

Dell™ PowerEdge™ Systems

Oracle Database 10g – Linux Deployment Guide Version 1.1

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.

© 2004 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, *Dell OpenManage*, *PowerEdge*, and *PowerVault* are trademarks of Dell Inc.; *EMC*, *PowerPath*, and *Navisphere* are registered trademarks of EMC Corporation; *Intel* is a registered trademark and *Xeon* is a trademark of Intel Corporation; *Red Hat* is a registered trademark 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.

Contents

| | |
|--|-----------|
| Software and Hardware Requirements | 7 |
| License Agreements | 9 |
| Important Documentation | 9 |
| Installing and Configuring Red Hat Enterprise Linux | 9 |
| 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 | 11 |
| SCSI Cluster Setup | 14 |
| Configuring Networking and Storage for Oracle RAC 10g | 16 |
| Configuring the Public and Private Networks | 16 |
| Verifying the Storage Configuration | 19 |
| Configuring Shared Storage for CRS | 21 |
| Configuring Shared Storage for the Database | 22 |
| Installing Oracle RAC 10g | 24 |
| Installing CRS | 24 |
| Installing the Oracle Database 10g Software | 25 |
| Configuring the Listener | 26 |
| Creating the Seed Database | 27 |
| Setting the oracle User Password | 30 |
| Configuring and Deploying Oracle Database 10g (Single Node) | 31 |
| Configuring the Public Network | 31 |
| Installing Oracle Database 10g | 31 |
| Configuring the Listener | 32 |
| Configuring Database Storage | 32 |
| Creating the Seed Database | 32 |
| Setting the oracle User Password | 33 |
| Adding and Removing Nodes | 34 |
| Adding a New Node to the Network Layer | 34 |
| Configuring the Shared Storage on the New Node | 34 |
| Adding a New Node to the Clusterware Layer | 36 |
| Adding a New Node to the Database Layer | 36 |

| | |
|--|-----------|
| Adding a New Node to the Database Instance Layer | 37 |
| Removing a Node From the Cluster | 39 |
| Reinstalling the Software | 40 |
| Additional Information | 41 |
| Supported Software Versions | 41 |
| Configuring Automatic Reboot for a Hung Operating System | 42 |
| Determining the Private Network Interface | 43 |
| Troubleshooting | 44 |
| Getting Help | 50 |
| Obtaining and Using Open Source Files | 51 |
| Index | 53 |

Figures

| | | |
|-------------|--|----|
| Figure 1-1. | Hardware Connections for a Fibre Channel Cluster | 12 |
| Figure 1-2. | Hardware Connections for a SCSI Cluster | 14 |

Tables

| | | |
|------------|---|----|
| Table 1-1. | Software Requirements | 7 |
| Table 1-2. | Minimum Hardware Requirements— Fibre Channel Cluster | 8 |
| Table 1-3. | Minimum Hardware Requirements— SCSI Cluster | 8 |
| Table 1-4. | Minimum Hardware Requirements— Single Node | 9 |
| Table 1-5. | Fibre Channel Hardware Interconnections | 12 |
| Table 1-6. | SCSI Hardware Interconnections | 14 |
| Table 1-7. | NIC Port Assignments | 16 |

| | | |
|-------------|---------------------------------------|----|
| Table 1-8. | Supported Software Versions | 41 |
| Table 1-9. | Integrated NICs | 43 |
| Table 1-10. | Troubleshooting | 44 |

This document provides information for installing, configuring, reinstalling, and using your Oracle 10g software following Dell's Supported Configurations for Oracle. The following topics are covered:

- Software and hardware requirements
- Installing and configuring Red Hat® Enterprise Linux
- Verifying cluster hardware and software configurations
- Configuring networking and storage for Oracle Real Application Clusters (RAC) 10g
- Installing Oracle RAC 10g
- Configuring and deploying 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's Supported Configurations for Oracle, see www.dell.com/oracle.

Software and Hardware Requirements




NOTE: You should have downloaded the Dell Deployment CD image from www.dell.com/oracle and created a CD. You must also have your Oracle and Red Hat original CDs.

Table 1-1 lists basic software requirements for Dell's Supported Configurations for Oracle. Table 1-2 through Table 1-4 list the hardware requirements. For detailed information on the minimum software versions for drivers and applications, see "Supported Software Versions" on page 41.

