each node is configured, the device names may differ between nodes. Note that you will clear PVIDs later in this procedure.

8. If the device names are the same on all nodes, then enter commands similar to the following on all nodes to change the owner, group, and permissions on the character raw device files for the disk devices you want to use with ASM datafiles:

```
# chown oracle:dba /dev/rhdiskn
# chmod 660 /dev/rhdiskn
```

9. If PVIDs are configured, and the device name associated with the PVID for a disk that you want to use is different on any node, then you must create a new device file for the disk on each of the nodes using a common unused name.

For the new device files, choose an alternative device file name that identifies the purpose of the disk device. Table 5–4 suggests alternative device file names for each file. For database files, replace *dbname* in the alternative device file name with the name that you chose for the database in step 1.

> **Note:** Alternatively, you can choose a name that contains a number that will never be used on any of the nodes, for example `hdisk99`.

To create a new common device file for a disk device on all nodes, follow these steps on each node:

a. Enter the following command to determine the device major and minor numbers that identify the disk device, where *n* is the disk number for the disk device on this node:

```
# ls -alF /dev/*hdiskn
```

The output from this command is similar to the following:

```
brw------- 1 root system    24,8192 Dec 05 2001  /dev/hdiskn
crw------- 1 root system    24,8192 Dec 05 2001  /dev/rhdiskn
```

In this example, the device file `/dev/rhdisk`*n* represents the character raw device, `24` is the device major number, and `8192` is the device minor number.

b. Enter a command similar to the following to create the new device file, specifying the new device file name and the device major and minor numbers that you identified in the previous step:

Note: In the following example, you must specify the character `c` to create a character raw device file.

```
# mknod /dev/ora_ocr_raw_256m c 24 8192
```

c. Enter commands similar to the following to change the owner, group, and permissions on the character raw device file for the disk:

– OCR:

```
# chown root:oinstall /dev/ora_ocr_raw_256m
# chmod 640 /dev/ora_ocr_raw_256m
```

– Voting disk or database files:

```
# chown oracle:dba /dev/ora_vote_raw_256m
# chmod 660 /dev/ora_vote_raw_256m
```

d. Enter a command similar to the following to verify that you have created the new device file successfully:

```
# ls -alF /dev | grep "24,8192"
```

The output should be similar to the following:

```
brw------- 1 root   system   24,8192 Dec 05 2001  /dev/hdiskn
crw-r----- 1 root   oinstall 24,8192 Dec 05 2001  /dev/ora_ocr_raw_256m
crw------- 1 root   system   24,8192 Dec 05 2001  /dev/rhdiskn
```

10. To enable simultaneous access to a disk device from multiple nodes, you must set the appropriate Object Data Manager (ODM) attribute listed in the following table to the value shown, depending on the disk type:

Disk Type	Attribute	Value
SSA or FAStT disks	reserve_lock	no
ESS, EMC, HDS, CLARiiON, or MPIO-capable disks	reserve_policy	no_reserve

To determine whether the attribute has the correct value, enter a command similar to the following on all cluster nodes for each disk device that you want to use:

```
# /usr/sbin/lsattr -E -l hdiskn
```

If the required attribute is not set to the correct value on any node, enter a command similar to one of the following on that node:

■ SSA and FAStT devices:

```
# /usr/sbin/chdev -l hdiskn  -a reserve_lock=no
```

■ ESS, EMC, HDS, CLARiiON, and MPIO-capable devices:

```
# /usr/sbin/chdev -l hdiskn  -a reserve_policy=no_reserve
```

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

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

12. Enter commands similar to the following on every node to change the owner, group, and permissions on the character raw device file for each disk that you want to add to the disk group:

```
# chown oracle:dba /dev/rhdiskn
# chmod 660 /dev/rhdiskn
```

> **Note:** If you are using a multi-pathing disk driver with ASM, then ensure that you set the permissions only on the correct logical device name for the disk.
>
> The device name associated with a disk may be different on other nodes. Ensure that you specify the correct device name on each node.

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

```
/dev/rhdisk10
```

When you have completed creating and configuring Automatic Storage Management with raw partitions, proceed to Chapter 6, "Installing Oracle Clusterware" on page 6-1

Configuring Database File Storage on Raw Devices

The following subsections describe how to configure raw partitions for database files.

- Identifying Required Raw Partitions for Database Files
- Configuring Raw Disk Devices for Database File Storage Without HACMP or GPFS
- Configuring Raw Logical Volumes for Database File Storage
- Creating a Volume Group for Database Files
- Creating Database File Raw Logical Volumes in the New Volume Group
- Importing the Database File Volume Group on the Other Cluster Nodes
- Activating the Database File Volume Group in Concurrent Mode on All Cluster Nodes

Identifying Required Raw Partitions for Database Files

Table 5–4 lists the number and size of the raw partitions that you must configure for database files.

> **Note:** Because each file requires exclusive use of a complete disk device, Oracle recommends that, if possible, you use disk devices with sizes that closely match the size requirements of the files that they will store. You cannot use the disks that you choose for these files for any other purpose.

