

Dell™ PowerEdge™ Systems
Oracle Database 10g R2 on
Red Hat® Enterprise Linux® 4
Advanced Server x86-64

Deployment Guide Version 2.3

Notes and Notices



NOTE: A NOTE indicates important information that helps you make better use of your computer.



NOTICE: A NOTICE indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

Information in this document is subject to change without notice.

© 2007 Dell Inc. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, and *PowerVault* are trademarks of Dell Inc.; *EMC*, *PowerPath*, and *Navisphere* are registered trademarks of EMC Corporation; *Intel* and *Xeon* are registered trademarks of Intel Corporation; *Red Hat* and *Red Hat Enterprise Linux* are registered trademarks of Red Hat, Inc.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

April 2007


Rev. A02

Contents

Oracle 10g RAC Deployment Service	5
Terminology Used in this Document	5
Software and Hardware Requirements	6
Minimum Software Requirements	6
Minimum Hardware Requirements	6
Installing and Configuring Red Hat Enterprise Linux	8
Before You Begin	8
Installing Red Hat Enterprise Linux Using the Deployment CDs	9
Configuring Red Hat Enterprise Linux	10
Updating Your System Packages Using Red Hat Network	11
Verifying Cluster Hardware and Software Configurations	11
Fibre Channel Cluster Setup	12
Cabling Your Fibre Channel Storage System	14
Direct-attached Fibre Channel Configuration	14
SAN-Attached Fibre Channel Configuration	15
SAS Cluster Setup with a PowerVault MD3000	16
Cabling Your SAS Storage System	17
Installing the SAS 5/E Adapter Driver	18
Installing the Host-Based Software Needed for Storage	18
Installing Multi-path Driver Software for MD3000	18
Post Installation Tasks	19
Configuring Storage and Networking for Oracle 10g RAC	19
Configuring the Public and Private Networks	19
Verifying the Storage Configuration	22
Configuring Shared Storage for Oracle Clusterware and the Database Using OCFS2	26
Configuring Shared Storage for Oracle Clusterware and the Database Using ASM	28
Installing Oracle 10g RAC	31
Before You Begin	31
Installing Oracle Clusterware	31
Installing the Oracle Database 10g Software	34

Installing the Oracle Database 10g 10.2.0.3 Patchset	35
Configuring the Listener	38
Creating the Seed Database Using OCFS2	38
Creating the Seed Database Using ASM	40
Securing Your System	42
Disabling rsh	42
Setting the Password for the User Oracle	43
Configuring and Deploying Oracle Database 10g (Single Node)	43
Configuring the Public Network	43
Configuring Database Storage	43
Installing Oracle Database 10g	46
Installing the Oracle Database 10g 10.2.0.3 Patchset	47
Configuring the Listener	48
Creating the Seed Database	48
Adding and Removing Nodes	51
Adding a New Node to the Network Layer	51
Configuring Shared Storage on the New Node Using ASM	51
Adding a New Node to the Oracle Clusterware Layer	52
Adding a New Node to the Database Layer	53
Reconfiguring the Listener	53
Adding a New Node to the Database Instance Layer	54
Removing a Node From the Cluster.	55
Additional Information	57
Supported Software Versions	57
Changing Boot Options	58
Determining the Private Network Interface	58
Troubleshooting	59
Getting Help	65
Dell Support	65
Oracle Support	66
Obtaining and Using Open Source Files.	66
 Index	 67


This document provides information about installing, configuring, reinstalling, and using Oracle Database 10g Enterprise Edition with the Oracle Real Application Clusters (RAC) software on your Dell|Oracle supported configuration. Use this document in conjunction with the *Dell Deployment*, *Red Hat Enterprise Linux*, and *Oracle 10g RAC* software CDs to install your software.

 **NOTE:** If you install your operating system using only the operating system CDs, the steps in this document may not be applicable.

This document covers the following topics:

- Software and hardware requirements
- Installing and configuring Red Hat® Enterprise Linux®
- Verifying cluster hardware and software configurations
- Configuring storage and networking for Oracle RAC
- Installing Oracle RAC
- Configuring and installing Oracle Database 10g (single node)
- Adding and removing nodes
- Reinstalling the software
- Additional information
- Troubleshooting
- Getting help
- Obtaining and using open source files

For more information on Dell supported configurations for Oracle, see the Dell|Oracle Tested and Validated Configurations website at www.dell.com/10g.

 **NOTE:** For Oracle product information, see the *How to Get Started* guide in the Oracle CD kit.

Oracle 10g RAC Deployment Service

If you purchased the Oracle 10g RAC Deployment Service, your Dell Professional Services representative will assist you with the following:

- Verifying cluster hardware and software configurations
- Configuring storage and networking
- Installing Oracle 10g RAC Release 2

Terminology Used in this Document

This document uses the terms *logical unit number* (LUN) and *virtual disk*. These terms are synonymous and can be used interchangeably. The term *LUN* is commonly used in a Dell|EMC Fibre-Channel storage system environment and *virtual disk* is commonly used in a Direct-attached SAS (Dell MD3000) storage environment.

Software and Hardware Requirements

The following sections describe the minimum software and hardware requirements for Dell's Supported Configurations for Oracle. For more information on the minimum software versions for drivers and applications, see "Supported Software Versions."

Minimum Software Requirements

Table 1-1 lists basic software requirements for Dell-supported configurations for Oracle.



 **NOTE:** Your Dell configuration includes a 30-day trial license of Oracle software. If you do not have a license for this product, contact your Dell sales representative.

Table 1-1. Software Requirements

Software Component	Configuration
Red Hat Enterprise Linux AS x86-64 (Version 4)	Update 4
Oracle Database 10g R2	Version 10.2 <ul style="list-style-type: none">• Enterprise Edition, including the RAC option for clusters• Enterprise Edition for single-node configuration
EMC® PowerPath®	Version 4.5.1

Minimum Hardware Requirements

The following sections list the hardware requirements for fibre channel and direct-attached SAS configurations. For more information on specific hardware components, see the documentation included with your system.

 **NOTE:** Depending on the number of users, the applications you use, your batch processes, and other factors, you may need a system that exceeds the minimum hardware requirements to achieve desired performance.

Minimum Hardware Requirements for a Fibre Channel Configuration

Table 1-2 through Table 1-3 indicate the minimum hardware requirements for a fibre channel cluster and single node configuration.


 **NOTE:** The hardware configuration of all the nodes in a cluster must be identical.

Table 1-2. Minimum Hardware Requirements for a Fibre Channel Cluster

Hardware Component	Configuration
Dell™ PowerEdge™ system (two to eight nodes using Automatic Storage Management [ASM])	Intel® Xeon® processor family 1 GB of RAM with Oracle Cluster File System Version 2 (OCFS2) or Automated Storage Management (ASM) PowerEdge Expandable RAID Controller (PERC) for internal hard drives Two 73-GB hard drives (RAID 1) connected to PERC Three Gigabit Network Interface Controller (NIC) ports Two optical host bus adapter (HBA) ports
Dell EMC Fibre Channel storage system	See the Dell Oracle Tested and Validated Configurations website at www.dell.com/10g for information on supported configurations
Gigabit Ethernet switch (two required)	See the Dell Oracle Tested and Validated Configurations website at www.dell.com/10g for information on supported configurations
Dell EMC Fibre Channel switch (two required)	Eight ports for two to six nodes 16 ports for seven or eight nodes

Table 1-3. Minimum Hardware Requirements for a Fibre Channel Single Node

Hardware Component	Configuration
PowerEdge system	Intel Xeon processor family 1 GB of RAM Two 73-GB hard drives (RAID 1) connected to PERC Two NIC ports
Dell EMC Fibre Channel storage system (optional)	See the Dell Oracle Tested and Validated Configurations website at www.dell.com/10g for information on supported configurations
Dell EMC Fibre Channel switch (optional)	Eight ports
Gigabit Ethernet switch (two required)	See www.dell.com/10g for information on supported configurations.

Minimum Hardware Requirements for a Direct-attached SAS Configuration

Table 1-4 through Table 1-5 indicate the minimum hardware requirements for a direct-attached SAS cluster and single node configuration.



NOTE: The hardware configuration of all the nodes in a cluster must be identical.


Table 1-4. Minimum Hardware Requirements for a Direct-attached SAS Cluster

Hardware Component	Configuration
Dell™ PowerEdge™ system (two to eight nodes using Automatic Storage Management [ASM])	Intel® Xeon® processor family 1 GB of RAM with Oracle Cluster File System Version 2 (OCFS2) or Automated Storage Management (ASM) PowerEdge Expandable RAID Controller (PERC) for internal hard drives Two 73-GB hard drives (RAID 1) connected to PERC Three Gigabit Network Interface Controller (NIC) ports Two optical host bus adapter (HBA) ports
Dell™ PowerVault™ MD3000 storage system.	See your Dell PowerVault MD3000 storage system documentation for more details.
Gigabit Ethernet switch (two required)	See www.dell.com/10g for information on supported configurations.

Table 1-5. Minimum Hardware Requirements for a Direct-attached SAS Single Node

Hardware Component	Configuration
PowerEdge system	Intel Xeon processor family 1 GB of RAM Two 73-GB hard drives (RAID 1) connected to PERC Two NIC ports
Dell™ PowerVault™ MD3000 storage system.	See your Dell PowerVault MD3000 storage system documentation for more details.
Gigabit Ethernet switch (two required)	See www.dell.com/10g for information on supported configurations.

Installing and Configuring Red Hat Enterprise Linux

 **NOTICE:** To ensure that the operating system is installed correctly, disconnect all external storage devices from the system *before* you install the operating system.

This section describes the installation of the Red Hat Enterprise Linux AS operating system and the configuration of the operating system for Oracle Database deployment.

Before You Begin

Before you install the Red Hat Enterprise Linux operating system, download the Red Hat Enterprise Linux Quarterly Update ISO images from the Red Hat Network website at rhn.redhat.com and burn these images to CDs.

To download the ISO images, perform the following steps:

- 1 Navigate to the Red Hat Network website at rhn.redhat.com.
- 2 Click **Channels**.
- 3 In the left menu, click **Easy ISOs**.
- 4 In the **Easy ISOs** page left menu, click **All**.
The ISO images for all Red Hat products appear.
- 5 In the **Channel Name** menu, click the appropriate ISO image for your Red Hat Enterprise Linux software.
- 6 Download the ISOs for your Red Hat Enterprise Linux software as listed in your Solution Deliverable List (SDL) from the Dell|Oracle Tested and Validated Configurations website at www.dell.com/10g.
- 7 Burn the ISO images to CDs.

Installing Red Hat Enterprise Linux Using the Deployment CDs

- 1 Disconnect all external storage devices from the system.
- 2 Locate your *Dell Deployment* CD and the *Red Hat Enterprise Linux 4 AS x86-64* CDs.
- 3 Insert the *Dell Deployment* CD 1 into the CD drive and reboot the system.
The system boots to the *Dell Deployment* CD.
A menu asking for the deployment image source appears.
- 4 Type 5 to select **Oracle 10g R2 EE on Red Hat Enterprise Linux 4 x86_64 Update 4**.
- 5 Type 1 to select **Copy solution by Deployment CD**.
- 6 When prompted, insert the *Dell Deployment* CD 2 and each *Red Hat* installation CD into the CD drive.





NOTE: This procedure may take several minutes to complete.

- 6 When prompted, insert the *Dell Deployment* CD 2 and each *Red Hat* installation CD into the CD drive.

A deployment partition is created, and the contents of the CDs are copied to it. When the copy operation completes, the system automatically ejects the last CD and boots to the deployment partition.

The system automatically reboots and the Red Hat Setup Agent appears.

- 7 In the **Red Hat Setup Agent Welcome** window, click **Next** to configure your operating system settings. Do not create any operating system users at this time.
- 8 When prompted, specify a *root* password.



- 9 When the **Network Setup** window appears, click **Next**. You will configure network settings later. A **Security Level** window appears.
- 10 Disable the firewall.
 **NOTE:** You may enable the firewall after completing the Oracle deployment.
- 11 Click through the remaining screens, keeping the default settings.
 **NOTE:** Do not create users in the **Create User** window.
- 12 Log in as `root`.

Configuring Red Hat Enterprise Linux

- 1 Log in as `root` and perform the following steps on all nodes in the cluster.
- 2 Insert the *Dell Deployment* CD 2 into the CD drive and type the following commands:

```
mount /dev/cdrom  
/media/cdrom/install.sh
```

The contents of the CD are copied to the `/usr/lib/dell/dell-deploy-cd` directory. When the copy procedure complete is complete, type `umount /dev/cdrom` and remove the CD from the CD drive.

 **NOTE:** On some systems the CD drive may be recognized as `cdrecorder`.
- 3 Type `cd /dell-oracle-deployment/scripts/standard` to navigate to the directory containing the scripts installed from the *Dell Deployment* CD.
 **NOTE:** Scripts discover and validate installed component versions and, when required, update components to supported levels.
- 4 Type `./005-oraclesetup.py` to configure the Red Hat Enterprise Linux for Oracle installation.
- 5 Type `source /root/.bash_profile` to start the environment variables.
- 6 Type `./010-hwCheck.py` to verify that the CPU, RAM, and disk sizes meet the minimum Oracle Database installation requirements.

If the script reports that a parameter failed, update your hardware configuration and run the script again (see "Minimum Hardware Requirements" on page 6 for updating your hardware configuration).
- 7 Set the password for the user `oracle`.
- 8 Connect the external storage device.

9 Do one of the following depending on whether you are using a fibre channel or serial-attached SCSI (SAS) configuration:

– **Fibre Channel**

Reload the HBA driver(s) using **rmmod** and **modprobe** commands.

Example:

For the Emulex HBAs, reload the lpfc driver by entering:

```
rmmod lpfc
```

```
modprobe lpfc
```

For the QLA HBAs, identify the drivers that are loaded (`lsmod | grep qla`), and reload these drivers.


– **SAS**

For information on setting up a SAS cluster configuration, see "Setting Up Your SAS Cluster with a PowerVault MD3000" on page 15.