Table 1-1. Software Requirements

| Software Component | Configuration |
|--|--|
| Red Hat Enterprise Linux AS (version 3) operating system | Quarterly update 2 |
| Oracle 10g | Version 10.1.0 <ul style="list-style-type: none"> • Enterprise Edition, including the RAC option for clusters • Enterprise Edition for single-node configuration |
| EMC® PowerPath® (Fibre Channel clusters only) | Version 3.0.6 |

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


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

Table 1-2. Minimum Hardware Requirements—Fibre Channel Cluster

| Hardware Component | Configuration |
|--|---|
| Dell™ PowerEdge™ 1750, 1850, 2600, 2650, 2850, 4600, 6600, or 6650 system (two to eight nodes using Oracle Cluster File System [OCFS] or Automatic Storage Management [ASM]) | Intel® Xeon™ processor. 1 GB of RAM. PERC controller for internal hard drives. Two 36-GB hard drives (RAID 1) connected to a PERC controller. Three Gigabit NIC ports. Two QLogic optical HBA ports. |
| Dell EMC CX200, CX300, CX400, CX500, CX600, or CX700 Fibre Channel storage system | See www.dell.com/oracle for information on supported configurations. |
| Gigabit Ethernet switch (two) | See www.dell.com/oracle for information on supported configurations. |
| Dell EMC Fibre Channel switch (two) | Eight ports for two to six nodes. Sixteen ports for seven or eight nodes. |

Table 1-3. Minimum Hardware Requirements—SCSI Cluster

| Hardware Component | Configuration |
|---|---|
| Dell PowerEdge 1750, 1850, 2600, 2650, 2850, 4600, 6600, or 6650 system (two nodes) | Intel Xeon processor. 1 GB of RAM. PERC controller for internal hard drives. Two 36-GB hard drives (RAID 1) connected to a PERC controller. Three Gigabit NIC ports. Additional PERC controller for shared storage only. |
| Dell PowerVault™ 22.xS storage system | Two enclosure management modules (EMMs). At least one logical drive configured as RAID 0, RAID 1, RAID 5, or RAID 10. Dell recommends RAID 10. |
| Gigabit Ethernet switch | See www.dell.com/oracle for information on supported configurations. |

Table 1-4. Minimum Hardware Requirements—Single Node

| Hardware Component | Configuration |
|--|---|
| Dell PowerEdge 1750, 1850, 2600, 2650, 2850, 4600, 6600, or 6650 system | <p>Intel Xeon processor.</p> <p>1 GB of RAM.</p> <p>Two 36-GB hard drives (RAID 1) connected to a PERC controller.</p> <p>Two NIC ports.</p> <p>PERC controller for Oracle data files (optional).</p> |
| Dell PowerVault 22.xS storage system (optional) | <p>One EMM.</p> <p>At least one logical drive configured as RAID 0, RAID 1, RAID 5, or RAID 10. Dell recommends RAID 10.</p> |
| Dell EMC CX200, CX300, CX400, CX500, CX600, or CX700 Fibre Channel storage system (optional) | See www.dell.com/oracle for information on supported configurations. |
| Dell EMC Fibre Channel switch (optional) | Eight ports. |

License Agreements



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

Important Documentation

For more information on specific hardware components, see the documentation that came with your system.

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

Installing and Configuring Red Hat Enterprise Linux



NOTICE: You must disconnect all external storage 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 deployment.

Installing Red Hat Enterprise Linux Using the Deployment CDs

- 1 Disconnect all external storage from the system.
- 2 Locate your Dell *Deployment* CD and original Red Hat Enterprise Linux AS with update 2 CDs.

- 3 Insert the *Deployment CD* into the CD drive and reboot the system.
The system boots to the *Deployment CD*.
- 4 When prompted, type `1` and press `<Enter>` to select **Red Hat Enterprise Linux 3 AS for Oracle Database**.
If asked if you want to create a blank utility partition, type `yes`.
- 5 Insert each Red Hat *Installation CD* into the CD drive when prompted.
A deployment partition is created and the contents of the CDs are copied to it. When the copy operation is complete, the system automatically ejects the last CD and boots to the deployment partition.
When the installation is complete, the system automatically reboots and the Red Hat Setup Agent appears.
- 6 Click **Next** at the **Red Hat Setup Agent Welcome** window to configure your operating system settings.
 - When prompted, specify a `root` password.
 - When the **Network Setup** window appears, click **Next**. You will configure network settings later.
 - When the **Security Level** window appears, disable the firewall. You may enable the firewall after completing Oracle deployment.
- 7 Log in as `root`.