Table 5–4   Raw Partitions Required for Database Files on AIX

Number	Partition Size (MB)	Purpose and Sample Alternative Device File Name
1	500	SYSTEM tablespace: *dbname*_system_raw_500m
1	300 + (Number of instances * 250)	SYSAUX tablespace: *dbname*_sysaux_raw_800m
Number of instances	500	UNDOTBS*n* tablespace (One tablespace for each instance, where *n* is the number of the instance): *dbname*_undotbs*n*_raw_500m
1	250	TEMP tablespace: *dbname*_temp_raw_250m
1	160	EXAMPLE tablespace: *dbname*_example_raw_160m
1	120	USERS tablespace: *dbname*_users_raw_120m
2 * number of instances	120	Two online redo log files for each instance (where *n* is the number of the instance and *m* is the log number, 1 or 2): *dbname*_redo*n_m*_raw_120m
2	110	First and second control files: *dbname*_control{1|2}_raw_110m
1	5	Server parameter file (SPFILE): *dbname*_spfile_raw_5m
1	5	Password file: *dbname*_pwdfile_raw_5m

> **Note:** If you prefer to use manual undo management, instead of automatic undo management, then, instead of the UNDOTBS*n* raw devices, you must create a single rollback segment tablespace (RBS) raw device that is at least 500 MB in size.

Configuring Raw Disk Devices for Database File Storage Without HACMP or GPFS

If you are installing Oracle RAC on an AIX cluster without HACMP or GPFS, you can use shared raw disk devices for database file storage. However, Oracle recommends that you use Automatic Storage Management to store database files in this situation. This section describes how to configure the shared raw disk devices for database files.

To configure shared raw disk devices for Oracle Clusterware files, database files, or both:

1. If you intend to use raw disk devices for database file storage, then specify a name for the database that you want to create.

   The name that you specify must start with a letter and have no more than four characters. For example: `orcl`.

2. Identify or configure the required disk devices.

   The disk devices must be shared on all of the cluster nodes.

3. As the `root` user, enter the following command on any node to identify the device names for the disk devices that you want to use:

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

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

   ```
   hdisk17          0009005fb9c23648                    None
   ```

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

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

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

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

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

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

   ```
   hdisk18          0009005fb9c23648                    None
   ```

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

6. If the device names are the same on all nodes, then enter commands similar to the following on all nodes to change the owner, group, and permissions on the character raw device files for the disk devices you want to use for database files:

   ```
   # chown oracle:dba /dev/rhdiskn
   # chmod 660 /dev/rhdiskn
   ```

7. If the device name associated with the PVID for a disk that you want to use is different on any node, then you must create a new device file for the disk on each of the nodes using a common unused name.

   For the new device files, choose an alternative device file name that identifies the purpose of the disk device. Table 5–3 suggests alternative device file names for each file. For database files, replace *dbname* in the alternative device file name with the name that you chose for the database in step 1.

   > **Note:** Alternatively, you could choose a name that contains a number that will never be used on any of the nodes, for example `hdisk99`.

   To create a new common device file for a disk device on all nodes, perform these steps on each node:

a. Enter the following command to determine the device major and minor numbers that identify the disk device, where *n* is the disk number for the disk device on this node:

```
# ls -alF /dev/*hdiskn
```

The output from this command is similar to the following:

```
brw------- 1 root system    24,8192 Dec 05 2001  /dev/hdiskn
crw------- 1 root system    24,8192 Dec 05 2001  /dev/rhdiskn
```

In this example, the device file /dev/rhdisk*n* represents the character raw device, 24 is the device major number, and 8192 is the device minor number.

b. Enter a command similar to the following to create the new device file, specifying the new device file name (in this example, using an alternative device file name from Table 5–3) and the device major and minor numbers that you identified in the previous step:

Note: As the following example illustrates, you must specify the character c to create a character raw device file.

```
# mknod dbname_example_raw_160m c 24 8192
```

c. Enter a command similar to the following to verify that you have created the new device file successfully:

```
# ls -alF /dev | grep "24,8192"
```

The output should be similar to the following:

```
brw------- 1 root   system   24,8192 Dec 05 2001  /dev/hdiskn
crw-r----- 1 root   oinstall 24,8192 Dec 05 2001  /dev/ora_ocr_raw_256m
crw------- 1 root   system   24,8192 Dec 05 2001  /dev/rhdiskn
```

8. To enable simultaneous access to a disk device from multiple nodes, you must set the appropriate Object Data Manager (ODM) attribute listed in the following table to the value shown, depending on the disk type:

Disk Type	Attribute	Value
SSA, FAStT, or non-MPIO-capable disks	reserve_lock	no
ESS, EMC, HDS, CLARiiON, or MPIO-capable disks	reserve_policy	no_reserve

To determine whether the attribute has the correct value, enter a command similar to the following on all cluster nodes for each disk device that you want to use:

```
# /usr/sbin/lsattr -E -l hdiskn
```

If the required attribute is not set to the correct value on any node, then enter a command similar to one of the following on that node:

- SSA and FAStT devices

  ```
  # /usr/sbin/chdev -l hdiskn  -a reserve_lock=no
  ```

- ESS, EMC, HDS, CLARiiON, and MPIO-capable devices

  ```
  # /usr/sbin/chdev -l hdiskn  -a reserve_policy=no_reserve
  ```