 **NOTE:** If you are unable to see the attached storage LUNs after configuring Linux, reboot the system.

Updating Your System Packages Using Red Hat Network

Red Hat periodically releases software updates to fix bugs, address security issues, and add new features. You can download these updates through the Red Hat Network (RHN) service. See the Dell|Oracle Tested and Validated Configurations website at www.dell.com/10g for the latest supported configurations before you use RHN to update your system software to the latest revisions.

 **NOTE:** If you are deploying Oracle Database on a single node, skip the following sections and see "Configuring and Deploying Oracle Database 10g (Single Node)."

Verifying Cluster Hardware and Software Configurations

Before you begin cluster setup, verify the hardware installation, communication interconnections, and node software configuration for the entire cluster. This section provides setup information for hardware and software cluster configurations.

- Each node must include the minimum hardware peripheral components as described in "Minimum Hardware Requirements" on page 6.
- Each node must have the following software installed:
 - Red Hat Enterprise Linux software (see Table 1-1)
 - The Fibre Channel HBA driver (for a fibre channel cluster)
 - Correct version of the Multi-Path driver (for direct-attached SAS). For more information, see the documentation that came with your Dell MD3000 storage system.

- The Fibre Channel or Direct-attached SAS storage system must be configured with the following:
 - A minimum of three LUNs/virtual disks created and assigned to the cluster storage group (see Table 1-6)
 - A minimum LUN/virtual disk size of 5 GB


 **NOTE:** For additional information regarding requirements for assigning virtual disks to a node in a direct-attached SAS configuration, see the documentation that came with your Dell MD3000 storage system.

Table 1-6. LUNs/Virtual Disks for the Cluster Storage Group

LUN	Minimum Size	Number of Partitions	Used For
First LUN/virtual disk	1024 MB	Three of 300 MB each	Voting disk, Oracle Cluster Registry (OCR), and Oracle Server Parameter (SP) file
Second LUN/virtual disk	Larger than the size of your database	One	Database
Third LUN/virtual disk	Minimum twice the size of your second LUN/virtual disk	One	Flash Recovery Area

Fibre Channel Cluster Setup

Your Dell Professional Services representative completed the setup of your Fibre Channel cluster. Verify the hardware connections and the hardware and software configurations as described in this section. Figure 1-1 and Figure 1-3 show an overview of the connections required for the cluster, and Table 1-7 summarizes the cluster connections.

Verify that the following tasks are completed for your cluster:

- All hardware is installed in the rack.
- All hardware interconnections are set up as shown in Figure 1-1 and Figure 1-3, and listed in Table 1-7.
- All Logical Unit Numbers (LUNs), Redundant Array of Independent Disks (RAID) groups, and storage groups are created on the Dell|EMC Fibre Channel storage system.
- Storage groups are assigned to the nodes in the cluster.

Before continuing with the following sections, visually inspect all hardware and interconnections for correct installation.

Figure 1-1. Hardware Connections for a Fibre Channel Cluster

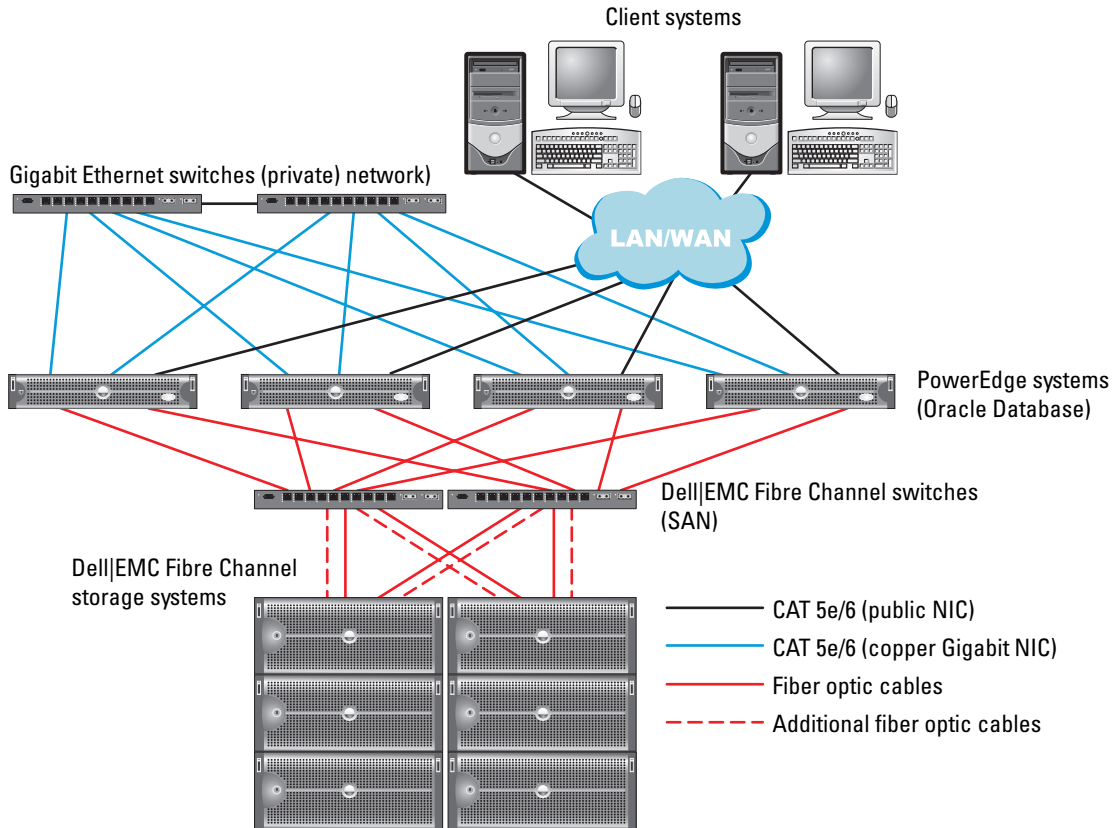


Table 1-7. Fibre Channel Hardware Interconnections

Cluster Component	Connections
Each PowerEdge system node	<p>One Category 5 enhanced (CAT 5e) or CAT 6 cable from the public NIC to the Local Area Network (LAN)</p> <p>One CAT 5e or CAT 6 cable from the private Gigabit NIC to the Gigabit Ethernet switch</p> <p>One CAT 5e or CAT 6 cable from a redundant private Gigabit NIC to a redundant Gigabit Ethernet switch</p> <p>One fiber optic cable from HBA 0 to Fibre Channel switch 0</p> <p>One fiber optic cable from HBA 1 to Fibre Channel switch 1</p>

Table 1-7. Fibre Channel Hardware Interconnections (continued)

Cluster Component	Connections
Each Dell EMC Fibre Channel storage system	Two CAT 5e or CAT 6 cables connected to the LAN One to four fiber optic cable connections to each Fibre Channel switch. For example, for a four-port configuration: <ul style="list-style-type: none">• One fiber optic cable from SPA port 0 to Fibre Channel switch 0• One fiber optic cable from SPA port 1 to Fibre Channel switch 1• One fiber optic cable from SPB port 0 to Fibre Channel switch 1• One fiber optic cable from SPB port 1 to Fibre Channel switch 0
Each Dell EMC Fibre Channel switch	One to four fiber optic cable connections to the Dell EMC Fibre Channel storage system One fiber optic cable connection to each PowerEdge system's HBA
Each Gigabit Ethernet switch	One CAT 5e or CAT 6 connection to the private Gigabit NIC on each PowerEdge system One CAT 5e or CAT 6 connection to the remaining Gigabit Ethernet switch

Cablings Your Fibre Channel Storage System

Depending on your needs, you can configure your Oracle fibre channel cluster storage system in one of the following configurations:

- Direct-attached fibre channel (see Figure 1-2)
- Four-port SAN-attached fibre channel (Figure 1-3)


The following sections describe the cabling requirements for these configurations.

Direct-attached Fibre Channel Configuration

To configure your nodes in a Direct-attached fibre channel configuration (see Figure 1-2), perform the following steps:

- 1** Connect one optical cable from HBA 0 on node 1 to port 0 of SP-A.
- 2** Connect one optical cable from HBA 1 on node 1 to port 0 of SP-B.
- 3** Connect one optical cable from HBA 0 on node 2 to port 1 of SP-A.
- 4** Connect one optical cable from HBA 1 on node 2 to port 1 of SP-B.

Figure 1-2. Cabling in a Direct-attached Fibre Channel Cluster




SAN-Attached Fibre Channel Configuration

To configure your nodes in a four-port SAN-attached configuration (see Figure 1-3), perform the following steps:

- 1 Connect one optical cable from SP-A port 0 to Fibre Channel switch 0.
- 2 Connect one optical cable from SP-A port 1 to Fibre Channel switch 1.
- 3 Connect one optical cable from SP-A port 2 to Fibre Channel switch 0.
- 4 Connect one optical cable from SP-A port 3 to Fibre Channel switch 1.
- 5 Connect one optical cable from SP-B port 0 to Fibre Channel switch 1.
- 6 Connect one optical cable from SP-B port 1 to Fibre Channel switch 0.
- 7 Connect one optical cable from SP-B port 2 to Fibre Channel switch 1.
- 8 Connect one optical cable from SP-B port 3 to Fibre Channel switch 0.
- 9 Connect one optical cable from HBA 0 on node 1 to Fibre Channel switch 0.
- 10 Connect one optical cable from HBA 1 on node 1 to Fibre Channel switch 1.
- 11 Connect one optical cable from HBA 0 on node 2 to Fibre Channel switch 0.
- 12 Connect one optical cable from HBA 1 on node 2 to Fibre Channel switch 1.

Figure 1-3. Cabling in a SAN-Attached Fibre Channel Cluster



SAS Cluster Setup with a PowerVault MD3000

To configure your PowerEdge Systems and PowerVault MD3000 hardware and software to function in a Oracle Real Application Cluster environment, verify the following hardware connections and the hardware and software configurations as described in this section using Figure 1-4, Table 1-8, Figure 1-5 and Table 1-6.

Figure 1-4. Cabling the SAS Cluster and PowerVault MD3000




Table 1-8. SAS Cluster Hardware Interconnections

Cluster Component	Connections
Each PowerEdge system node	One CAT 5e/6 cable from public NIC to the local area network (LAN) One CAT 5e/6 cable from private Gigabit NIC to Gigabit Ethernet switch (private network) One CAT 5e/6 cable from redundant private Gigabit NIC to redundant Gigabit Ethernet switch (private network) Two SAS connections to PowerVault MD3000 system node via SAS 5/E; See “Cabling Your Storage System for a SAN-attached Configuration”
Each Dell PowerVault MD3000	Two CAT 5e/6 cables connected to LAN (one from each storage processor module). Two SAS connections to each PowerEdge system node via SAS 5/E; See “Cabling Your Storage System for a SAN-attached Configuration”
Each Gigabit Ethernet switch	One Cat 5e/6 connection to the private Gigabit NIC on each PowerEdge system One Cat 5e/6 connection to the other Gigabit Ethernet switch


Cabling Your SAS Storage System

SAS clusters can only be installed in a direct-attached cluster, hence they are limited to two nodes only. To configure your nodes in a direct-attached configuration (see Figure 1-5), complete the following steps:

- 1 Connect one SAS cable from a port of the SAS controller of node 1 to the In-0 port of RAID controller 0.
- 2 Connect one SAS cable from the other port of SAS controller of node 1 to the In-0 port of RAID controller 1.

- 3 Connect one SAS cable from a port of the SAS controller of node 2 to the In-1 port of RAID controller 0.
- 4 Connect one SAS cable from the other port of SAS controller of node 2 to the In-1 port of RAID controller 1.

Figure 1-5. Cabling in a Direct-attached SAS Cluster



Installing the SAS 5/E Adapter Driver

Follow the documentation that came with the MD3000 and the SAS HBAs to install drivers on both nodes of the cluster.

NOTE: Confirm that the version of the driver on the *MD3000 Resource* CD is the same as that on the Dell Oracle Solutions Deliverables List (SDL).

Installing the Host-Based Software Needed for Storage


To install the necessary storage software for MD3000, follow the documentation accompanying the storage device. Use the Dell software CD that came with your MD3000 system on one of the two nodes of your cluster. Perform a typical (full installation) on the first node. On the other node, install the only the host-based software.

Installing Multi-path Driver Software for MD3000

- 1 On node 1, install Multi-path driver software.
 - NOTE:** For more information, see the documentation that came with your Dell MD3000 storage system.
- 2 When the installation procedure is completed, restart your system.
- 3 Repeat step 1 and step 2 on the remaining nodes.

Post Installation Tasks

After installing the drivers and the software, perform the post installation tasks listed in the *MD3000 Installation Guide* to create the environment shown in Table 1-6 on page 12.

 **NOTE:** Dell Best Practices dictate that you configure the disks for the LUNS in a RAID 10 configuration.

Configuring Storage and Networking for Oracle 10g RAC

This section provides information and procedures for setting up a Fibre Channel or Direct-attached SAS cluster running a seed database:


- Configuring the public and private networks
- Configuring shared storage for Oracle Clusterware and the database using OCFS2 or ASM

Oracle 10g RAC is a complex database configuration that requires an ordered list of procedures.

To configure networks and storage in a minimal amount of time, perform the following procedures in order.

Configuring the Public and Private Networks

This section presents steps to configure the public and private cluster networks.


 **NOTE:** Each node requires a unique public and private internet protocol (IP) address and an additional public IP address to serve as the virtual IP address for the client connections and connection failover. The virtual IP address must belong to the same subnet as the public IP. All public IP addresses, including the virtual IP address, should be registered with Domain Naming Service (DNS) and routable.

Depending on the number of NIC ports available, configure the interfaces as shown in Table 1-9.

Table 1-9. NIC Port Assignments

NIC Port	Three Ports Available	Four Ports available
1	Public IP and virtual IP	Public IP
2	Private IP (bonded)	Private IP (bonded)
3	Private IP (bonded)	Private IP (bonded)
4	NA	Virtual IP

Configuring the Public Network

 **NOTE:** Ensure that your public IP address is a valid, routable IP address.

If you have not already configured the public network, do so by performing the following steps on *each node*:

- 1 Log in as `root`.
- 2 Edit the network device file `/etc/sysconfig/network-scripts/ifcfg-eth#`

where `#` is the number of the network device.