Configuring Red Hat Enterprise Linux

- 1 Log in as `root`.
- 2 Insert the Dell *Deployment CD* into the CD drive and type the following commands:


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

The contents of the CD are copied to the `/usr/lib/dell/dell-deploy-cd` directory. When the copy is complete, type `umount /dev/cdrom` and remove the CD from the CD drive.
- 3 Type `cd /dell-oracle-deployment/scripts/standard` to navigate to the directory containing the scripts installed from the Dell *Deployment CD*.
Scripts discover and validate installed component versions and, when required, update components to supported levels.
- 4 Type `./005-oraclesetup.py` to configure Red Hat Enterprise Linux for Oracle installation.
- 5 Type `source /root/.bash_profile`.

- 6 Type `./010-hwCheck.py` to verify that CPU, RAM, and disk size meet minimum requirements for Oracle installation.
If the script reports that a parameter failed, update your hardware configuration and run the script again.
- 7 Type `./275-rpms_dkms.py` to install the Dynamic Kernel Module Support (DKMS) driver.
- 8 Type `../custom/335-rpms_apps.py` to install PERC utility RPMs.
- 9 Type `./340-rpms_ocfs.py` to install the OCFS RPMs (optional for single node).

You should now connect the external storage.

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 www.dell.com/oracle for the latest supported configurations before you use RHN to update your system software to the latest revisions.



NOTE: If you are deploying Oracle on a single node, skip the following sections and see "Configuring and Deploying Oracle Database 10g (Single Node)" on page 31.

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. The following sections provide information on setup for both Fibre Channel and SCSI cluster configurations.

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 shows an overview of the connections required for the cluster, and Table 1-5 summarizes the cluster connections.

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

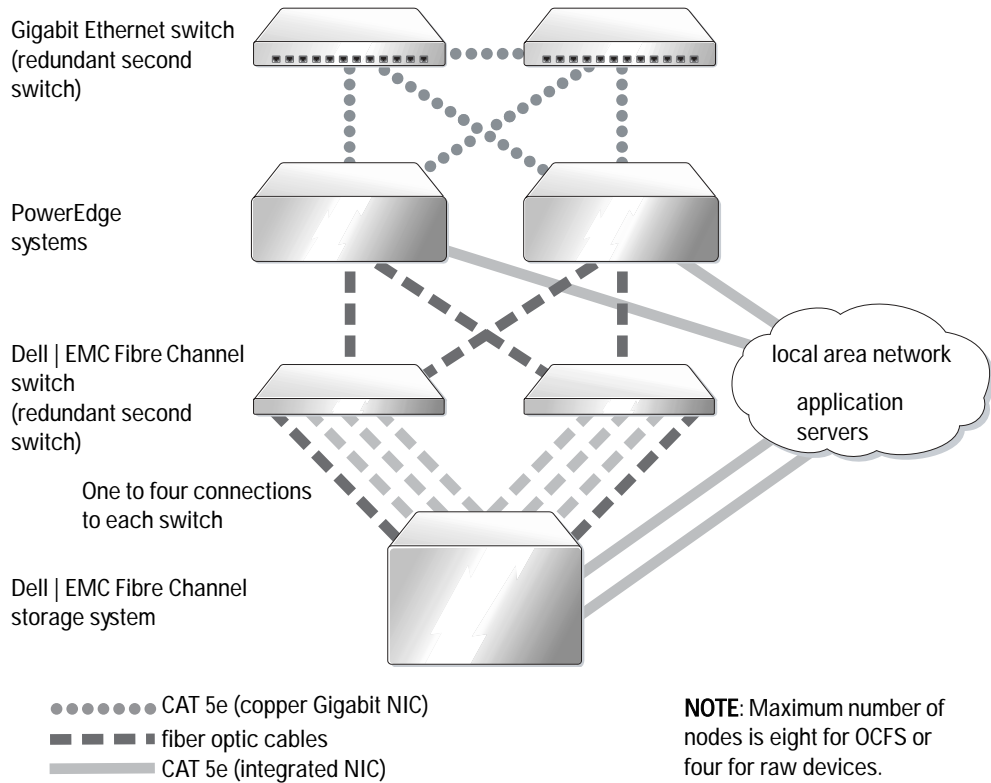


Table 1-5. Fibre Channel Hardware Interconnections

| Cluster Component | Connections |
|----------------------------|--|
| Each PowerEdge system node | <p>One CAT 5e from public NIC to LAN</p> <p>One CAT 5e from private Gigabit NIC to Gigabit Ethernet switch</p> <p>One CAT 5e from redundant private Gigabit NIC to redundant Gigabit Ethernet switch</p> <p>One optical cable from optical HBA 0 to Fibre Channel switch 0, and one optical cable from HBA 1 to switch 1</p> |