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

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

10. If you are using raw disk devices for database files, then follow these steps to create the Database Configuration Assistant raw device mapping file:

> **Note:** You must complete this procedure only if you are using raw devices for database files. The Database Configuration Assistant raw device mapping file enables Database Configuration Assistant to identify the appropriate raw disk device for each database file. You do not specify the raw devices for the Oracle Clusterware files in the Database Configuration Assistant raw device mapping file.

a. Set the ORACLE_BASE environment variable to specify the Oracle base directory that you identified or created previously:

– Bourne, Bash, or Korn shell:

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

– C shell:

```
% setenv ORACLE_BASE /u01/app/oracle
```

b. Create a database file subdirectory under the Oracle base directory and set the appropriate owner, group, and permissions on it:

```
# mkdir -p $ORACLE_BASE/oradata/dbname
# chown -R oracle:oinstall $ORACLE_BASE/oradata
# chmod -R 775 $ORACLE_BASE/oradata
```

In this example, *dbname* is the name of the database that you chose previously.

c. Change directory to the `$ORACLE_BASE/oradata/dbname` directory.

d. Using any text editor, create a text file similar to the following that identifies the disk device file name associated with each database file.

Oracle recommends that you use a file name similar to *dbname*`_raw.conf` for this file.

> **Note:** The following example shows a sample mapping file for a two-instance Oracle RAC cluster. Some of the devices use alternative disk device file names. Ensure that the device file name that you specify identifies the same disk device on all nodes.

```
system=/dev/rhdisk11
sysaux=/dev/rhdisk12
example=/dev/rhdisk13
users=/dev/rhdisk14
temp=/dev/rhdisk15
undotbs1=/dev/rhdisk16
undotbs2=/dev/rhdisk17
redo1_1=/dev/rhdisk18
```

```
redo1_2=/dev/rhdisk19
redo2_1=/dev/rhdisk20
redo2_2=/dev/rhdisk22
control1=/dev/rhdisk23
control2=/dev/rhdisk24
spfile=/dev/dbname_spfile_raw_5m
pwdfile=/dev/dbname_pwdfile_raw_5m
```

In this example, *dbname* is the name of the database.

Use the following guidelines when creating or editing this file:

– Each line in the file must have the following format:

`database_object_identifier=device_file_name`

The alternative device file names suggested in Table 5–4 include the database object identifier that you must use in this mapping file. For example, in the following alternative disk device name, `redo1_1` is the database object identifier:

`/dev/rac_redo1_1_raw_120m`

– For an Oracle RAC database, the file must specify one automatic undo tablespace datafile (`undotbsn`) and two redo log files (`redon_1`, `redon_2`) for each instance.

– Specify at least two control files (`control1`, `control2`).

– To use manual instead of automatic undo management, specify a single RBS tablespace datafile (`rbs`) instead of the automatic undo management tablespace data files.

e. Save the file and note the file name that you specified.

f. When you are configuring the `oracle` user's environment later in this chapter, set the DBCA_RAW_CONFIG environment variable to specify the full path to this file.

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

`/dev/rhdisk10`

Configuring Raw Logical Volumes for Database File Storage

> **Note:** To use raw logical volumes for database file storage, HACMP must be installed and configured on all cluster nodes.

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

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

■ You must use concurrent-capable volume groups for database files.

- Oracle Clusterware files require less than 560 MB of disk space, with external redundancy. To make efficient use of the disk space in a volume group, Oracle recommends that you use the same volume group for the logical volumes for both the Oracle Clusterware files and the database files.

- If you are upgrading a database, then you must also create a new logical volume for the SYSAUX tablespace. Refer to the "Configuring Raw Logical Volumes for Database File Storage" section on page 5-19 for more information about the requirements for the SYSAUX logical volumes.

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

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

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

- All volume groups that you intend to use for database files must be activated in concurrent mode before you start the installation.

- The procedures in this section describe how to create basic volumes groups and volumes. If you want to configure more complex volumes (for example, using mirroring), then use this section in conjunction with the HACMP documentation.

Creating a Volume Group for Database Files

To create a volume group for the Oracle Database files:

1. If necessary, install the shared disks that you intend to use.

2. To ensure that the disks are available, enter the following command on every node:

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

   The output from this command is similar to the following:

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

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

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

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

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

   The output from this command is similar to the following:

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

   For each disk, this command shows:

- The disk device name

- Either the 16 character physical volume identifier (PVID) if the disk has one, or `none`

- Either the volume group to which the disk belongs, or `none`

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

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

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

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

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

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

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

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

7. Identify an appropriate major number that is unused on all nodes in the cluster.

8. To create a volume group, enter a command similar to the following, or use SMIT (`smit mkvg`):

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

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

Command Option	SMIT Field	Sample Value and Description
`-y VGname`	**VOLUME GROUP name**	`oracle_vg1` Specify the name for the volume group. The name that you specify could be a generic name, as shown, or it could specify the name of the database that you intend to create.