Configure the file as follows:

```
DEVICE=eth0
ONBOOT=yes
IPADDR=<Public IP Address>
NETMASK=<Subnet mask>
BOOTPROTO=static
HWADDR=<MAC Address>
SLAVE=no
```

- 3 Edit the `/etc/sysconfig/network` file, and, if necessary, replace `localhost.localdomain` with the fully qualified public node name.

For example, the line for node 1 would be as follows:

```
hostname=node1.domain.com
```

- 4 Type:

```
service network restart
```
- 5 Type `ifconfig` to verify that the IP addresses are set correctly.
- 6 To check your network configuration, ping each public IP address from a client on the LAN outside the cluster.
- 7 Connect to each node to verify that the public network is functioning and type `ssh <public IP>` to verify that the secure shell (`ssh`) command is working.

Configuring the Private Network Using Bonding

Before you deploy the cluster, configure the private cluster network to allow the nodes to communicate with each other. This involves configuring network bonding and assigning a private IP address and hostname to each node in the cluster.

To set up network bonding for Broadcom or Intel NICs and configure the private network, perform the following steps on *each node*:

- 1 Log in as `root`.
- 2 Add the following line to the `/etc/modprobe.conf` file:

```
alias bond0 bonding
```
- 3 For high availability, edit the `/etc/modprobe.conf` file and set the option for link monitoring. The default value for `miimon` is 0, which disables link monitoring. Change the value to 100 milliseconds initially, and adjust it as needed to improve performance as shown in the following example.

Type:

```
options bonding miimon=100 mode=1
```

- 4 In the `/etc/sysconfig/network-scripts/` directory, create or edit the `ifcfg-bond0` configuration file. For example, using sample network parameters, the file would appear as follows:

```
DEVICE=bond0
IPADDR=192.168.0.1
NETMASK=255.255.255.0
NETWORK=192.168.0.0
BROADCAST=192.168.0.255
ONBOOT=yes
BOOTPROTO=none
USERCTL=no
```

The entries for `NETMASK`, `NETWORK`, and `BROADCAST` are optional.

`DEVICE=bond n` is the required name for the bond, where n specifies the bond number.

`IPADDR` is the private IP address.

To use `bond0` as a virtual device, you must specify which devices will be bonded as slaves.

- 5 For each device that is a bond member, perform the following steps:
 - a In the directory `/etc/sysconfig/network-scripts/`, edit lines in the `ifcfg-eth n` file as follows:

```
DEVICE=eth $n$ 
HWADDR=<MAC ADDRESS>
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
MASTER=bond0
SLAVE=yes
BOOTPROTO=none
```

b Type

`service network restart` and ignore any warnings.

- 6** On *each node*, type `ifconfig` to verify that the private interface is functioning. The private IP address for the node should be assigned to the private interface `bond0`.
- 7** When the private IP addresses are set up on every node, ping each IP address from one node to ensure that the private network is functioning.
- 8** Connect to each node and verify that the private network and `ssh` are functioning correctly by typing:
`ssh <private IP>`
- 9** On *each node*, modify lines in the `/etc/hosts` file by typing:

```
127.0.0.1      localhost.localdomain  localhost
<private IP node1> <private hostname node1>
<private IP node2> <private hostname node2>
```

```
<public IP node1> <public hostname node1>
<public IP node2> <public hostname node2>
```

```
<virtual IP node1> <virtual hostname node1>
<virtual IP node2> <virtual hostname node2>
```



NOTE: The examples in this and the following step are for a two-node configuration. Add lines for each additional node.

- 10** On *each node*, create or modify the `/etc/hosts.equiv` file by listing all of your public IP addresses or host names. For example, if you have one public hostname, one virtual IP address, and one virtual hostname for each node, add the following lines:

```
<public hostname node1> oracle
<public hostname node2> oracle
```

```
<virtual IP or hostname node1> oracle
<virtual IP or hostname node2> oracle
```

- 11** Log in as `oracle`, connect to each node to verify that the **remote shell (rsh)** command is working by typing:

```
rsh <public hostname nodex>
```

where *x* is the node number.

Verifying the Storage Configuration

The following sections describe how to create and adjust the disk partitions for fibre channel or direct-attached SAS storage.

Creating Disk Partitions on your Storage

While configuring the clusters, create partitions on your Fibre Channel or Direct-attached SAS storage system. To create the partitions, all the nodes must be able to detect the external storage devices.



 **NOTE:** The procedure in this section describes how to deploy Oracle for Direct-attached SAS Storage as well as Fibre-Channel Storage. For illustration purposes, Fibre-Channel Storage nomenclature has been used. If Direct-attached SAS Storage (MD3000) is used, please use the following reference table to translate fibre-channel nomenclature to MD3000 nomenclature.

Table 1-10. Fibre Channel and Direct-attached SAS Nomenclature

Fibre Channel Storage	Direct-attached SAS (MD3000)
LUNs	Virtual disks
/dev/emcpower(X)	/dev/sd(X)
PowerPath	Multi-Path

To verify that each node can detect each storage LUN or logical disk, perform the following steps:

- 1 For Dell|EMC Fibre Channel storage system, verify that the EMC Navisphere[®] agent and the correct version of PowerPath[®] are installed on each node, and that each node is assigned to the correct storage group in your EMC Navisphere software. See the documentation that came with your Dell|EMC Fibre Channel storage system for instructions.

 **NOTE:** The Dell Professional Services representative who installed your cluster performed this step. If you reinstall the software on a node, you must perform this step.

- 2 Visually verify that the storage devices and the nodes are connected correctly to the Fibre Channel switch (see Figure 1-1 and Table 1-7).
- 3 Verify that you are logged in as `root`.
- 4 On *each node*, type:

```
more /proc/partitions
```

The node detects and displays the LUNs or logical disks, as well as the partitions created on those external devices.

 **NOTE:** The listed devices vary depending on how your storage system is configured.

A list of the LUNs or logical disks that are detected by the node is displayed, as well as the partitions that are created on those external devices. PowerPath pseudo devices appear in the list, such as `/dev/emcpowera`, `/dev/emcpowerb`, and `/dev/emcpowerc`.

In the case of a direct-attached SAS configuration, the virtual disks will appear as `/dev/sdb` and `/dev/sdc` and so on.

5 In the `/proc/partitions` file, ensure that:

- All PowerPath pseudo devices appear in the file with similar device names across all nodes.
For example, `/dev/emcpowera`, `/dev/emcpowerb`, and `/dev/emcpowerc`.
- In the case of MD3000, all the virtual disks appear in the file with similar device names across all nodes.
For example, `/dev/sdb`, `/dev/sdc`, and `/dev/sdd`
- The external storage logical volumes appear as SCSI devices, and each node is configured with the same number of LUNs/virtual disks.

For example, if the node is configured with a SCSI drive or RAID container attached to a Fibre Channel storage device with three logical disks, `sda` identifies the node's RAID container or internal drive, and `emcpowera`, `emcpowerb`, and `emcpowerc` identify the LUNs (or PowerPath pseudo devices).

If the node is configured with a SCSI drive or RAID container attached to a Direct-attached SAS storage device with three virtual disks, `sda` identifies the node's RAID container or internal drive, and `sdb`, `sdc`, and `sdd` identify the external storage logical volumes.


6 If the external storage devices do not appear in the `/proc/partitions` file, reboot the node.

Adjusting Disk Partitions for Linux Systems

In Linux, align the partition table before data is written to the LUN/virtual disk, as the partition map will be rewritten and all data on the LUN/virtual disk destroyed.

EXAMPLE: fdisk Utility Arguments

The following example indicates the arguments for the `fdisk` utility. In this example, the LUN is mapped to `/dev/emcpowera`, and the LUN stripe element size is 128 blocks.

 **NOTE:** In this example, the disk `/dev/emcpowera` already has a primary partition `/dev/emcpowera1` created. In the case of the MD3000 this process would need to be performed on `/dev/sdb1`.

```
fdisk /dev/emcpowera
x # expert mode
b # adjust starting block number
l # choose partition 1
128 # set it to 128, (This is the default stripe element size on
Clariion)
w # write the new partition
```

This method is preferable to the LUN alignment offset method for LUNs that will have a snapshot, clone, or MirrorView image made of them. It is also preferred for SAN Copy sources and targets.

Procedure: Using the fdisk Utility to Adjust a Disk Partition

Use the following procedure to use the **fdisk** utility to adjust a disk partition.

- 1 At the command prompt, enter the following:

```
fdisk <PartitionName>
```

where *<PartitionName>* is the name of the partition that you are adjusting. For example, if the partition name is */dev/emcpowera*, you would enter:

```
fdisk /dev/emcpowera
```

The system displays the following message:

```
The number of cylinders for this disk is set to 8782.
```

```
There is nothing wrong with that, but this is larger than 1024,  
and could in certain setups cause problems with:
```

- 1) software that runs at boot time (e.g., old versions of LILO)
- 2) booting and partitioning software from other OSs
(e.g., DOS FDISK, OS/2 FDISK)

- 2 At the command prompt, enter the following **fdisk** utility argument:

```
x
```

- 3 At the command prompt, enter the following **fdisk** utility argument:

```
b
```

- 4 When prompted for the partition number, enter the partition number at the command prompt. For example:

```
1
```

- 5 Specify the new location on the disk partition for the beginning of data. For example:

```
128
```

- 6 At the command prompt, enter the following **fdisk** utility argument:

```
w
```

The system displays the following message:

```
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
```

```
Syncing disks.
```

- 7 Repeat step 1 through step 6 for all Oracle data LUN's.

Configuring Shared Storage for Oracle Clusterware and the Database Using OCFS2

Before You Begin Using OCFS2

- 1 Log in as `root`.
- 2 Navigate to the directory containing the scripts installed from the *Dell Deployment* CD by typing:

```
cd /dell-oracle-deployment/scripts/standard
```
- 3 Install all OCFS packages, by typing:


```
./340-rpms-ocfs.py
```
- 4 Repeat step 1 and step 2 on all other nodes.

Configure Storage Using OCFS2

On the *first node*:

- 1 Log in as `root`.
 - 2 Perform the following steps:
 - a Start the X Window System by typing:

```
startx
```
 - b Generate the OCFS2 configuration file `/etc/ocfs2/cluster.conf` with a default cluster name of `ocfs2` by typing the following in a terminal:

```
ocfs2console
```
 - c From the menu, click **Cluster**→**Configure Nodes**.
If the cluster is offline, the console will start it. A message window appears displaying that information. Close the message window.
The **Node Configuration** window appears.
 - d To add nodes to the cluster, click **Add**. Enter the node name (same as the host name) and the private IP. Retain the default value of the port number. After entering all the details, click **OK**. Repeat this step to add all the nodes to the cluster.
 - e When all the nodes are added, click **Apply** and then click **Close** in the **Node Configuration** window.
-  **NOTE:** If you get the error message: `Unable to access cluster service, delete the file:`

```
/etc/ocfs2/cluster.conf
```


and try again.

f From the menu, click **Cluster**→**Propagate Configuration**.

Propagate Cluster Configuration window appears. Wait until the message **Finished** appears in the window and then click **Close**.

g Select **File**→**Quit**.

3 On *all the nodes*, enable the cluster stack on startup by typing:

```
/etc/init.d/o2cb enable
```

4 Change the `O2CB_HEARTBEAT_THRESHOLD` value on all the nodes using the following steps:

a Stop the O2CB service on all the nodes by typing:

```
/etc/init.d/o2cb stop
```

b Edit the `O2CB_HEARTBEAT_THRESHOLD` value in `/etc/sysconfig/o2cb` to 81 on all the nodes.

c Start the O2CB service on all the nodes by typing:

```
/etc/init.d/o2cb start
```

5 On the *first node*, for a Fibre Channel cluster, create one partition on each of the other two external storage devices with **fdisk**:

a Create a primary partition for the entire device by typing:

```
fdisk /dev/emcpowerX
```

Type **h** for help within the **fdisk** utility.

b Verify that the new partition exists by typing:

```
cat /proc/partitions
```

c If you do not observe the new partition, type:

```
sfdisk -R /dev/<device name>
```



NOTE: The following steps use the sample values:

- `/u01`, `/u02`, and `/u03` for mount points
- `u01`, `u02` & `u03` as labels
- `emcpowera`, `emcpowerb`, and `emcpowerc` as devices for Fibre Channel Storage

6 On *any one node*, format the external storage devices with 4 K block size, 128 K cluster size, and 4 node slots (node slots refer to the number of cluster nodes) using the command line utility **mkfs.ocfs2** as follows:

```
mkfs.ocfs2 -b 4K -C 128K -N 4 -L u01 /dev/emcpowera1
```

```
mkfs.ocfs2 -b 4K -C 128K -N 4 -L u02 /dev/emcpowerb1
```

```
mkfs.ocfs2 -b 4K -C 128K -N 4 -L u03 /dev/emcpowerc1
```



NOTE: For more information about setting the format parameters of clusters, see http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html.

7 On *each node*, perform the following steps:

- a** Create mount points for each OCFS2 partition. To perform this procedure, create the target partition directories and set the ownerships by typing:

```
mkdir -p /u01 /u02 /u03
chown -R oracle.dba /u01 /u02 /u03
```

- b** On each node, modify the `/etc/fstab` by adding the lines for each device:

```
/dev/emcpowera1 /u01 ocfs2 _netdev,datavolume,nointr 0 0
/dev/emcpowerb1 /u02 ocfs2 _netdev,datavolume,nointr 0 0
/dev/emcpowerc1 /u03 ocfs2 _netdev,datavolume,nointr 0 0
```

If the PowerPath pseudo devices do not appear with exactly the same device name across all the nodes, modify the `/etc/fstab` file on each node to ensure that all the shared directories on each node access the same disks.

Make appropriate entries for all OCFS2 volumes.

- c** On *each node*, type the following to mount all the volumes listed in the `/etc/fstab` file:

```
mount -a -t ocfs2
```

- d** On *each node*, add the following command to the `/etc/rc.local` file:

```
mount -a -t ocfs2
```

Configuring Shared Storage for Oracle Clusterware and the Database Using ASM

Configuring Shared Storage for Oracle Clusterware

This section provides instructions for configuring shared storage for Oracle Clusterware.