Table 1-5. Fibre Channel Hardware Interconnections (*continued*)

| Cluster Component | Connections |
|---|---|
| Dell EMC Fibre Channel storage system | Two CAT 5e cables connected to LAN One to four optical connections to each Fibre Channel switch. For example, for a four-port configuration: One optical cable from SPA port 0 to Fibre Channel switch 0 One optical cable from SPA port 1 to Fibre Channel switch 1 One optical cable from SPB port 0 to Fibre Channel switch 1 One optical cable from SPB port 1 to Fibre Channel switch 0 |
| Dell EMC Fibre Channel switch | One to four optical connections to the Dell EMC Fibre Channel storage system One optical connection to each PowerEdge system's HBA |
| Each Gigabit Ethernet switch | One CAT 5e connection to private Gigabit NIC on each PowerEdge system One CAT 5e connection to other Gigabit Ethernet switch |

Verify that the following tasks have been completed for your cluster:

- All hardware is installed in the rack.
- All hardware interconnections are set up as shown in Figure 1-1 and listed in Table 1-5.
- All logical unit numbers (LUNs), 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.

Fibre Channel Hardware and Software Configurations

- Each node must include the following minimum hardware peripheral components:
 - One or two hard drives (36-GB minimum) in the internal hard-drive bay
 - Three Gigabit NIC ports
 - Two QLogic HBAs
- Each node must have the following software installed:
 - Red Hat Enterprise Linux software (see Table 1-1)
 - QLogic driver
 - OCFS for cluster configuration using OCFS

- The Fibre Channel storage must be configured with the following:
 - A minimum of three LUNs created and assigned to the cluster
 - A minimum LUN size of 5 GB

SCSI Cluster Setup

Figure 1-2 shows an overview of connections required for your SCSI cluster. The following sections describe the connections and setup of your cluster hardware, and Table 1-6 summarizes the cluster connections.

Figure 1-2. Hardware Connections for a SCSI Cluster




Table 1-6. SCSI Hardware Interconnections

| Cluster Component | Connections |
|----------------------------|---|
| Each PowerEdge system node | One CAT 5e from public NIC to LAN One CAT 5e from private Gigabit NIC to Gigabit Ethernet switch One SCSI cable from RAID controller to PowerVault storage system |
| PowerVault storage system | SCSI cable to RAID controller on each PowerEdge system |
| Gigabit Ethernet switch | One CAT 5e connection to private Gigabit NIC on each PowerEdge system |

To complete the hardware connections as shown in Figure 1-2 and listed in Table 1-6:

- 1 Use a CAT 5e cable to connect the private Gigabit NIC of each PowerEdge system to the Gigabit Ethernet switch.
- 2 Ensure that the PowerVault SCSI disk storage enclosure split bus module is set to cluster mode.
- 3 Attach a SCSI cable from Channel 1 on the RAID controller of each PowerEdge system to an EMM on the PowerVault SCSI disk storage enclosure.
- 4 Ensure that both EMMs in the PowerVault SCSI disk storage enclosure are identical.

Setting Up the PowerEdge Nodes

To set up PowerEdge nodes, you must enable cluster mode on each system's RAID controller, set a different SCSI ID for each controller, and configure disk volumes.

Perform the following steps to set up both PowerEdge systems:

- 1 Boot the first PowerEdge system.
- 2 Press < Ctrl> < m> during the RAID controller's boot to start the RAID controller configuration.
- 3 Set the controller to **Cluster Mode**.

Each PERC card that is used to connect to a shared storage enclosure must have cluster mode enabled using the PERC card's BIOS configuration utility. Enabling cluster mode implements the additional functionality required for the controller to operate in a cluster environment. For more information about enabling cluster mode, see the PERC documentation, which includes information on the SCSI host adapter.