Command Option	SMIT Field	Sample Value and Description
`-y VGname`	**VOLUME GROUP name**	`oracle_vg1`
		Specify the name for the volume group. The name that you specify could be a generic name, as shown, or for a database volume group, it could specify the name of the database that you intend to create.
`-B`	**Create a big VG format Volume Group**	Specify this option to create a big VG format volume group.
		Note: If you are using SMIT, then choose **yes** for this field.
`-s PPsize`	**Physical partition SIZE in megabytes**	`32`
		Specify the size of the physical partitions for the database. The sample value shown enables you to include a disk up to 32 GB in size (32 MB * 1016).
`-V Majornum`	**Volume Group MAJOR NUMBER**	`46`
		Specify the device major number for the volume group that you identified in Step 7.
`-n`	**Activate volume group AUTOMATICALLY at system restart**	Specify this option to prevent the volume group from being activated at system restart.
		Note: If you are using SMIT, then choose **no** for this field.
`-C`	**Create VG Concurrent Capable**	Specify this option to create a concurrent capable volume group.
		Note: If you are using SMIT, then choose **yes** for this field.
`PhysicalVolumes`	**PHYSICAL VOLUME names**	`hdisk3 hdisk4`
		Specify the device names of the disks that you want to add to the volume group.

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

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

Creating Database File Raw Logical Volumes in the New Volume Group

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

1. Choose a name for the database that you want to create.

The name that you choose must start with a letter and have no more than four characters, for example, orcl.

2. Identify the logical volumes that you must create.

Table 5–5 lists the number and size of the logical volumes that you must create for database files.

Table 5–5   Raw Logical Volumes Required for Database Files

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	500	SYSTEM tablespace:
		dbname_system_raw_500m
1	500	SYSAUX tablespace:
		dbname_sysaux_raw_500m
1	300 + (Number of instances * 250)	SYSAUX tablespace:
		dbname_sysaux_raw_800m
1	500	UNDOTBS1 tablespace:
		dbname_undotbs1_raw_500m
Number of instances	500	UNDOTBSn tablespace (One tablespace for each instance, where n is the number of the instance):
		dbname_undotbsn_raw_500m
1	250	TEMP tablespace:
		dbname_temp_raw_250m
1	160	EXAMPLE tablespace:
		dbname_example_raw_160m
1	120	USERS tablespace:
		dbname_users_raw_120m
2	120	Two online redo log files (where m is the log number, 1 or 2):
		dbname_redo1_m_raw_120m
2 * number of instances	120	Two online redo log files for each instance (where n is the number of the instance and m is the log number, 1 or 2):
		dbname_redon_m_raw_120m
2	110	First and second control files:
		dbname_control{1|2}_raw_110m
1	5	Server parameter file (SPFILE):
		dbname_spfile_raw_5m

Table 5–5   (Cont.) Raw Logical Volumes Required for Database Files

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	5	Password file: *dbname*_pwdfile_raw_5m

3. To create each required logical volume for data files, Oracle recommends that you use a command similar to the following to create logical volumes with a zero offset:

```
# /usr/sbin/mklv -y LVname -T O -w n -s n -r n VGname NumPPs
```

In this example:

- *LVname* is the name of the logical volume that you want to create

- The `-T O` option specifies that the device subtype should be z, which causes Oracle to use a zero offset when accessing this raw logical volume

- *VGname* is the name of the volume group where you want to create the logical volume

- *NumPPs* is the number of physical partitions to use

   To determine the value to use for *NumPPs*, divide the required size of the logical volume by the size of the physical partition and round the value up to an integer. For example, if the size of the physical partition is 32 MB and you want to create a 500 MB logical volume, then you should specify 16 for the *NumPPs (500/32 = 15.625).*

Using a zero offset improves database performance and fixes the issues described in Oracle bug 2620053.

Note: On raw logical volumes, if you create tablespaces in datafiles that are not created in this way, a message is recorded in the alert.log file.

If you prefer, you can also use the command `smit mklv` to create raw logical volumes.

The following example shows the command used to create a logical volume for the SYSAUX tablespace of the `test` database in the `oracle_vg1` volume group with a physical partition size of 32 MB (800/32 = 25):

```
# /usr/sbin/mklv -y test_sysaux_raw_800m -T O -w n -s n -r n oracle_vg1 25
```

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

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

```
# chown oracle:dba /dev/rdbname*
# chmod 660 /dev/rdbname*
```

Importing the Database File Volume Group on the Other Cluster Nodes

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

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

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

2. Note the PVIDs of the physical devices used by the volume group.

3. To vary off the volume group that you want to use, enter a command similar to the following on the node where you created it:

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

4. On each cluster node, complete the following steps:

   a. Enter the following command to determine the physical volume names associated with the PVIDs you noted previously:

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

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

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

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

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

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

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

      ```
      # chown oracle:dba /dev/rdbname*
      # chmod 660 /dev/rdbname*
      ```

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

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

Activating the Database File Volume Group in Concurrent Mode on All Cluster Nodes

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

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

Creating the Database Configuration Assistant Raw Device Mapping File

> **Note:** You must complete this procedure only if you are using raw devices for database files. You do not specify the raw devices for the Oracle Clusterware files in the Database Configuration Assistant raw device mapping file.