Configuring Shared Storage Using the RAW Device Interface

- 1** On the *first node*, create three partitions on an external storage device with the `fdisk` utility:

Type:

```
fdisk /dev/emcpowerX
```

and create three partitions of 300 MB each for the Oracle Cluster Repository (OCR), Voting Disk, and the Oracle system parameter file.

- 2** Verify the new partitions by typing:

```
more /proc/partitions
```


On all the nodes, if the new partitions do not appear in the `/proc/partitions` file, type:

```
sfdisk -R /dev/<device name>
```

- 3 On all the nodes in a Fibre Channel cluster, perform the following steps:

Edit the `/etc/sysconfig/rawdevices` file and add the following lines:

```
/dev/raw/votingdisk          /dev/emcpowera1
/dev/raw/ocr.dbf             /dev/emcpowera2
/dev/raw/spfile+ASM.ora     /dev/emcpowera3
```

 **NOTE:** If the three partitions on PowerPath pseudo devices are not consistent across the nodes, modify your `/dev/sysconfig/rawdevices` configuration file accordingly.

- 4 Type `chkconfig networkwait off`.


Configuring Shared Storage for the Database Using ASM

To configure your cluster using ASM, perform the following steps on *all nodes*:

- 1 Log in as root.
- 2 On all the nodes, create one partition on each of the other two external storage devices with the `fdisk` utility:

- a Create a primary partition for the entire device by typing:

```
fdisk /dev/emcpowerX
```

 **NOTE:** Type `h` for help within the `fdisk` utility.


- b Verify that the new partition exists by typing:

```
cat /proc/partitions
```

If you do not see the new partition, type:

```
sfdisk -R /dev/<device name>
```

- 3 Type `chkconfig networkwait off`.

 **NOTE:** Shared storage configuration using ASM can be done either using the RAW device interface or the Oracle ASM library driver.

Configuring Shared Storage Using the RAW Device Interface

- 1 Edit the `/etc/sysconfig/rawdevices` file and add the following lines:

```
/dev/raw/ASM1          /dev/emcpowerb1
/dev/raw/ASM2          /dev/emcpowerc1
```

- 2 Create the RAW devices by typing:

```
udevstart
```

- 3 Restart the RAW Devices Service by typing:

```
service rawdevices restart
```

- 4 To add an additional ASM disk (for example, ASM3), edit the `/etc/udev/scripts/raw-dev.sh` file on all the nodes and add the appropriate bold entries as shown below:

```
MAKEDEV raw
mv /dev/raw/raw1 /dev/raw/votingdisk
mv /dev/raw/raw2 /dev/raw/ocr.dbf
mv /dev/raw/raw3 /dev/raw/spfile+ASM.ora
mv /dev/raw/raw4 /dev/raw/ASM1
mv /dev/raw/raw5 /dev/raw/ASM2
mv /dev/raw/raw6 /dev/raw/ASM3
chmod 660 /dev/raw/raw*
/dev/raw/{votingdisk,ocr.dbf,spfile+ASM.ora,ASM1,ASM2,ASM3}
chown oracle.dba /dev/raw/*
/dev/raw/{votingdisk,ocr.dbf,spfile+ASM.ora,ASM1,ASM2,ASM3}
```

To add additional ASM disks type `udevstart` on all the nodes and repeat step 4.

Configuring Shared Storage Using the ASM Library Driver

- 1 Log in as root.
- 2 Open a terminal window and perform the following steps on all nodes:
 - a Type `service oracleasm configure`
 - b Type the following inputs for all the nodes:
Default user to own the driver interface []: oracle
Default group to own the driver interface []: dba
Start Oracle ASM library driver on boot (y/n) [n]: y
Fix permissions of Oracle ASM disks on boot (y/n) [y]: y

- 3 On the *first node*, in the terminal window, type the following and press <Enter>:

```
service oracleasm createdisk ASM1 /dev/emcpowerb1
service oracleasm createdisk ASM2 /dev/emcpowerc1
```

- 4 Repeat step 3 for any additional ASM disks that need to be created.
- 5 Verify that the ASM disks are created and marked for ASM usage.

In the terminal window, type the following and press <Enter>:

```
service oracleasm listdisks
```

The disks that you created in step 3 appear.

For example:

```
ASM1
```

```
ASM2
```

- 6 Ensure that the remaining nodes are able to access the ASM disks that you created in step 3.

On each remaining node, open a terminal, type the following, and press <Enter>:

```
service oracleasm scandisks
```

Installing Oracle 10g RAC

This section describes the steps required to install Oracle 10g RAC, which involves installing CRS and installing the Oracle Database 10g software. Dell recommends that you create a seed database to verify that the cluster works correctly before you deploy it in a production environment.

Before You Begin

The following sections describe activities that you should complete before installing Oracle 10g RAC.

Download the Red Hat and Dell Deployment CD Images

Before you install the Oracle RAC software on your system:

- Download the *Red Hat* CD images from the Red Hat website at rhn.redhat.com.
- Locate your Oracle CD kit.
- Download the *Dell Deployment* CD images that are appropriate for the solution being installed from the Dell|Oracle Tested and Validated Configurations website at www.dell.com/10g. Burn all these downloaded CD images to CDs.

Configure the System Clock Settings for All Nodes

To prevent failures during the installation procedure, configure all the nodes with identical system clock settings.

Synchronize your node system clock with a Network Time Protocol (NTP) server. If you cannot access an NTP server, perform one of the following procedures:

- Ensure that the system clock on the Oracle Database software installation node is set to a later time than the remaining nodes.
- Configure one of your nodes as an NTP server to synchronize the remaining nodes in the cluster.

Installing Oracle Clusterware

- 1 Log in as `root`.
- 2 Start the X Window System by typing:

```
startx
```
- 3 Open a terminal window and type:

```
xhost +
```
- 4 Mount the *Oracle Clusterware* CD.

5 Type:

```
su - oracle
```

6 Type:

```
<CD_mountpoint>/cluvfy/runcluvfy.sh stage -pre crsinst  
-n node1,node2 -r 10gR2 -verbose
```

where *node1* and *node2* are the public host names.

If your system is *not* configured correctly, troubleshoot the issues and then repeat the `runcluvfy.sh` command, above.

If your system is configured correctly, the following message appears:

```
Pre-check for cluster services setup was successful on all the nodes.
```

7 Type the following commands to start the Oracle Universal Installer:

```
unset ORACLE_HOME  
<CD_mountpoint>/runInstaller
```

The following message appears:

```
Has 'rootpre.sh' been run by root? [y/n] (n)
```

8 Type `y` to proceed.

9 In the **Welcome** window, click **Next**.

10 In the **Specify Home Details** window, change the Oracle home path to `/crs/oracle/product/10.2.0/crs` and click **Next**.

11 In the **Product-Specific Prerequisite Checks** window, ensure that **Succeeded** appears in the **Status** column for each system check, and then click **Next**.

12 In the **Specify Cluster Configuration** window, add the nodes that will be managed by Oracle Clusterware.

a Click **Add**.

b Enter a name for the **Public Node Name**, **Private Node Name**, and **Virtual Host Name**, and then click **OK**.

c Repeat step a and step b for the remaining nodes.

d In the **Cluster Name** field, type a name for your cluster.

The default cluster name is `crs`.

e Click **Next**.

- 13 In the **Specify Network Interface Usage** window, ensure that the public and private interface names are correct.

To modify an interface, perform the following steps:

- a Select the interface name and click **Edit**.
 - b In the **Edit private interconnect type** window in the **Interface Type** box, select the appropriate interface type and then click **OK**.
 - c In the **Specify Network Interface Usage** window, ensure that the public and private interface names are correct, and then click **Next**.
- 14 In the **Specify Oracle Cluster Registry (OCR) Location** window, perform the following steps:

- a In the **OCR Configuration** box, select **External Redundancy**.

- b In the **Specify OCR Location** field, type:

```
/dev/raw/ocr.dbf
```

or

```
/u01/ocr.dbf if using OCFS2.
```

- c Click **Next**.

- 15 In the **Specify Voting Disk Location** window, perform the following steps:

- a In the **OCR Configuration** box, select **External Redundancy**.

- b In the **Specify OCR Location** field, type:

```
/dev/raw/votingdisk
```

or

```
/u01/votingdisk if using OCFS2.
```


- c Click **Next**.


- 16 In the **Summary** window, click **Install**.

Oracle Clusterware is installed on your system.

When complete, the **Execute Configuration scripts** window appears.

- 17 Follow the instructions in the window and then click **OK**.

 **NOTE:** If `root.sh` hangs while formatting the Voting disk, apply Oracle patch 4679769 and then repeat this step.

 **NOTE:** Run `root.sh` sequentially on all nodes starting with the primary node. For example, run `root.sh` on node 1. When `root.sh` completes on node 1, continue with node 2 and so on. Do not run `root.sh` on more than one node at a time.

- 18 In the **Configuration Assistants** window, ensure that **Succeeded** appears in the **Status** column for each tool name.

Next, the **End of Installation** window appears.

- 19 Click **Exit**.
- 20 On *all nodes*, perform the following steps:
 - a Verify the Oracle Clusterware installation by typing the following command:

```
olsnodes -n -v
```

A list of the public node names of all nodes in the cluster appears.
 - b Type:

```
crs_stat -t
```

All running Oracle Clusterware services appear. Ensure that they are all online.

Installing the Oracle Database 10g Software

- 1 As user `root`, type:

```
xhost +
```
- 2 As user `root`, mount the *Oracle Database 10g* CD.
- 3 Log in as `oracle`, and type:

```
cluvfy stage -pre dbinst -n node1,node2 -r 10gR2 -verbose
```

where *node1* and *node2* are the public host names.


If your system is *not* configured correctly, see "Troubleshooting" for more information.

If your system is configured correctly, the following message appears:


```
Pre-check for database installation was successful.
```
- 4 Type:

```
<CD_mountpoint>/runInstaller
```

The Oracle Universal Installer starts.
- 5 In the **Welcome** window, click **Next**.
- 6 In the **Select Installation Type** window, select **Enterprise Edition** and click **Next**.
- 7 In the **Specify Home Details** window in the **Path** field, verify that the complete Oracle home path is `/opt/oracle/product/10.2.0/db_1` and click **Next**.

 **NOTE:** The Oracle home name in this step must be different from the Oracle home name that you identified during the CRS installation. You cannot install the Oracle 10g Enterprise Edition with RAC into the same home name that you used for CRS.
- 8 In the **Specify Hardware Cluster Installation Mode** window, click **Select All** and click **Next**.

- 9 In the **Product-Specific Prerequisite Checks** window, ensure that **Succeeded** appears in the **Status** column for each system check, and then click **Next**.

 **NOTE:** In some cases, a warning may appear regarding swap size. Ignore the warning and click **Yes** to proceed.

- 10 In the **Select Configuration Option** window, select **Install database Software only** and click **Next**.
- 11 In the **Summary** window, click **Install**.
The Oracle Database software is installed on your cluster.
Next, the **Execute Configuration Scripts** window appears.
- 12 Follow the instructions in the window and click **OK**.
- 13 In the **End of Installation** window, click **Exit**.

Installing the Oracle Database 10g 10.2.0.3 Patchset

Downloading and Extracting the Installation Software

- 1 On the *first node*, log in as `oracle`.
- 2 Create a folder for the patches and utilities at `/opt/oracle/patches`.
- 3 Open a web browser and navigate to the Oracle Support website at metalink.oracle.com.
- 4 Log in to your Oracle Metalink account.
- 5 Search for the patch number 5337014 with Linux x86-64 as the platform.
- 6 Download the patch to the `/opt/oracle/patches` directory.
- 7 To unzip the downloaded zip file, type the following in a terminal window and press <Enter>:

```
unzip p5337014_10203_LINUX-x86-64.zip
```

Upgrading Oracle Clusterware Installation

- 1 On the *first node*, log in as `root`.
- 2 Shut down Oracle Clusterware by typing the following in the terminal window and pressing <Enter>:

```
crsctl stop crs
```
- 3 On the remaining nodes, open a terminal window and repeat step 1 and step 2.
- 4 On the *first node*, log in as `oracle`.
- 5 In the terminal window, type the following and press <Enter>:

```
export ORACLE_HOME=/crs/oracle/product/10.2.0/crs
```


- 6 Start the Oracle Universal Installer. To do so, type the following in the terminal window and press <Enter>:

```
cd /opt/oracle/patches/Disk1/  
./runInstaller
```


The **Welcome** screen appears.

- 7 Click **Next**.
- 8 In the **Specify Home Details** screen, click **Next**.
- 9 In the **Specify Hardware Cluster Installation Mode** screen, click **Next**.
- 10 In the **Summary** screen, click **Install**.

The Oracle Universal Installer scans your system, displays all the patches that are required to be installed, and installs them on your system. When the installation completes, the **End of Installation** screen appears.

 **NOTE:** This procedure may take several minutes to complete.

- 11 Read all the instructions that are displayed in the message window, which appears.

 **NOTE:** Do not shut down the Oracle Clusterware daemons, as you already performed this procedure in step 1 and step 2.

- 12 Open a terminal window.
- 13 Log in as `root`.
- 14 Type the following and press <Enter>:


```
$ORA_CRS_HOME/install/root102.sh
```
- 15 Repeat step 12 through step 14 on the remaining nodes, one node at a time.
- 16 On the *first node*, return to the **End of Installation** screen.
- 17 Click **Exit**.
- 18 Click **Yes** to exit the Oracle Universal Installer.

Upgrading the RAC Installation

- 1 On the *first node*, open a terminal window.
- 2 Log in as `oracle`.
- 3 Shut down the Oracle Clusterware node applications on all nodes.

In the terminal window, type the following and press <Enter>:

```
$ORACLE_HOME/bin/srvctl stop nodeapps -n <nodename>
```

 **NOTE:** Ignore any warning messages that may appear.

- 4 Repeat step 3 on the remaining nodes and change the *nodename* of that given node.
- 5 On the *first node*, open a new terminal window.