- 4 Change the SCSI ID for this controller to 6.
- 5 Reboot the system.
- 6 Press < Ctrl> < m> during the RAID controller's boot to start the RAID controller configuration.
- 7 Configure the volumes for the disk drives on the PowerVault SCSI disk storage enclosure. Create one logical drive of at least 36 GB and initialize.
- 8 Reboot the system.
- 9 Restart the second PowerEdge system.
- 10 Press < Ctrl> < m> during the RAID controller's boot to start the RAID controller configuration.
- 11 Set the controller to **Cluster Mode** (leave the SCSI ID for this controller as 7).
- 12 Reboot the system.
- 13 Press < Ctrl> < m> during the RAID controller's boot to start the RAID controller configuration.

- 14 Select **Objects** —> **Logical Drives** and verify that the controller can see the configured volumes.
- 15 Restart the system.

Both PowerEdge systems should now see the logical drives created on the SCSI disk enclosure.

SCSI Hardware and Software Configuration


- Each node must include the following minimum hardware peripheral components:
 - One or two hard drives (36-GB minimum) in the internal hard-drive bay
 - Integrated SCSI controller or integrated PERC 3/Di controller connected to the internal hard drives
 - Three NICs
 - One PERC 3/DC or PERC 4/DC controller connected to the external storage
- Each node must have Red Hat Enterprise Linux software (see Table 1-1) installed.
- The PowerVault SCSI enclosure must be configured with the following:
 - One logical drive created and initialized
 - A minimum of two 36-GB hard drives

Configuring Networking and Storage for Oracle RAC 10g

The following sections describe steps to set up a Fibre Channel or SCSI cluster running a seed database. Because of the complexity of the Oracle RAC 10g configuration, you must carefully complete a number of steps before you install Oracle and create a database. Follow the steps in their exact order and perform them as presented to have your system configured and running in the least amount of time.

Configuring the Public and Private Networks

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

 **NOTE:** For each node, you need an unused public and private 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 DNS.

Depending on the number of NIC ports available, configure the interfaces as shown in Table 1-7.

Table 1-7. 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) |

Table 1-7. NIC Port Assignments (*continued*)

| NIC Port | Three Ports Available | Four Ports available |
|----------|-----------------------|----------------------|
| 3 | Private IP (bonded) | Private IP (bonded) |
| 4 | NA | Virtual IP |

Configuring the Public Network

If you have not already done so, configure the public network 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, and 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` at the command prompt.
- 5 Type `ifconfig` at the command prompt to verify that the IP addresses are set correctly.
- 6 To check for proper network setup, 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 that `ssh` is working by typing `ssh <public IP>`.

Configuring the Private Network Using Bonding

Before you deploy the cluster, you must set up 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/modules.conf` file:

```
alias bond0 bonding
```

- 3 For high availability, edit the `/etc/modules.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.

```
options bonding miimon=100
```

- 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 member of the bond, perform the following steps:
 - a In the directory `/etc/sysconfig/network-scripts/`, edit the `ifcfg-eth n` file, containing the following lines:

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

- b Type `service network restart`.

Ignore any warnings observed.

- 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 to verify that the private network is functioning and that ssh is working by typing `ssh <private IP>`.

9 On *each node*, modify the `/etc/hosts` file by adding the following lines.



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

```
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>
```

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 As the user `oracle`, connect to each node to verify that rsh is working by typing `rsh <public hostname nodex>`, where `x` is the node number.

Verifying the Storage Configuration

During the cluster configuration described in this document, you will create partitions on your Fibre Channel storage or PowerVault SCSI enclosure. In order to create the partitions, all cluster nodes must be able to detect the external storage devices. To verify that each node can detect each storage LUN or logical disk, perform the following steps:

1 For Dell | EMC Fibre Channel storage, verify that the EMC Navisphere[®] agent and the correct version of PowerPath (see Table 1-8) 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 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 complete this step.


2 Visually verify that the storage and each node are connected correctly to the Fibre Channel switch (see Figure 1-1 and Table 1-5) or SCSI enclosure (see Figure 1-2 and Table 1-6).

3 Verify that you are logged in as `root`.

- 4 On each node, type `more /proc/partitions` at a command prompt.