To enable Database Configuration Assistant to identify the appropriate raw device for each database file, you must create a raw device mapping file, as follows:

1. Set the ORACLE_BASE environment variable to specify the Oracle base directory that you identified or created previously:

   - Bourne, Bash, or Korn shell:

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

   - C shell:

     ```
     % setenv ORACLE_BASE /u01/app/oracle
     ```

2. Create a database file subdirectory under the Oracle base directory and set the appropriate owner, group, and permissions on it:

   ```
   # mkdir -p $ORACLE_BASE/oradata/dbname
   # chown -R oracle:oinstall $ORACLE_BASE/oradata
   # chmod -R 775 $ORACLE_BASE/oradata
   ```

   In this example, *dbname* is the name of the database that you chose previously.

3. Change directory to the $ORACLE_BASE/oradata/*dbname* directory.

4. Enter the following command to create a text file that you can use to create the raw device mapping file:

   ```
   # find /dev -user oracle -name 'r*' -print > dbname_raw.conf
   ```

5. Edit the *dbname_raw.conf* file in any text editor to create a file similar to the following:

   > **Note:** The following example shows a sample mapping file for a two-instance Oracle RAC cluster.

   ```
   system=/dev/rdbname_system_raw_500m
   sysaux=/dev/rdbname_sysaux_raw_800m
   example=/dev/rdbname_example_raw_160m
   users=/dev/rdbname_users_raw_120m
   temp=/dev/rdbname_temp_raw_250m
   undotbs1=/dev/rdbname_undotbs1_raw_500m
   undotbs2=/dev/rdbname_undotbs2_raw_500m
   redo1_1=/dev/rdbname_redo1_1_raw_120m
   redo1_2=/dev/rdbname_redo1_2_raw_120m
   redo2_1=/dev/rdbname_redo2_1_raw_120m
   redo2_2=/dev/rdbname_redo2_2_raw_120m
   control1=/dev/rdbname_control1_raw_110m
   control2=/dev/rdbname_control2_raw_110m
   spfile=/dev/rdbname_spfile_raw_5m
   pwdfile=/dev/rdbname_pwdfile_raw_5m
   ```

In this example, *dbname* is the name of the database.

Use the following guidelines when creating or editing this file:

- Each line in the file must have the following format:

  `database_object_identifier=logical_volume`

  The logical volume names suggested in this manual include the database object identifier that you must use in this mapping file. For example, in the following logical volume name, `redo1_1` is the database object identifier:

  `/dev/rrac_redo1_1_raw_120m`

- The file must specify one automatic undo tablespace datafile (`undotbs`*n*) and two redo log files (`redo`*n*`_1`, `redo`*n*`_2`) for each instance.

- Specify at least two control files (`control1`, `control2`).

- To use manual instead of automatic undo management, specify a single RBS tablespace datafile (`rbs`) instead of the automatic undo management tablespace data files.

6. Save the file and note the file name that you specified.

7. When you are configuring the `oracle` user's environment later in this chapter, set the DBCA_RAW_CONFIG environment variable to specify the full path to this file.

Desupport of the Database Configuration Assistant Raw Device Mapping File

With the release of Oracle Database 11g and Oracle RAC release 11g, configuring raw devices using Database Configuration Assistant is not supported.

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 AIX. 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

Verifying Oracle Clusterware Requirements with CVU

Using the following command syntax, log in as the installation owner user (`oracle` or `crs`), and start the 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 CVU 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.

Interpreting CVU Messages About Oracle Clusterware Setup

If the CVU 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 the CVU command 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: The CVU 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 Standard Configuration Operating System Groups and Users" in Chapter 3, and "Configuring SSH on All Cluster Nodes" in Chapter 2 on page 2-18 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 CVU 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 CVU again to determine if the user equivalency error is resolved.

If `ssh` is in a location other than the default, `/usr/bin`, then CVU 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 the CVU check 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 2-23 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 2-24 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 `/usr/sbin/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

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 check 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:** 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. If you use vendor clusterware, then use your vendor documentation to complete setup of your public and private domain addresses.

When you enter the public node name, use the primary host name of each node. In other words, use the name displayed by the `hostname` command. 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

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

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.

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 AIX*. If you intend to use Oracle Clusterware by itself, then refer to the single-instance Oracle Database installation guide.

> **See Also:** *Oracle Database Oracle Clusterware and 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

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 needs 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.

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.

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
```

7

Oracle Clusterware Postinstallation Procedures

This chapter describes how to complete the postinstallation tasks after you have installed the Oracle Clusterware software. It contains the following sections:

- Required Postinstallation Tasks
- Recommended Post-Installation Tasks

Required Postinstallation Tasks

You must perform the following tasks after completing your installation:

- Back Up the Voting Disk After Installation
- Download and Install Patch Updates

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.exe` utility.

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.

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**.

4. On the Select a Patch Search Area page, click **New MetaLink Patch Search**.

5. On the Simple Search page, click **Advanced**.

6. On the Advanced Search page, click the search icon next to the Product or Product Family field.

7. In the Search and Select: Product Family field, enter RDBMS Server in the **For** field, and click **Go**.

8. Select **RDBMS Server** under the Results heading, and click **Select**.

   RDBMS Server appears in the Product or Product Family field. The current release appears in the Release field.

9. Select your platform from the list in the Platform field, and click **Go**.

10. Any available patch updates appear under the Results heading.

11. Click the number of the patch that you want to download.

12. 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.

13. Return to the Patch Set page, click **Download**, and save the file on your system.

14. 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.

Recommended Post-Installation Tasks

Oracle recommends that you complete the following tasks after installing Oracle Clusterware:

- Back Up the root.sh Script
- Run CVU Postinstallation Check

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.

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

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

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
```

Removing Oracle Clusterware

The scripts `rootdelete.sh` and `rootdeinstall.sh` remove Oracle Clusterware from your system.