6 Log in as `oracle`.

7 Type the following and press <Enter>:

```
export ORACLE_HOME=/opt/oracle/product/10.2.0/db_1
```

8 Start the Oracle Universal Installer by typing the following in the terminal window, and pressing <Enter>:

```
cd /opt/oracle/patches/Disk1/  
./runInstaller
```

The **Welcome** screen appears.

9 Click **Next**.

10 In the **Specify Home Details** screen, click **Next**.

11 In the **Specify Hardware Cluster Installation Mode** screen, click **Next**.

12 In the **Summary** screen, click **Install**.

The Oracle Universal Installer scans your system, displays all the patches that are required to be installed, and installs them on your system. When the installation completes, the **End of Installation** screen appears.

A message window appears, prompting you to run `root.sh` as user `root`.


13 Open a terminal window.

14 Type the following and press <Enter>:

```
/opt/oracle/product/10.2.0/db_1/root.sh
```

15 Repeat step 13 and step 14 on the remaining nodes, one node at a time.

When the installation completes, the **End of Installation** screen appears.

 **NOTE:** This procedure may take several minutes to complete.

16 In the **End of Installation** screen, click **Exit**.

17 Click **Yes** to exit the Oracle Universal Installer.

18 On the *first node*, open a terminal window.

19 Log in as `oracle`.

20 Type the following and press <Enter>:

```
srvctl start nodeapps -n <nodename>
```

Where <nodename> is the public host name of the node.

Reconfiguring the CSS Mismatch for Proper EMC PowerPath Failover

When an HBA, switch, or EMC SP failure occurs, the total PowerPath failover time to an alternate device may exceed 105 seconds. To ensure that the PowerPath failover procedure functions correctly, increase the CSS time-out to 120 seconds.

For more information, see *Oracle Metalink Note 294430.1* on the Oracle Metalink website at metalink.oracle.com.

To increase the CSS time-out:

- 1 Shut down the CRS on all nodes except one.
- 2 On the running node, log in as user `root` and type:

```
crsctl set css misscount 120
```
- 3 Reboot all nodes for the CSS setting to take effect.

Configuring the Listener

This section describes the steps to configure the listener, which is required for remote client connection to a database.

On *one node only*, perform the following steps:

- 1 Log in as `root`.
- 2 Start the X Window System by typing:

```
startx
```
- 3 Open a terminal window and type:

```
xhost +
```
- 4 As the user `oracle`, type `netca` to start the Net Configuration Assistant.
- 5 Select **Cluster Configuration** and click **Next**.
- 6 In the **TOPSNodes** window, click **Select All Nodes** and click **Next**.
- 7 In the **Welcome** window, select **Listener Configuration** and click **Next**.
- 8 In the **Listener Configuration**→**Listener** window, select **Add** and click **Next**.
- 9 In the **Listener Configuration**→**Listener Name** window, type `LISTENER` in the **Listener Name** field and click **Next**.
- 10 In the **Listener Configuration**→**Select Protocols** window, select **TCP** and click **Next**.
- 11 In the **Listener Configuration**→**TCP/IP Protocol** window, select **Use the standard port number of 1521** and click **Next**.
- 12 In the **Listener Configuration**→**More Listeners?** window, select **No** and click **Next**.
- 13 In the **Listener Configuration Done** window, click **Next**.
- 14 Click **Finish**.

Creating the Seed Database Using OCFS2

- 1 Log in as `oracle` on the first node and type:

```
cluvfy stage -pre dbcfg -n node1,node2 -d $ORACLE_HOME -verbose
```

where `node1` and `node2` are the public host names.

If your system is *not* configured correctly, see "Troubleshooting" for more information.

If your system is configured correctly, the following message appears:

```
Pre-check for database configuration was successful.
```

- 2 On the *first node*, as user `oracle`, start the Database Configuration Assistant (DBCA) by typing:

```
dbca -datafileDestination /u02
```
- 3 In the **Welcome** window, select **Oracle Real Application Cluster Database** and click **Next**.
- 4 In the **Operations** window, click **Create a Database** and click **Next**.
- 5 In the **Node Selection** window, click **Select All** and click **Next**.
- 6 In the **Database Templates** window, click **Custom Database** and click **Next**.
- 7 In the **Database Identification** window, enter a **Global Database Name** such as `racdb` and click **Next**.
- 8 In the **Management Options** window, click **Next**.
- 9 In the **Database Credentials** window:
 - a Click **Use the same password for all accounts**.
 - b Complete password selections and entries.
 - c Click **Next**.
- 10 In the **Storage Options** window, select **Cluster File System** and click **Next**.
- 11 In the **Database File Locations** window, click **Next**.
- 12 In the **Recovery Configuration** window:
 - a Click **Specify Flash Recovery Area**.
 - b Click **Browse** and select `/u03`.
 - c Specify the flash recovery size.
 - d Select **Enable Archiving**.
 - e Click **Next**.
- 13 In the **Database Content** window, click **Next**.
- 14 In the **Database Services** window, click **Next**.
- 15 In the **Initialization Parameters** window, if your cluster has more than four nodes, change the **Shared Pool** value to `500 MB`, and click **Next**.
- 16 In the **Database Storage** window, click **Next**.
- 17 In the **Creation Options** window, select **Create Database** and click **Finish**.
- 18 In the **Summary** window, click **OK** to create the database.



NOTE: The seed database may take more than an hour to create.

 **NOTE:** If you receive an **Enterprise Manager Configuration Error** during the seed database creation, click **OK** to ignore the error.

When the database creation completes, the **Password Management** window appears.

19 Click **Exit**.

A message appears indicating that the cluster database is starting on all the nodes.

20 On *each node*, perform the following steps:


a Determine the database instance that exists on that node by typing:

```
srvctl status database -d <database name>
```

b Add the `ORACLE_SID` environment variable entry in the user profile `oracle` by typing:

```
echo "export ORACLE_SID=racdbx" >> /home/oracle/.bash_profile
source /home/oracle/.bash_profile
```

where *racdbx* is the database instance identifier assigned to the node.

 **NOTE:** This example assumes that *racdb* is the global database name that you defined in DBCA.

Creating the Seed Database Using ASM

This section contains procedures for creating the seed database using Oracle ASM and for verifying the seed database.

Perform the following steps:

1 Log in as `oracle`, and type:

```
cluvfy stage -pre dbcfg -n node1,node2 -d $ORACLE_HOME -verbose
```

where `node1` and `node2` are the public host names.

If your system is *not* configured correctly, see "Troubleshooting" for more information.

If your system is configured correctly, the following message appears:

```
Pre-check for database configuration was successful.
```

2 On the *first node*, as the user `oracle`, type `dbca &` to start the Oracle Database Creation Assistant (DBCA).

3 In the **Welcome** window, select **Oracle Real Application Cluster Database** and click **Next**.

4 In the **Operations** window, click **Create a Database** and click **Next**.

5 In the **Node Selection** window, click **Select All** and click **Next**.

6 In the **Database Templates** window, click **Custom Database** and click **Next**.

7 In the **Database Identification** window, enter a **Global Database Name**, such as *racdb*, and click **Next**.

8 In the **Management Options** window, click **Next**.

9 In the **Database Credentials** window, select a password option, enter the appropriate password information (if required), and click **Next**.

- 10** In the **Storage Options** window, click **Automatic Storage Management (ASM)** and click **Next**.
- 11** In the **Create ASM Instance** window, perform the following steps:
 - a** In the **SYS password** field, type a password.
 - b** Select **Create server parameter file (SPFILE)**.
 - c** In the **Server Parameter Filename** field, type:
`/dev/raw/spfile+ASM.ora`
 - d** Click **Next**.
- 12** When a message appears indicating that DBCA is ready to create and start the ASM instance, click **OK**.
- 13** Under **ASM Disk Groups**, click **Create New**.
- 14** In the **Create Disk Group** window, perform the following steps:
 - a** Enter a name for the disk group to be created, such as `databaseDG`, select **External Redundancy**, and then select the disks to include in the disk group.
If you are using the RAW device interface, select `/dev/raw/ASM1`.
A window appears indicating that disk group creation is in progress.
 - b** If you are using the ASM library driver and you cannot access candidate disks, click **Change Disk Discovery String**, type `ORCL:*` as the string, and then select `ORCL:ASM1`.
 - c** Click **OK**.
The first ASM disk group is created on your cluster.
Next, the **ASM Disks Groups** window appears.
- 15** Repeat step 14 for the remaining ASM disk group, using `flashbackDG` as the disk group name.
- 16** In the **ASM Disk Groups** window, select the disk group that you would like to use for Database Storage (for example, `databaseDG`) and click **Next**.
- 17** In the **Database File Locations** window, select **Use Oracle-Managed Files** and click **Next**.
- 18** In the **Recovery Configuration** window, click **Browse**, select the flashback group that you created in step 15 (for example, `flashbackDG`), change the Flash Recovery Area size as needed.
- 19** Select **Enable Archiving** and click **Next**.
- 20** In **Database Content** window, click **Next**.
- 21** In **Database Services** window, configure your services (if required) and then click **Next**.
- 22** In the **Initialization Parameters** window, perform the following steps:
 - a** Select **Custom**.
 - b** In **Shared Memory Management**, select **Automatic**.


c In the **SGA Size** and **PGA Size** windows, enter the appropriate information.

d Click **Next**.

23 In the **Database Storage** window, click **Next**.

24 In the **Creation Options** window, select **Create Database** and click **Finish**.

25 In the **Summary** window click **OK** to create the database.

 **NOTE:** This procedure may take an hour or more to complete.

When the database creation completes, the **Database Configuration Assistant** window appears.

26 Click **Password Management** to assign specific passwords to authorized users (if required). Otherwise, click **Exit**.

A message appears indicating that the cluster database is being started on all nodes.

27 Perform the following steps on *each node*:

a Determine the database instance that exists on that node by typing:

```
srvctl status database -d <database name>
```

b Type the following commands to add the `ORACLE_SID` environment variable entry in the *oracle* user profile:

```
echo "export ORACLE_SID=racdbx" >> /home/oracle/.bash_profile  
source /home/oracle/.bash_profile
```

where *racdbx* is the database instance identifier assigned to the node.

This example assumes that *racdb* is the global database name that you defined in DBCA.

28 On *one node*, type:

```
srvctl status database -d dbname
```

where *dbname* is the global identifier name that you defined for the database in DBCA.

If the database instances are running, confirmation appears on the screen.

If the database instances are *not* running, type:

```
srvctl start database -d dbname
```

where *dbname* is the global identifier name that you defined for the database in DBCA.

Securing Your System

Disabling rsh

To prevent unauthorized users from accessing your system, Dell recommends that you disable **rsh** after you install the Oracle software.

To disable `rsh`, type:

```
chkconfig rsh off
```

Setting the Password for the User Oracle

Dell strongly recommends that you set a password for the user *oracle* to protect your system. Complete the following steps to create the *oracle* password:

- 1 Log in as `root`.
- 2 Type `passwd oracle` and follow the instructions on the screen to create the *oracle* password.



NOTE: Additional security setup may be performed according to the site policy, provided the normal database operation is not disrupted.

Configuring and Deploying Oracle Database 10g (Single Node)

This section provides information about completing the initial setup or completing the reinstallation procedures as described in "Installing and Configuring Red Hat Enterprise Linux." This section covers the following topics:

- Configuring the Public Network
- Configuring Database Storage
- Configuring Database Storage Using the Oracle ASM Library Driver
- Installing the Oracle Database 10g
- Installing the Oracle Database 10g 10.2.0.3 Patchset
- Configuring the Listener
- Creating the Seed Database

Configuring the Public Network

Ensure that your public network is functioning and that an IP address and host name are assigned to your system.

Configuring Database Storage

Configuring Database Storage Using ext3 File System

If you have additional storage device, perform the following steps:

- 1 Log in as `root`.
- 2 Type:

```
cd /opt/oracle
```

3 Type:

```
mkdir oradata recovery
```

Using the **fdisk** utility, create a partition where you want to store your database files.

For example:

sdb1 if your storage device is **sdb**

Using the **fdisk** utility, create a partition where you want to store your recovery files.

For example:

sdcl if your storage device is **sdcl**

4 Verify the new partition by typing:

```
cat /proc/partitions
```

If you do not detect the new partition, type:

```
sfdisk -R /dev/sdb
```

```
sfdisk -R /dev/sdc
```

5 Type:

```
mke2fs -j /dev/sdb1
```

```
mke2fs -j /dev/sdc1
```

6 Edit the **/etc/fstab** file for the newly created file system by adding entries such as:

```
/dev/sdb1 /opt/oracle/oradata ext3 defaults 1 2
```

7 Type:

```
chown -R oracle.dba oradata recovery
```

```
/dev/sdc1 /opt/oracle/recovery ext3 defaults 1 2
```

8 Type:

```
mount /dev/sdb1 /opt/oracle/oradata
```

```
mount /dev/sdc1 /opt/oracle/recovery
```

Configuring Database Storage Using Oracle ASM

The following example assumes that you have two storage devices (**sdb** and **sdcl**) available to create a disk group for the database files, and a disk group to be used for flash back recovery and archive log files, respectively.