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

For a Fibre Channel cluster, ensure that you see the same three PowerPath pseudo devices (`/dev/emcpowera`, `dev/emcpowerb`, and `/dev/emcpowerc`, for example). For a SCSI cluster, ensure that you see the logical drive (`/dev/sdb`, for example).

 **NOTE:** In the following sections, it is assumed that `sdb` is a logical drive on the external storage array. If this is not the case for your hardware configuration, substitute the appropriate device name when performing the procedures for SCSI cluster configuration.

The listed devices vary depending on how your storage is configured. The primary SCSI drive or RAID container (array) on each node will be listed as `sda` and be partitioned. If you have any other SCSI disks or RAID containers on the node, they will be listed as `sdb`, `sdc`, and so on. The LUNs on the Fibre Channel storage system or SCSI enclosure should also be seen as SCSI devices. For example, if you have one RAID container in the node and three logical disks on the storage, the node should see the node RAID container or internal disk as `sda` and the logical disks as `sdb`, `sdc`, and `sdd`. If you have three LUNs on the Fibre Channel storage, the node should see the node RAID container as `sda`, and the Fibre Channel LUNs as `emcpowera`, `emcpowerb`, and `emcpowerc`. Ensure that each node in the cluster sees the same number of Fibre Channel LUNs.

If you do not see the external storage devices, perform the following steps:

- 1 For a Fibre Channel storage system, stop the PowerPath service on all nodes by typing:

```
service naviagent stop
service PowerPath stop
```

- 2 For a Fibre Channel storage system, reload the HBA driver on all nodes to synchronize the kernel's partition tables on all nodes by typing:

```
rmmod qla2300
modprobe qla2300
```

- 3 For a Fibre Channel storage system, restart the PowerPath service on all nodes by typing:

```
service PowerPath start
service naviagent start
```

- 4 For a PowerVault SCSI enclosure, reboot both nodes.
- 5 Confirm that all nodes see the external storage devices by typing:

```
more /proc/partitions
```

Configuring Shared Storage for CRS

This section provides instructions for configuring the shared storage for the Cluster Ready Services (CRS).

- 1 On the first node, create three partitions on an external storage device with **fdisk**:
Type `fdisk /dev/emcpowerx` (`sdb` for a SCSI cluster) at the command prompt and create three partitions of 150 MB each. One will be used for the Cluster Repository, one as the voting disk, and the other for the Oracle system parameter file.
- 2 For a SCSI enclosure, create two additional partitions of at least 10 GB each.
One partition is for database files and the other for flash recovery files.
- 3 Verify the new partitions by typing `more /proc/partitions`.
If you do not observe the new partitions, type `sfdisk -R /dev/<device name>`.
- 4 On *each node*, perform the following steps:

- a Type the following commands to change the names of the raw character devices to make them identifiable:

```
mv /dev/raw/raw1 /dev/raw/votingdisk
mv /dev/raw/raw2 /dev/raw/ocr.dbf
mv /dev/raw/raw3 /dev/raw/spfile+ASM.ora
```

- b Type the following commands to set user oracle ownership of the Cluster Repository and voting disk:

```
chown oracle.dba /dev/raw/votingdisk
chown oracle.dba /dev/raw/ocr.dbf
chown oracle.dba /dev/raw/spfile+ASM.ora
```

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

For a Fibre Channel cluster:

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

For a SCSI cluster:

```
/dev/raw/votingdisk /dev/sdb1
/dev/raw/ocr.dbf /dev/sdb2
/dev/raw/spfile+ASM.ora /dev/sdb3
```

- d Type: `service rawdevices restart`.

Configuring Shared Storage for the Database

This section provides procedures for configuring the shared storage using either Oracle Cluster File System (OCFS) or Automatic Storage Management (ASM).

Configuring the Shared Storage Using OCFS

- 1 Log in as `root`.
- 2 On *each node*, perform the following steps:
 - a Start the X Window System, if it is not running, by typing `startx` at a command prompt.
 - b Type `ocfstool` at a command prompt.
 - c From the menu, click **Tasks** and click **Generate Config**.
 - d Enter the private NIC device name and private host name of the node and click **OK**.
 - e Click **Exit**.
- 3 For a Fibre Channel cluster, on the first node, create one partition on each of the other two external storage devices with **fdisk**:
 - a Type `fdisk /dev/emcpowerx` at the command prompt and create a primary partition for the entire device.
Type `h` for help within the `fdisk` utility.
 - b Verify the new partition by typing `cat /proc/partitions`.
If you do not observe the new partition, type `sfdisk -R /dev/<device name>`.