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 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 filesystem. 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
```

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
```

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.

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

This appendix contains the following topics:

- Install OS Watcher and RACDDT
- General Installation Issues
- Performing Cluster Diagnostics During Oracle Clusterware Installations
- Interconnect Errors

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 needs 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.

General Installation Issues

The following is a list of examples of types of errors that can occur during installation. It contains the following issues:

- An error occurred while trying to get the disks

- Failed to connect to server, Connection refused by server, or Can't open display

- Nodes unavailable for selection from the OUI Node Selection screen

- Node nodename is unreachable

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

- Time stamp is in the future

- 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.

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 2-18.

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 post-installation 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
```

Performing Cluster Diagnostics During Oracle Clusterware Installations

If Oracle Universal Installer (OUI) does not display the Node Selection page, then perform clusterware diagnostics by running the `olsnodes -v` command from the binary directory in your Oracle Clusterware home *(CRS_home*/bin on Linux and UNIX-based systems, and *CRS_home*\BIN on Windows-based systems) and analyzing

its output. Refer to your clusterware documentation if the detailed output indicates that your clusterware is not running.

In addition, use the following command syntax to check the integrity of the Cluster Manager:

```
cluvfy comp clumgr -n node_list -verbose
```

In the preceding syntax example, the variable *node_list* is the list of nodes in your cluster, separated by commas.

Interconnect 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 before initiating an Oracle Clusterware process, it includes information about how to stop processes in Oracle Real Application Clusters (Oracle RAC) databases.

This appendix contains the following topics:

> **Note:** You can use the procedures in this chapter to prepare to perform rolling upgrades of Oracle Clusterware from any Oracle Clusterware 10*g* release 10.2 or Oracle Clusterware 11*g* installation to the latest patch 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 Patches and Upgrades
- How to Stop Processes in an Existing Oracle Database
- How to Perform Oracle Clusterware Rolling Upgrades for Patches
- How to Perform Rolling Upgrades From an Earlier Release to 11g

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.

Restrictions for Clusterware Upgrades to Oracle Clusterware 11*g*

To upgrade existing Oracle Clusterware or Oracle Cluster Ready Services installations to Oracle Clusterware 11*g*, you must first upgrade the existing installations to the most recent patch update. The following sections provide information about Oracle Clusterware upgrades.

Upgrading from Oracle Clusterware 10*g* Release 2 to Oracle Clusterware 11*g*

Complete the following procedure to prepare an existing Oracle Clusterware release 10.2 for rolling upgrade to release 11*g*:

1. Upgrade the Oracle Clusterware home (or CRS home) to Oracle Clusterware release 10.2.0.3, or 10.2.0.2 with CRS Bundle Patch #2 (reference Bug 5256865). Upgrading the database Oracle home is not required.

2. Upgrade Oracle Clusterware from release 10.2.0.3 to release 11.1.

Upgrading from Oracle Cluster Ready Services 10*g* Release 1 to Oracle Clusterware 11*g*

Complete the following procedure to prepare an existing Oracle Cluster Ready Services release 10.1 for rolling upgrade to Oracle Clusterware release 11*g*:

1. Upgrade the Oracle Cluster Ready Services home (or CRS home) to release 10.1.0.3. Upgrading the database Oracle home is not required.

2. Upgrade Oracle Cluster Ready Services from release 10.1.0.3 to Oracle Clusterware release 11.1.

Verify System Readiness for Patches and Upgrades

If you are completing a patch 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 Cluster Verification Utility to assist you with system checks in preparation for starting a database patch or upgrade.

> **See Also:** *Oracle Database Upgrade Guide*

How to Stop Processes in an Existing Oracle Database

To stop process in an existing Oracle RAC database, where you want to shut down the entire database in preparation for an Oracle Clusterware upgrade, complete the following steps.

Shut Down Oracle RAC Databases

Shut down any existing Oracle Database instances on each node, with normal or immediate priority.

If Automatic Storage Management (ASM) is running, then shut down all databases that use ASM, and then shut down the ASM instance on each node of the cluster.

> **Note:** To upgrade using Oracle Clusterware, you must shut down all Oracle Database instances on all cluster nodes before modifying the Oracle software. If you are performing a patch update, review the instructions in the Patch Set Notes for detailed instructions.

Stop All Oracle Processes

Stop all listener and other processes running in the Oracle home directories where you want to modify the database software.

> **Note:** If you shut down ASM instances, then you must first shut down all database instances that use ASM, even if these databases run from different Oracle homes.

Stop Oracle Clusterware Processes

If you are modifying an Oracle Clusterware installation, then shut down the following Oracle Database 10*g* services.

> **Note:** You must perform these steps in the order listed.

1. Shut down any processes in the Oracle home on each node that might be accessing a database; for example, shut down Oracle Enterprise Manager Database Control.

   > **Note:** 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.

2. Shut down all Oracle RAC instances on all nodes. To shut down all Oracle RAC instances for a database, enter the following command, where *db_name* is the name of the database:

   ```
   $ oracle_home/bin/srvctl stop database -d db_name
   ```

3. Shut down all ASM instances on all nodes. To shut down an ASM instance, enter the following command, where *node* is the name of the node where the ASM instance is running:

   ```
   $ oracle_home/bin/srvctl stop asm -n node
   ```

4. 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
   ```