1 Log in as **root**.

2 Create a primary partition for the entire device by typing:

```
fdisk /dev/sdb
```

3 Create a primary partition for the entire device by typing:

```
fdisk /dev/sdc
```

Configuring ASM Storage Using the RAW Device Interface

- 1 Edit the `/etc/sysconfig/rawdevices` file and add the following lines:
`/dev/raw/ASM1 /dev/sdb1`
`/dev/raw/ASM2 /dev/sdc1`
- 2 Create the RAW Devices by typing:
`udevstart`
- 3 Restart the RAW Devices Service by typing:
`service rawdevices restart`

Configuring Database Storage Using the Oracle ASM Library Driver

- 1 Open a terminal window and log in as `root`.
- 2 Perform the following steps:
 - a Type:
`service oracleasm configure`
 - b Type the following input for all the nodes:
Default user to own the driver interface []: `oracle`
Default group to own the driver interface []: `dba`
Start Oracle ASM library driver on boot (y/n) [n]: `y`
Fix permissions of Oracle ASM disks on boot (y/n) [y]: `y`
- 3 In the terminal window, type the following:
`service oracleasm createdisk ASM1 /dev/sdb1`
`service oracleasm createdisk ASM2 /dev/sdc1`
- 4 Repeat step 1 through step 3 for any additional ASM disks that you need to create.
- 5 Verify that the ASM disks are created and marked for ASM usage.
In the terminal window, type the following and press <Enter>:
`service oracleasm listdisks`
The disks you created in step 3 are listed in the terminal window.
For example:
`ASM1`
`ASM2`

Installing Oracle Database 10g

Perform the following steps to install Oracle 10g:

- 1 Log in as `root`.
- 2 As the user `root`, mount the *Oracle Database 10g* CD.
- 3 Start the X Window System by typing:
`startx`
- 4 Open a terminal window and type:
`xhost +`
- 5 Log in as `oracle`.
- 6 Start the Oracle Universal Installer.
In the terminal window, type the following and press <Enter>:
`<CD_mountpoint>/runInstaller`
- 7 In the **Select Installation Method** window, click **Advanced Installation** and then click **Next**.
- 8 In the **Select Installation Type** window, click **Enterprise Edition** and click **Next**.
- 9 In the **Specify Home Details** window in the **Path** field, ensure that the path is:
`/opt/oracle/product/10.2.0/db_1`
- 10 Click **Next**.
- 11 In the **Product-Specific Prerequisite Checks** window, click **Next**.
- 12 When the **Warning** message appears, ignore the message and click **Yes**.
- 13 In the **Select Configuration Option** window, click **Install Database Software Only**.
- 14 In the **Summary** window, click **Install**.
- 15 When prompted, open a terminal window and run `root.sh`.
A brief progress window appears, followed by the **End of Installation** window.
- 16 Click **Exit** and confirm by clicking **Yes**.
- 17 Log in as `root`.
- 18 Type:
`/opt/oracle/product/10.2.0/db_1/bin/localconfig add`
The installation procedure completes.

Installing the Oracle Database 10g10.2.0.3 Patchset

Downloading and Extracting the Installation Software


- 1 Log in as `oracle`.
- 2 Create a folder for the patches and utilities at `/opt/oracle/patches`.
- 3 Open a web browser and navigate to the Oracle Metalink website at metalink.oracle.com.
- 4 Log in to your Oracle Metalink account.
- 5 Search for the patch number 5337014 with Linux x86-64 as the platform.
- 6 Download the patch to the `/opt/oracle/patches` directory.
- 7 To unzip the downloaded zip file, type the following in a terminal window and press <Enter>:

```
unzip p5337014_10203_LINUX-x86-64.zip
```

Upgrading the Database Software

- 1 As the user `root`, stop the `cssd` process. To do so, type the following and press <Enter>:

```
/etc/init.d/init.cssd stop
```

 **NOTE:** This procedure may take a few minutes to complete.

- 2 Open a terminal window.
- 3 Log in as `oracle`.
- 4 Ensure that `ORACLE_HOME` is set to `/opt/oracle/product/10.2.0/db_1`.
- 5 Start the Oracle Universal Installer. To do so, type the following in the terminal window, and press <Enter>:

```
/opt/oracle/patches/Disk1/runInstaller
```

The **Welcome** screen appears.

- 6 Click **Next**.
- 7 In the **Specify Home Details** screen, click **Next**.
- 8 In the **Specify Hardware Cluster Installation Mode** screen, click **Next**.
- 9 In the **Summary** screen, click **Install**.


The Oracle Universal Installer scans your system, displays all the patches that are required to be installed, and installs them on your system. When the installation completes, the **End of Installation** screen appears.

A message window appears, prompting you to run `root.sh` as user `root`.

- 10 In a terminal window, type the following and press <Enter>:

```
/opt/oracle/product/10.2.0/db_1/root.sh
```

11 Press <Enter> to accept the default answers to the questions generated by `root.sh`.

 **NOTE:** This procedure may take several minutes to complete.

12 When you complete executing `root.sh`, go back to **Execute Configuration Scripts** window and click **OK**.

13 In the **End of Installation** screen, click **Exit**.

14 Click **Yes** to exit the Oracle Universal Installer.

15 Restart the `cssd` process as `root`. To do so, type the following and press <Enter>:

```
/etc/init.d/init.cssd start
```

Configuring the Listener

1 Log in as `root`.

2 Start the X Window System by typing:

```
startx
```

3 Open a terminal window and type:

```
xhost +
```

4 Log in as `oracle`.

5 Type `netca` to start the Oracle Net Configuration Assistant.

6 Accept the default settings and click **Next** on all the screens to complete the listener configuration.

Creating the Seed Database

Creating the Seed Database Using ext3 File System

Perform the following steps to create a seed database with the DBCA:

1 Log in as `oracle`.

2 Start the Oracle DBCA by typing:

```
dbca
```

3 In the **Welcome** window, click **Next**.

4 In the **Operations** window, click **Create a Database** and click **Next**.

5 In the **Database Templates** window, click **Custom Database** and click **Next**.

6 In the **Database Identification** window, type the name of the database that you are creating in the **Global Database Name** and the **SID Prefix** fields, and click **Next**.

7 In the **Management Options** window, click **Next**.

8 In the **Database Credentials** window, complete password selections and entries and click **Next**.

9 In the **Storage Options** window, select **File System** and click **Next**.

10 In the **Database File Locations** window, click **Next**.

- 11 In the **Recovery Configuration** window, click **Browse** and select the flashback recovery area that you created in "Configuring Database Storage Using ext3 File System" (for example, `/opt/oracle/recovery`). Change the Flash Recovery Area size as needed. Select **Enable Archiving** and click **Next**.
- 12 In the **Database Content** window, click **Next**.
- 13 In the **Initialization Parameters** window, click **Next**.
- 14 In the **Database Storage** window, click **Next**.
- 15 In the **Creation Options** window, click **Create Database** and click **Finish**.
- 16 In the **Confirmation** window, click **OK** to create the database.



NOTE: The seed database creation may take more than an hour to complete.

When the database creation procedure completes, the **Password Management** window appears.

- 17 Click **Exit**.

- 18 Type:

```
export ORACLE_SID=dbname
```

where *dbname* is the global identifier name that you defined for the database in DBCA.

- 19 To verify that the database is operating, perform the following steps:

- a Type `sqlplus "/ as sysdba"` to display the SQL> prompt.

- b Type the following query at the SQL> prompt:

```
SELECT * FROM v$instance;
```

- c If the database is not running and you receive an error message, type `startup` at the SQL> prompt to start the database instance on the node.

Creating the Seed Database Using Oracle ASM

If you configured your storage using Oracle ASM, perform the following steps to create a seed database with the DBCA:

- 1 As the user `oracle`, start DBCA by typing:

```
dbca &
```

- 2 In the **Welcome** window, click **Next**.

- 3 In the **Operations** window, click **Create a Database** and click **Next**.

- 4 In the **Database Templates** window, click **Custom Database** and click **Next**.

- 5 In the **Database Identification** window, enter a **Global Database Name** such as `oraadb` and click **Next**.

- 6 In the **Management Options** window, click **Next**.

- 7 In the **Database Credentials** window, click **Use the Same Password for All Accounts**, complete password entries, and click **Next**.

- 8 In the **Storage Options** window, click **ASM** and click **Next**.
- 9 In the **Create ASM Instance** window, enter the password for user `sys` and click **Next**.
- 10 When a message appears indicating that DBCA is ready to create and start the ASM instance, click **OK**.
- 11 In the **ASM Disk Groups** window, under **Available Disk Groups**, click **Create New**.
- 12 In the **Create Disk Group** window, enter the storage information for the database files and click **OK**.
 - a Enter a name for the disk group to be created, such as `databaseDG`, select **External Redundancy**, and select the disks to include in the disk group.
 - b If you are using the RAW device interface, select `/dev/raw/ASM1`.
 - c If you are using the ASM library driver and you cannot access the candidate disks, click **Change Disk Discovery String** and type `ORCL:*` as the string, and then select **ASM1**.
A window appears indicating that disk group creation is in progress.
 - d If you are using the ASM library driver and the candidate disks are not listed, click **Change Disk Discover String** and enter `ORCL:*` as the string.
- 13 Under **Available Disk Groups**, click **Create New**.
- 14 In the **Disk Group** window, enter the information for the flashback recovery files and click **OK**.
 - a Enter a name for the disk group to be created, such as `flashbackDG`, select **External Redundancy**, and select the disks to include in the disk group.
 - b If you are using the RAW device interface, select `/dev/raw/ASM2`.
A window appears indicating that disk group creation is in progress.
 - c If you are using the ASM library driver and you cannot access the candidate disks, click **Change Disk Discovery String**, type `ORCL:*` as the string, and then select **ASM2**.
- 15 In the **ASM Disk Groups** window, check the disk group that you would like to use for Database Storage (for example, `databaseDG`) and click **Next**.
- 16 In the **Database File Locations** window, check **Use Common Location for All Database Files**, and click **Next**.
- 17 In the **Recovery Configuration** window, click **Browse**. Select the flashback group that you created in step 14 (for example, `flashbackDG`) and change the Flash Recovery Area size as needed.
- 18 Select **Enable Archiving** and click **Next**.
- 19 In the **Database Content** window, click **Next**.
- 20 In the **Initialization Parameters** window, select **Typical** and click **Next**.
- 21 In the **Database Storage** window, click **Next**.
- 22 In the **Creation Options** window, select **Create Database** and click **Finish**.
- 23 In the **Confirmation** window click **OK** to create the database.



NOTE: Creating the seed database may take more than an hour.


When the database creation completes, the **Password Management** window appears.

24 Click Exit.

25 When database creation completes, type the following commands to add the ORACLE_SID environment variable entry in the *oracle* user profile:


```
echo "export ORACLE_SID=oradb" >> /home/oracle/.bash_profile
source /home/oracle/.bash_profile
```

This example assumes that *oradb* is the global database name that you defined in DBCA.

 **NOTE:** See the section "Securing Your System" and follow the steps for additional security setup.

Adding and Removing Nodes

This section describes the steps to add a node to an existing cluster and the steps to remove a node from a cluster.

 **NOTE:** The new node must have the same hardware and operating system configuration as the existing node(s).

To add a node to an existing cluster:

- Add the node to the network layer.
- Configure shared storage.
- Add the node to Oracle Clusterware, database, and the database instance layers.

To remove a node from an existing cluster, reverse the process by removing the node from the database instance, the database, and the Oracle Clusterware layers.

For more information about adding an additional node to an existing cluster, see the *Oracle Real Application Clusters 10g Administration* document on the Oracle website at www.oracle.com.

Adding a New Node to the Network Layer

To add a new node to the network layer:

- 1 Install the Red Hat Enterprise Linux operating system on the new node. See "Installing and Configuring Red Hat Enterprise Linux."
- 2 Configure the public and private networks on the new node. See "Configuring the Public and Private Networks."
- 3 Verify that each node can detect the storage LUNs or logical disks. See "Verifying the Storage Configuration."

Configuring Shared Storage on the New Node Using ASM

To extend an existing RAC database to your new nodes, configure storage for the new nodes so that the storage is the same as on the existing nodes. If you are using ASM, ensure that the new nodes can access the ASM disks with the same permissions as the existing nodes. This section provides the appropriate procedure for ASM.

To configure the ASM disks:

- 1 Log in as `root`.
- 2 Open a terminal window and log in as `root`.
- 3 Copy the `/etc/sysconfig/rawdevices` file from one of the existing nodes to the same location on the new node.
- 4 If you are using the RAW device interface for ASM, type `service rawdevices restart` to restart the RAW Devices Service.
- 5 Open a terminal window and perform the following steps on the new node:
 - a Type `service oracleasm configure`
 - b Type the following inputs for all the nodes:
Default user to own the driver interface []: `oracle`
Default group to own the driver interface []: `dba`
Start Oracle ASM library driver on boot (y/n) [n]: `y`
Fix permissions of Oracle ASM disks on boot (y/n) [y]: `y`
 - c Ensure that the new node can access the ASM disks.
 - d In the terminal, type the following and press <Enter>:
`service oracleasm scandisks`
- 6 Ensure that the ASM disks are available on the new node.
In the terminal window, type the following and press <Enter>:
`service oracleasm listdisks`
All available disks on the remaining nodes are listed.
For example:
`ASM1`
`ASM2`

Adding a New Node to the Oracle Clusterware Layer

- 1 Log in as `oracle` into one of the existing nodes.
- 2 From the `/crs/oracle/product/10.2.0/crs/oui/bin` directory of this existing node, type `./addNode.sh` to start the Oracle Universal Installer.
- 3 In the **Welcome** window, click **Next**.
- 4 In the **Specify Cluster Nodes to Add to Installation** window, enter the public and private node names for the new node and click **Next**.
If all the network and storage verification checks pass, the **Cluster Node Addition Summary** window appears.

- 5 Click **Install**.

The **Cluster Node Addition Progress** window displays the status of the node addition process.

- 6 When prompted to run `rootaddnode.sh` on the local node and `root.sh` on the new node as user `root`, run `/crs/oracle/product/10.2.0/crs/install/rootaddnode.sh`.
- 7 When `root.sh` finishes running in the **Execute Configuration Scripts** window, click **OK**.
- 8 In the **End of Cluster Node Addition** window, click **Exit** and click **Yes** in the **Exit** window.

Adding a New Node to the Database Layer

- 1 Log in as `oracle` into one of the existing nodes.
- 2 From the `/opt/oracle/product/10.2.0/db_1/oui/bin` directory of this existing node, type `./addNode.sh` to start the Oracle Universal Installer.
- 3 In the **Welcome** window, click **Next**.
- 4 In the **Specify Cluster Nodes for Node Addition** window, verify that the new node is selected and click **Next**.

If all the verification checks pass, the **Cluster Node Addition Summary** window appears.

- 5 Click **Install**.

The **Cluster Node Addition Progress** window displays the status of the node addition process.