5. Log in as the root user, and shut down the Oracle Clusterware process by entering the following command on all nodes:

   ```
   # CRS_home/bin/crsctl stop crs
   ```

Stop Oracle Database 10g Processes Before Adding Products or Upgrading

This section provides an overview of what needs to be done before adding additional products to Oracle Database 11*g* release 1 (11.1). If you are performing a patch upgrade, then refer to the Database Patch Set Notes for the patch for additional instructions.

> **Note:** You must perform these steps in the order listed.

1. Shut down any processes in the Oracle home on each node that can access a database; for example, shut down Oracle Enterprise Manager Database Control.

> **Note:** 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.

2. Shut down all Oracle RAC instances on all nodes. To shut down all Oracle RAC instances for a database, enter the following command, where *db_name* is the name of the database:

   ```
   $ oracle_home/bin/srvctl stop database -d db_name
   ```

3. Shut down all ASM instances on all nodes. To shut down an ASM instance, enter the following command, where *node* is the name of the node where the ASM instance is running:

   ```
   $ oracle_home/bin/srvctl stop asm -n node
   ```

How to Perform Oracle Clusterware Rolling Upgrades for Patches

To perform a rolling upgrade, complete all of the following steps in sequence.

- Copy Patch Software to the Primary Upgrade Node
- Shut Down Oracle RAC Instances on Upgrade Nodes
- Stop All Oracle Processes on Upgrade Nodes
- Start OUI and Complete Upgrade Processes on Upgrade Nodes

> **Note:** To perform rolling upgrades, the existing Oracle Clusterware home directory, sometimes referred to in Oracle documentation as CRS home, must be located on local directories on the node. You cannot perform rolling upgrades on a shared Oracle Clusterware home directory.
>
> Also note that Oracle does not support attempting to add additional nodes to a cluster during a rolling upgrade.

Copy Patch Software to the Primary Upgrade Node

Download the patch software to the primary node on the cluster (the node where you performed initial Oracle Clusterware and Oracle Database installation). Review the patch set Readme to confirm that your system meets the system requirements for the patch set updates, and complete any special instructions for particular environments or configurations.

To download patches, or download the patch note Readme:

1. Log in to Oracle*MetaLink* (https://metalink.oracle.com)

2. Click the **Patches & Updates** tab.

3. Search for the patch that you want to install.

4. Click the patch number to open the patch page

   From this location, you can download the patch binary, download the patch Readme, and obtain other information regarding the patch update.

5. Download the patch set installation archive to a directory inside the Oracle base directory that meets the following requirements:

   - It is not the existing Oracle home directory, or Oracle Clusterware home directory

   - It is not under an existing Oracle home directory, or Oracle Clusterware home directory

6. Extract the patch set installation archive

Shut Down Oracle RAC Instances on Upgrade Nodes

On each node on which you want to perform a rolling upgrade, shut down the Oracle Database instance, with normal or immediate priority.

If Automatic Storage Management (ASM) is running, then for each node that you intend to perform a rolling upgrade, shut down the database that uses ASM, and then shut down the ASM instance on the node.

> **Note:** To upgrade Oracle Clusterware, you must shut down all Oracle Database instances on all cluster nodes that you intend to upgrade before modifying the Oracle software. If you are performing a patch update, review the instructions in the Patch Set Notes for detailed instructions.
>
> You can use this procedure to shut down one node, and perform upgrades one node at a time, or to shut down groups of nodes, and upgrade groups of nodes at the same time.

Stop All Oracle Processes on Upgrade Nodes

On each node on which you want to perform a rolling upgrade, before you upgrade Oracle Clusterware, you must shut down Oracle Database services that use clusterware processes.

Complete the following steps:

> **Note:** You must perform these steps in the order listed.

1. Shut down any processes on each node you intend to upgrade that might be accessing a database, such as Oracle Enterprise Manager Database Control.

   > **Note:** 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.

2. Shut down all Oracle RAC instances on each node you intend to upgrade. To shut down Oracle RAC instances on individual nodes in the database, enter the

following command, where *db_name* is the name of the database, and *instance_name* is the name of the instance:

```
$ oracle_home/bin/srvctl stop instance -d db_name -i instance_name
```

Repeat this process on each node of the cluster on which you intend to perform the rolling upgrade.

3. Shut down ASM instances on each node on which you intend to perform the rolling upgrade.

> **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.

To shut down an ASM instance, enter the following command, where *node* is the name of the node where the ASM instance is running:

```
$ oracle_home/bin/srvctl stop asm -n node
```

4. Stop all node applications on each node on which you intend to perform the rolling upgrade. 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
```

5. On each node where you intend to perform the rolling upgrade, Initiate the upgrade state by using the following command as the root user:

```
# preupdate.sh -crshome CRS_home -crsuser crs_user
```

In the preceding syntax example, the variable *CRS_home* is the location of the Oracle Clusterware home, and the variable *crs_user* is the Oracle Clusterware installation owner, typically `crs` in Oracle documentation. For example, if the Oracle Clusterware home is `/u01/app/crs`, and the Oracle Clusterware software owner is `crs`, then enter the following command:

```
# preupdate.sh -crshome /u01/app/crs -crsuser crs
```

Start OUI and Complete Upgrade Processes on Upgrade Nodes

To complete the patch upgrade, use the Oracle Universal Installer (OUI) downloaded with the patch update.

Complete the following steps:

1. Start the patch set OUI. At the Welcome window, click **Next**.

2. On the Specify Home Details window, select the Oracle Clusterware home directory, and click **Next**.

3. On the Specify Hardware Cluster Installation Nodes window, select the nodes where you want to perform the upgrade, and click **Next**.

4. Follow further instructions as directed from the OUI windows.

5. At the end of the install, OUI instructs you to run the script `preupdate.sh`.

Open a terminal window, log in as the root user, and run the script as directed.

Wait for the response that the shutdown request has completed successfully.

6. From the terminal window you opened in the previous step, enter a command similar to the following:

    ```
    # CRS_home/install/rootupgrade
    ```

7. When the install script completes, it displays text similar to the following, where *patch_version* displays the patch version you are installing:

    ```
    patch_version patch successfully applied.
    clscfg -upgrade completed successfully
    ```

    This indicates that the upgrade process is complete. The upgraded Oracle Clusterware stack and AUTOSTART resources are started on the node.

    Repeat steps 4 and 5 for each node on which you are performing a rolling upgrade.

How to Perform Rolling Upgrades From an Earlier Release to 11g

Use the following procedure to upgrade Oracle Clusterware from an earlier release to a later release:

> **Note:** You cannot change the owner of the Oracle Clusterware home during an upgrade. You must use the same Oracle software owner that owns the existing Oracle Clusterware home.

1. Mount the Oracle Clusterware 11*g* release 1 (11.1) installation media.

2. Complete operating system prerequisite checks on each of the nodes that you intend to upgrade, to ensure that they meet the system prerequisites for Oracle Clusterware 11*g* release 1 (11.1).

3. If necessary, perform patch upgrades of the previous release Oracle Clusterware or Oracle Cluster Ready Services software to the most recent patch version, as described in the preceding section, "How to Perform Oracle Clusterware Rolling Upgrades for Patches" on page B-4. When you have completed patch upgrades, shut down processes on the nodes you intend to upgrade, as described in the preceding section, "How to Stop Processes in an Existing Oracle Database" on page B-2.

4. From the Oracle Clusterware home, run the following script:

    ```
    mountpoint/clusterware/upgrade/preupdate.sh
    ```

    This script shuts down the Oracle Clusterware software stack, and unlocks the Oracle Clusterware home (CRS home) for upgrading.

    Run this script on each node in the cluster that you want to upgrade.

5. Ensure that you are logged in as the user that you want to own the Oracle Clusterware installation, and run the Oracle Clusterware installation. For example, if the owner you want to use for Oracle Clusterware is `crs`, and the mountpoint for Oracle Clusterware 11*g* release 1 (11.1) is `/shiphome/clusterware`, then you could use the following command:

    ```
    $ whoami
      crs
    $ /shiphome/clusterware/runInstaller
    ```

    Provide information as prompted by the Installer.

> **Note:** You can upgrade one node, a subset of nodes, or all nodes, depending on your upgrade plan. However, note the following information about starting upgrades with OUI:
>
> - **If this is the initial upgrade**, then the node where OUI is running must be one of the set of nodes that is being upgraded.
>
> - **If this is the second or subsequent upgrade**, then the node where OUI is running must be on a node that has not been upgraded.

6. When prompted, in a separate terminal session, log in as root, and run the rootupgrade script.

> **Note:** You must install Oracle Clusterware into the existing Oracle Clusterware home on the local and remote node subset. When you run the root script, it starts the Oracle Clusterware 11g release 1 (11.1) stack on the subset cluster nodes, but lists it as an inactive version.

7. Start the second and subsequent installs of a rolling upgrade from a node that hasn't already been upgraded.

8. After installing the Oracle Clusterware upgrade, if you intend to upgrade Automatic Storage Management, or the Oracle Database software, complete preinstallation procedures for installing Oracle Database or Oracle RAC, as described in *Oracle Database Installation Guide* for your platform, or in this guide (for Oracle RAC preinstallation checks).

> **See Also:** *Oracle Database Upgrade Guide* for additional information about completing database upgrades

Index