- 6 When prompted, as user `root` run `/opt/oracle/product/10.2.0/db_1/root.sh` on the new node and press <Enter> when asked to enter the full path name of the local bin directory.
- 7 When `root.sh` finishes running in the **Execute Configuration Scripts** window, click **OK**.
- 8 In the **End of Installation** window, click **Exit** and click **Yes** when asked to confirm.

Reconfiguring the Listener

This section describes the steps to reconfigure the listener, which is required for remote client connection to a database.



NOTE: The steps below assume that you are willing to stop the listener to reconfigure the existing listener. Otherwise, the steps may be a little different than the steps below.

On *one node only*, perform the following steps:

- 1 Log in as `root`.
- 2 Start the X Window System by typing:
`startx`
- 3 Open a terminal window and type:
`xhost +`

- 4 As user `oracle`, stop the listener by typing:
`lsnrctl stop`
- 5 When this is successful, type `netca` to start the Net Configuration Assistant.
- 6 Select **Cluster Configuration** and click **Next**.
- 7 In the Real Application Clusters, **Active Nodes** window, click **Select All Nodes** and click **Next**.
- 8 In the **Welcome** window, select **Listener Configuration** and click **Next**.
- 9 In the **Listener Configuration**→**Listener** window, select **Reconfigure** and click **Next**.
- 10 In the **Listener Configuration**→**Select Listener** window, select **LISTENER** from the pull down menu and click **Next**.
- 11 In the **Listener Configuration**→**Select Protocols** window, select **TCP** and click **Next**.
- 12 In the **Listener Configuration**→**TCP/IP Protocol** window, select **Use the standard port number of 1521** and click **Next**.
- 13 In the **Listener Configuration**→**More Listeners?** window, select **No** and click **Next**.
- 14 In the **Listener Configuration Done** window, click **Next**.
- 15 Click **Finish**.

Adding a New Node to the Database Instance Layer

- 1 On one of the existing nodes, as user `oracle`, start DBCA by typing:
`dbca &`
- 2 In the **Welcome** window, click **Next**.
- 3 In the **Operations** window, click **Instance Management** and click **Next**.
- 4 In the **Instance Management** window, click **Add Instance** and click **Next**.
- 5 In the **List of Cluster Databases** window, select the existing database.
If your user name is not operating system-authenticated, the DBCA prompts you for a user name and password for a database user with SYSDBA privileges.
- 6 Enter the user name `sys` and the password, and click **Next**.
The **List of Cluster Database Instances** window appears, showing the instances associated with the RAC database that you selected and the status of each instance.
- 7 In the **List of Cluster Database Instances** window, click **Next**.
- 8 In the **Instance Naming and Node Selection** window, enter the instance name at the top of the window, select the new node name, and click **Next**.
- 9 In the **Instance Storage** window, click **Finish**.
- 10 In the **Summary** window click **OK** to add the database instance.
A progress bar appears, followed by a message asking if you want to extend ASM to the new node(s).

- 11 Click **Yes**.

The following message appears:

Do you want to perform another operation?

- 12 Click **No**.

- 13 On any node, determine that the instance is successfully added by typing:

```
srvctl status database -d <database name>
```



NOTE: See the section "Securing Your System" and follow the steps for additional security setup.

Removing a Node From the Cluster

When you perform the procedures in this section, ensure that you select and remove the correct node from the cluster.

Deleting the Node From the Database Instance Layer

- 1 Log in as `oracle`.

- 2 From one of the remaining nodes, type:

```
dbca &
```

- 3 In the **Welcome** window, click **Next**.

- 4 In the **Operations** window, click **Instance Management** and click **Next**.

- 5 In the **Instance Management** window, click **Delete an instance** and click **Next**.

- 6 In the **List of Cluster Databases** window, select a RAC database from which to delete an instance.

If your user name is not operating system-authenticated, the DBCA prompts you for a user name and password for a database user with SYSDBA privileges.

- 7 Enter the user name `sys` and the password, and click **Next**.

The **List of Cluster Database Instances** window appears, showing the instances associated with the RAC database that you selected and the status of each instance.

- 8 Select the instance to delete and click **Next**.

This instance cannot be the local instance from where you are running DBCA. If you select the local instance, the DBCA displays an **Error** dialog. If this occurs, click **OK**, select another instance, and click **Finish**.

If services are assigned to this instance, the **DBCA Services Management** window appears. Use this window to reassign services to other instances in the cluster database.

- 9 In the **Summary** window, click **OK**.

- 10 Verify the information about the instance deletion operation and click **OK**.

A progress bar appears while DBCA removes the instance and its Oracle Net configuration.

When the operation completes, a dialog prompts whether you want to perform another operation.

- 11 Click **No** to exit.
- 12 Verify that the node was removed by typing:

```
srvctl config database -d <database name>
```

Reconfiguring the Listener

- 1 Type `netca`.
- 2 In the **Real Application Clusters**→ **Configuration** window, select **Cluster Configuration** and click **Next**.
- 3 In the **Real Application Clusters**→ **Active Nodes** window, select the node that you want to delete and click **Next**.
- 4 In the **Welcome** window, select **Listener Configuration** and click **Next**.
- 5 In the **Listener Configuration**→ **Listener** window, select **Delete** and click **Next**.
- 6 In the **Listener Configuration**→ **Select Listener** window, select **LISTENER** and click **Next**.
When a message `Are you sure you want to delete listener LISTENER?` appears, click **Yes**.
- 7 In the **Listener Configuration**→ **Listener Deleted** window, click **Next**.
- 8 In the **Listener Configuration Done** window, click **Next**.
- 9 Click **Finish**.

To Stop and Remove ASM From the Node That is Deleted

On one of the remaining nodes, perform the following steps:

- 1 Open a terminal window.
- 2 Type:

```
srvctl stop asm -n <node_name>
```


where `<node_name>` is the node you want to remove from the cluster.
- 3 Type:

```
srvctl remove asm -n <node_name>
```


where `<node_name>` is the node you want to remove from the cluster.

Deleting a Node From the Database Layer

- 1 On the node being deleted, log in as `oracle`.
- 2 Type the following command, using the public name of the node you are deleting (for example, if you are removing `node3-pub`):

```
srvctl stop nodeapps -n node3-pub
```

Ignore error CRS-0210 which complains about the listener.

- 3 On the node being deleted, log in as `root`.
- 4 If you wish to remove the Oracle Database software, type the following command:

```
rm -rf /opt/oracle/product/10.2.0/db_1/*
```

Removing a Node From the Oracle Clusterware Layer

- 1 On the node that you are deleting, as user `root`, disable CRS by typing the following command:

```
/crs/oracle/product/10.2.0/crs/install/rootdelete.sh remote nosharedvar
```

- 2 On one of the remaining nodes, as user `root`, type the following command:

```
/crs/oracle/product/10.2.0/crs/install/rootdeletenode.sh  
<public nodename>, <node-number>
```

Where *<public-nodename>* is public name and *<node-number>* is the node number of the node being deleted.

To determine the node number of any node, type the following command:

```
/crs/oracle/product/10.2.0/crs/bin/olsnodes -n
```

- 3 On the node that you are deleting, if you wish to remove the Oracle CRS software, type the following command:

```
rm -rf /crs/oracle/product/10.2.0/crs/*
```

Additional Information

Supported Software Versions

For the latest supported hardware and software, see the Dell|Oracle Tested and Validated Configuration website at www.dell.com/10g.

Changing Boot Options

Grub, the method Red Hat Linux uses to launch operating systems on Dell PowerEdge servers, uses various command and configuration options that allow you to control the operating system boot process. Grub is started during the server boot process. To change Red Hat Linux boot options, complete the following steps:

- 1 Boot the server.
A prompt displays on the server console requiring the user to change boot options.
- 2 Press any console keyboard key to continue
- 3 Press `p` to issue the password command to Grub.
You are prompted for the password.
- 4 Enter the default password of `del11` or your chosen password if you have changed the default.
Grub can now display the current boot options and enables you to edit these options as required.

Determining the Private Network Interface

To determine which interface device name is assigned to each network interface, perform the following steps:

- 1 Determine the types of NICs in your system.
See Table 1-11 to identify which integrated NICs are present in your system.
For add-in NICs, you may have Intel PRO/100 family or PRO/1000 family cards or Broadcom NetXtreme Gigabit cards. You may have to open your system and view the add-in cards to identify your card.

Table 1-11. Integrated NICs

System	Integrated NICs	Driver Name
PowerEdge 1950	Broadcom NetXtreme II BCM5708	bnx2
PowerEdge 2950	Broadcom NetXtreme II BCM5708	bnx2
PowerEdge 2900	Broadcom NetXtreme II BCM5708	bnx2
PowerEdge 2970	Broadcom NetXtreme II BCM5708	bnx2
PowerEdge 1850	Intel PRO/1000	e1000
PowerEdge 2850	Intel PRO/1000	e1000
PowerEdge 6850	Broadcom NetXtreme BCM5704	tg3
PowerEdge 6950	Broadcom NetXtreme II BCM5708	bnx2

- 2 Verify that a Broadcom NetXtreme Gigabit or Intel PRO/1000 family NIC is connected with a Cat 5e cable to the Gigabit Ethernet switch, which is your private NIC.
- 3 Determine the driver module your private NIC uses (see Table 1-11).

- 4 View the `/etc/modprobe.conf` file by typing:

```
more /etc/modprobe.conf
```

Several lines appear with the format `alias ethx driver-module`, where *x* is the Ethernet interface number and *driver-module* is the module you determined in step 3.

For example, the line `alias eth1 tg3` appears if your operating system assigned `eth1` to a Broadcom NetXtreme Gigabit NIC.

- 5 Determine which Ethernet interfaces (`ethx`) are assigned to the type of Gigabit NIC that is connected to the Gigabit switch.

If only one entry exists in `/etc/modprobe.conf` for your driver module type, then you have successfully identified the private network interface.

- 6 If you have more than one of the same type of NIC in your system, experiment to determine which Ethernet interface is assigned to each NIC.

For each Ethernet interface, follow the steps in "Configuring the Private Network Using Bonding" for the correct driver module until you have identified the correct Ethernet interface.

Troubleshooting

Table 1-12 provides recommended actions for problems that you may encounter while deploying and using your Red Hat Enterprise Linux and Oracle software.

Table 1-12. Troubleshooting

Category	Problem / Symptom	Cause	Recommended Corrective Action
Database	Nodes that start up correctly with Patchset 10.2.0.1 may fail to start up with ORA-4031 errors in Patchset 10.2.0.2.	RAC instances with greater than 4 Gigabit <code>db_cache_sizes</code> .	Set the variable <code>set _ksmsg_granule_size=16777216</code> in the <code>init.ora</code> file.
Database	Lock Manager Service (LMS) crash with ORA-00600 error: internal error code, arguments: <code>[kclastf_1], [2], []</code> .	Due to Oracle bug 5071492. See the Oracle Metalink website at metalink.oracle.com .	Apply patch 5071492 available on the Oracle Metalink website at metalink.oracle.com .
Database	The instance can terminate with ORA-600 error <code>[kclcls_5]</code> in the RAC instance.	Due to Oracle bug 4639236. See the Oracle Metalink website at metalink.oracle.com .	Apply patch 4639236 available on the Oracle Metalink website at metalink.oracle.com .

Table 1-12. Troubleshooting (continued)

Category	Problem / Symptom	Cause	Recommended Corrective Action
Database	ERROR IN KQLMBIVG SEE LCK TRACE FILE" [LT] [LB] KJUSERCLIENT LOCK	Due to Oracle bug 4690794. See the Oracle Metalink website at metalink.oracle.com .	Apply patch 4690794 available on the Oracle Metalink website at metalink.oracle.com .
Database	LMD0 PROCESS RECEIVED OS SIGNAL #11	Due to Oracle bug 5036588. See the Oracle Metalink website at metalink.oracle.com .	Apply patch 5036588 available on the Oracle Metalink website at metalink.oracle.com .
Performance and stability	Red Hat Enterprise Linux exhibiting poor performance and instability. Excessive use of swap space.	The Oracle System Global Area (SGA) exceeds the recommended size.	<ul style="list-style-type: none"> • Ensure that the SGA size does not exceed 65% of total system RAM. • Type <code>free</code> at a command prompt to determine total RAM and reduce the values of <code>db_cache_size</code> and <code>shared_pool_size</code> parameters in the Oracle parameter file accordingly.
Enterprise Manager	The Enterprise Manager agent goes down or fails.	The Enterprise Manager repository is not populated.	Type the following to recreate the configuration file and repository for the DB Console: <code>emca -config dbcontrol db repos recreate</code> For detailed instructions, see Oracle Metalink Note 330976.1.
OCFS2	System hangs while mounting or unmounting OCFS partitions.	Two nodes tried to unmount OCFS partitions at exactly the same time.	Do not reboot more than one system at the same time. NOTE: It is not recommended to restart the network on a live node. When trying to restart the network service from any live node in the cluster, the node hangs indefinitely. This is an expected behavior of OCFS2.

Table 1-12. Troubleshooting (continued)

Category	Problem / Symptom	Cause	Recommended Corrective Action
Performance and stability	Unknown interface type warning appears in Oracle alert file. Poor system performance.	The public interface is configured as cluster communications (private interface).	Force cluster communications to the private interface by performing the following steps on <i>one node</i> : <ol style="list-style-type: none"> 1 Log in as <code>oracle</code>. 2 Type <code>sqlplus "/ as sysdba"</code> at the command prompt. The <code>SQL></code> prompt appears. 3 Enter the following lines at the <code>SQL></code> prompt: <pre>alter system set cluster_interconnects=<private IP address node1>' scope=spfile sid='<SID1>' alter system set cluster_interconnects='<private IP address node2>' scope=spfile sid='<SID2>'</pre> Continue entering lines for each node in the cluster. 4 Restart the database on all nodes by typing the following lines: <pre>srvctl stop database -d <dbname> srvctl start database -d <dbname></pre> 5 Open the <code>/opt/oracle/admin/<dbname>/bdump/alert_<SID>.log</code> file and verify that the private IP addresses are being used for all instances.
NETCA	NETCA fails, resulting in database creation errors.	The public network, hostname, or virtual IP is not listed in the <code>/etc/hosts.equiv</code> file.	Before launching <code>netca</code> , ensure that a hostname is assigned to the public network and that the public and virtual IP addresses are listed in the <code>/etc/hosts.equiv</code> file.
NETCA	NETCA cannot configure remote nodes or a RAW device validation error occurs while running DBCA.	The <code>/etc/hosts.equiv</code> file either does not exist or does not include the assigned public or virtual IP addresses.	Verify that the <code>/etc/hosts.equiv</code> file on each node contains the correct public and virtual IP address. Try to <code>rsh</code> to other public names and VIP addresses as the user <code>oracle</code> .
CRS	Frequent session open/close messages in <code>/var/log/messages</code>	The cause of this problem has been identified and verified in an unpublished Oracle Bug 5679560. See the Oracle Metalink website at metalink.oracle.com .	Apply Patch 5679560 available on the Oracle Metalink website at metalink.oracle.com .

Table 1-12. Troubleshooting (continued)

Category	Problem / Symptom	Cause	Recommended Corrective Action
CRS	CRS gives up prematurely when trying to start.	Due to Oracle bug 4698419. See the Oracle Metalink website at metalink.oracle.com .	Apply patch 4698419 available on the Oracle Metalink website at metalink.oracle.com .
CRS	The Oracle Clusterware installation procedure fails.	EMC PowerPath device names are not uniform across the nodes.	Before you install Oracle Clusterware, restart PowerPath and ensure that the PowerPath device names are uniform across the nodes.
CRS	CRS fails to start when you reboot the nodes, or type <code>/etc/init.d/init.crs start</code> .	The Cluster Ready Services CSS daemon cannot write to the quorum disk.	<ul style="list-style-type: none"> • Attempt to start the service again by rebooting the node or typing <code>root .sh</code> from <code>/crs/oracle/product/10.2.0/crs/</code>. • Verify that each node has access to the quorum disk and the user <code>root</code> can write to the disk. • Check the last line in the file <code>\$ORA_CRS_HOME/css/log/ocssd.log</code>. • If you see <code>clssnmvWriteBlocks: Failed to flush writes to (votingdisk)</code>, verify the following: <ul style="list-style-type: none"> – The <code>/etc/hosts</code> file on each node contains correct IP addresses for all node hostnames, including the virtual IP addresses. – You can ping the public and private hostnames. – The quorum disk is writable.
CRS	When you run <code>root.sh</code> , CRS fails to start.	Check and make sure you have public and private node names defined and that you can ping the node names.	Attempt to start the service again by rebooting the node or by running <code>root.sh</code> from <code>/crs/oracle/product/10.2.0/crs/</code> after correcting the networking issues.
CRS	When you run <code>root.sh</code> , CRS fails to start.	The OCR file and Voting disk are inaccessible.	Correct the I/O problem and attempt to start the service again by rebooting the node or by running <code>root.sh</code> from <code>/crs/oracle/product/10.2.0/crs/</code> .

Table 1-12. Troubleshooting (continued)

Category	Problem / Symptom	Cause	Recommended Corrective Action
CRS	When you run <code>root.sh</code> following reinstallation, CRS fails to start.	The OCR file and Voting disk have not been cleared and contain old information.	<ol style="list-style-type: none"> 1 Clear the OCR and Voting disks by typing the following lines: <pre>dd if=/dev/zero of=/dev/raw/ocr.dbf</pre> <pre>dd if=/dev/zero of=/dev/raw/votingdisk</pre> 2 Attempt to start the service again by rebooting the node or by running <code>root.sh</code> from <code>/crs/oracle/product/10.2.0/crs/</code>
CRS	When you run <code>root.sh</code> , CRS fails to start.	The user <code>oracle</code> does not have permissions on <code>/var/tmp</code> (specifically <code>/var/tmp/.oracle</code>).	<ol style="list-style-type: none"> 1 Make user <code>oracle</code> the owner of <code>/var/tmp/.oracle</code> by typing <code>chown oracle.oinstall /var/tmp/.oracle</code> 2 Attempt to start the service again by rebooting the node or by running <code>root.sh</code> from: <code>/crs/oracle/product/10.2.0/crs/</code>
CRS	When you run <code>root.sh</code> , CRS fails to start.	Other CRS troubleshooting steps are attempted without success.	<ol style="list-style-type: none"> 1 Enable debugging by adding the following line to <code>root.sh</code>: <pre>set -x</pre> 2 Attempt to start the service again by running <code>root.sh</code> from: <code>/crs/oracle/product/10.2.0/crs/</code> 3 Check log files in the following directories to diagnose the issue: <pre>\$ORA_CRS_HOME/crs/log</pre> <pre>\$ORA_CRS_HOME/crs/init</pre> <pre>\$ORA_CRS_HOME/css/log</pre> <pre>\$ORA_CRS_HOME/css/init</pre> <pre>\$ORA_CRS_HOME/evm/log</pre> <pre>\$ORA_CRS_HOME/evm/init</pre> <pre>\$ORA_CRS_HOME/srvm/log</pre> 4 Check <code>/var/log/messages</code> for any error messages regarding CRS init scripts. 5 Capture all log files for support diagnosis.
CRS	Node continually reboots.	The node does not have access to the quorum disk on shared storage.	<ol style="list-style-type: none"> 1 Start Linux in single user mode. 2 Type: <pre>/etc/inet.d/init.crs disable</pre> 3 Verify that the quorum disk is available and the private interconnect is alive. 4 Reboot and type <code>/etc/inet.d/init.crs enable</code>

Table 1-12. Troubleshooting (continued)

Category	Problem / Symptom	Cause	Recommended Corrective Action
CRS	Node continually reboots.	The private interconnect is down.	<ol style="list-style-type: none"> 1 Start Linux in single user mode. 2 Type: <code>/etc/inet.d/init.crs disable</code> 3 Verify that the node can ping over the private interconnect to the remaining nodes in the cluster. 4 Type: <code>/etc/inet.d/init.crs enable</code> 5 Reboot your system. 6 In some cases, the network has a latency of up to 30 seconds before it can ping the remaining nodes in the cluster after reboot. If this situation occurs, add the following line to the beginning of your <code>/etc/inet.d/init.crs</code> file and reboot your system: <code>/bin/sleep 30</code>
DBCA	There is no response when you click OK in the DBCA Summary window.	Java Runtime Environment timing issue.	Click again. If there is still no response, restart DBCA.
Software installation	You receive <code>dd</code> failure error messages while installing the software using <i>Dell Deployment CD 1</i> .	Using copies, rather than the original Red Hat CDs.	When burning the CD images (ISOs), use the proper options such as <code>-dao</code> if using <code>cdrecord</code> command.
Software installation	When connecting to the database as a user other than <i>oracle</i> , you receive the error messages <code>ORA01034: ORACLE not available and Linux Error 13: Permission denied</code> .	Required permissions are not set on the remote node.	On all remote nodes, as user <code>root</code> , type: <code>chmod 6751 \$ORACLE_HOME</code>

Table 1-12. Troubleshooting (continued)

Category	Problem / Symptom	Cause	Recommended Corrective Action
Software installation	Oracle software fails to install on the nodes.	The nodes system clocks are not identical.	Perform one of the following procedures: <ul style="list-style-type: none"> • Ensure that the system clock on the Oracle software installation node is set to a later time than the remaining nodes. • Configure one of your nodes as an NTP server to synchronize the remaining nodes in the cluster.
Software installation	When you run <code>root.sh</code> , the utility fails to format the OCR disk.	The utility fails to format the OCR disk. This issue is documented in Oracle Metalink under bug 4679769.	Download and apply Oracle patch 4679769, found on the Oracle Metalink website at metalink.oracle.com .
Networking	The cluster verification check fails.	Your public network IP address is not routable. For example: 192.168.xxx.xxx	Assign a valid, routable public IP address.
Fibre Channel storage system	You receive I/O errors and warnings when you load the Fibre Channel HBA driver module.	The HBA driver, BIOS, or firmware needs to be updated.	Check the Solution Deliverable List on the Dell Oracle Tested and Validated Configurations website at www.dell.com/10g for the supported versions. Update as required the driver, BIOS, and firmware for the Fibre Channel HBAs.
Operating System	When you add a new peripheral device to your PowerEdge system, the operating system does not recognize the device.	Kudzu is disabled.	Manually run Kudzu after you add the new peripheral to your system.

Getting Help

Dell Support

For detailed information on the use of your system, see the documentation that came with your system components.

For white papers, Dell supported configurations, and general information, visit the Dell|Oracle Tested and Validated Configurations website at www.dell.com/10g.

For Dell technical support for your hardware and operating system software and to download the latest updates for your system, visit the Dell Support website at support.dell.com. Information about contacting Dell is provided in your system's *Installation and Troubleshooting Guide*.

Dell Enterprise Training and Certification is now available; see www.dell.com/training for more information. This training service may not be offered in all locations.

Oracle Support

For training information on your Oracle software and application Clusterware, see the Oracle website at www.oracle.com or see your Oracle documentation for information about contacting Oracle.

Technical support, downloads, and other technical information are available at the Oracle MetaLink website at metalink.oracle.com.

Obtaining and Using Open Source Files

The software contained on the *Dell Deployment* CD is an aggregate of third-party programs as well as Dell programs. Use of the software is subject to designated license terms. All software that is designated as "under the terms of the GNU GPL" may be copied, distributed, and/or modified in accordance with the terms and conditions of the GNU General Public License, Version 2, June 1991. All software that is designated as "under the terms of the GNU LGPL" (or "Lesser GPL") may be copied, distributed, and/or modified in accordance with the terms and conditions of the GNU Lesser General Public License, Version 2.1, February 1999. Under these GNU licenses, you are also entitled to obtain the corresponding source files by contacting Dell at 1-800-WWW-DELL. Please see SKU 420-4534 when making such request. You may be charged a nominal fee for the physical act of transferring a copy.

Index

A

ASM, 19, 28-29, 44, 51
library driver, 30, 43, 45
remove from node, 56
using RAW, 45

C

cabling

fibre channel storage, 14
SAS storage, 17
storage
direct-attached, 15
direct-attached SAS
cluster, 18
SAN-attached cluster, 16

clusters

fibre channel setup, 12
hardware connections for, 13
removing nodes, 55
SAS setup, 16
verifying hardware and
software, 11

Clusterware, 28

adding
node, 52
configuring storage for, 19, 26
installation fails, 62
installing, 31
removing node, 57
upgrading, 35

CRS, 34

fails, 62
fails to start, 62-63
node continually
reboots, 63-64
Oracle Clusterware install
fails, 62

CSS miscount, 37

D

database

adding node
instance layer, 54
layer, 53
deleting node
instance layer, 55
layer, 57
error
LMDO, 60
error in KQLMBIVG, 60
instance terminated, 59

LMS

crash, 59
nodes
startup failure, 59

DBCA

no response, 64

Dell deployment CD

downloading images, 31

Dell|EMC Fibre Channel

hardware requirements, 7
storage, 7, 14
switch, 7, 14

deployment

CD, 31

disabling rsh, 42

E

EMC

Navisphere, 23
PowerPath, 23
failover, 37
software requirements, 6

Enterprise Manager

fails, 60

ethernet switch, 7, 14

F

Fibre Channel

cluster
hardware connections, 13
cluster setup, 12
direct-attached
configuration, 14
hardware
interconnections, 13
I/O errors when loading HBA
driver, 65
SAN-attached
configuration, 15

H

hardware

- Fibre Channel
 - cluster, 13
 - interconnections, 13
- requirements, 6

help, 65

- Dell support, 65
- Oracle support, 66

L

listener

- configuring, 38, 48
- reconfiguring, 53, 56

LUNs, 12

N

NETCA

- cannot configure remote nodes, 61
- fails, 61
- RAW device validation error, 61

networks

- adding nodes, 51
- cluster verification fails, 65
- configuring, 19
 - private, 19
 - bonding, 20
 - public, 19-20, 43
- determining
 - private, 58

NIC

- port assignments, 19

nodes

- adding
 - and removing, 51
 - to a network layer, 51
 - to Clusterware, 52
 - to database instance layer, 54
 - to database layer, 53
- configuring shared storage, 51
- deleting
 - database instance layer, 55
 - database layer, 57
- deleting ASM, 56
- removing
 - from cluster, 55
 - from Clusterware layer, 57

O

OCFS2, 19, 26

- creating seed database, 38

OCR, 28, 33

operating system

- does not recognize system, 65

Oracle 10g RAC

- deployment service, 5
- installing, 31

Oracle Database 10g

- configuring, 43
- deploying, 43
- installing, 46
 - patchset, 43, 47

software

- downloading
 - patchset, 43, 47
- installing, 34
- requirements, 6
- upgrading, 47

P

password

- setting, 43

patches

- downloading software, 35
- Oracle Database 10g, 35
- upgrading
 - Clusterware, 35
 - RAC, 36

PowerEdge

- hardware
 - interconnections, 13
 - requirements, 7-8
- hardware requirements, 7-8

R

RAW

- configuring ASM, 45
- device interface, 28
- device validation error, 61
- storage
 - configuring, 29

Red Hat

- configuring system packages, 11

Red Hat (*continued*)

- downloading
 - CD images, 31
 - ISO images, 8
- excessive use of swap space, 60
- instability, 60

Linux

- configuring, 10
- deployment CDs, 9
- installing, 9
 - and configuring, 8
- poor performance, 60
- software requirements, 6

rsh, disabling, 42

S

SAS

- cluster setup, 16

seed database, 48

- ASM, 40, 49
- ext3 file system, 48
- OCFS2, 38

software

- failure
 - to format OCR disk, 65
- installing
 - error messages, 64
 - failure, 64-65
 - Oracle Database 10g, 34
- requirements, 6
- supported versions, 57

source files

- obtaining and using, 66

storage

- configuring
 - ASM, 26, 28-30, 44-45, 51
 - Clusterware, 26
 - database, 26
 - ASM library driver, 43, 45
 - ext3 file system, 43
 - new node, 51
 - Oracle, 26
 - RAW, 28, 45

verifying

- configuration, 22

system

- clock
 - synchronizing, 31
- password, 43
- poor performance, 61
- securing, 42

T

troubleshooting, 59

V

Voting Disk, 28, 33