NOTE: The following steps use the sample values `/u01` and `/u02` for mount points and `u01` and `u02` as labels.

- 4 On *one node only*, format the external storage devices for OCFS by typing:


```
mkfs.ocfs -b 128 -F -u <ID of user oracle> -g <primary group ID of user oracle> -L <volume label> -m <mount point> -p <OCFS volume permissions> <PowerPath or SCSI device name>
```

To find the user ID and group ID, type `id oracle`.

- a For a Fibre Channel cluster, format the PowerPath pseudo device partitions by typing:


```
mkfs.ocfs -F -b 128 -L u01 -m /u01 -u 500 -g 500 -p 0775 /dev/emcpowerb1
mkfs.ocfs -F -b 128 -L u02 -m /u02 -u 500 -g 500 -p 0775 /dev/emcpowerc1
```

- b For a SCSI cluster, format the partitions to OCFS by typing the following:

```
mkfs.ocfs -F -b 128 -L u01 -m /u01 -u 500 -g 500 -p 0775
/dev/sdb5
mkfs.ocfs -F -b 128 -L u02 -m /u02 -u 500 -g 500 -p 0775
/dev/sdb6
```

- 5 On *each node*, perform the following steps:

- a Create mount points for each of the OCFS partitions by creating the directories for the locations where partitions will be mounted, and set the ownerships by typing:

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

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

For a Fibre Channel storage system:


```
/dev/emcpowerb1 /u01 ocfs _netdev 0 0
/dev/emcpowerc1 /u02 ocfs _netdev 0 0
```

For a SCSI enclosure:

```
LABEL=u01 /u01 ocfs _netdev 0 0
LABEL=u02 /u02 ocfs _netdev 0 0
```

Make appropriate entries for all OCFS volumes.

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

 **NOTE:** Ignore any warnings related to mismatch of OCFS module and kernel version.

```
/sbin/load_ocfs
mount -a -t ocfs
```

Configuring the Shared Storage Using ASM

This section describes how to set up the shared storage using ASM. To configure your cluster using ASM, perform the following steps on *all nodes*:

- 1 Log in as `root`.
- 2 Type the following commands to change the names of the raw character devices to make them identifiable:

```
mv /dev/raw/raw4 /dev/raw/ASM1
mv /dev/raw/raw5 /dev/raw/ASM2
```

- 3 Type the following commands to set user oracle ownership of the Cluster Repository and voting disk:

```
chown oracle.dba /dev/raw/ASM1
chown oracle.dba /dev/raw/ASM2
```

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

For a Fibre Channel cluster:

```
/dev/raw/ASM1          /dev/emcpowerb
/dev/raw/ASM2          /dev/emcpowerc
```

For a SCSI cluster:

```
/dev/raw/ASM1          /dev/sdb5
/dev/raw/ASM2          /dev/sdb6
```

- 5 Type: `service rawdevices restart`.

Installing Oracle RAC 10g

This section describes the steps required to install Oracle RAC 10g, 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.

Installing CRS

- 1 Log in as root.
- 2 Start the X Window System by typing `startx` at a command prompt, and type `xhost +` in a terminal window.
- 3 Mount the *Oracle Cluster Ready Services* CD.
- 4 Type `su - oracle`.
- 5 Type the following commands at a command prompt:

```
unset ORACLE_HOME
/mnt/cdrom/runInstaller.
```

The Oracle Universal Installer starts.

- 6 In the **Welcome** window, click **Next**.
- 7 In the **Specify File Locations** window, verify that the Oracle home path is `/opt/oracle/product/10.1.0/crs_1` and click **Next**.
- 8 In the **Language Selection** window, select a language and click **Next**.
- 9 In the **Cluster Configuration** window, enter a global cluster name or accept the default name **crs**, enter the public and private node names for each node, and click **Next**.

The cluster name must be unique throughout the enterprise.

- 10 In the **Private Interconnect Enforcement** window, click each interface type and select **public**, **private**, or **Do not use**, and then click **Next**.



NOTE: The NIC designations that you select in this step must be available on all nodes. If you select eth0, for example, as public, each of the other nodes must also have a public NIC with device name eth0.

- 11 In the **Oracle Cluster Registry** window, enter a complete path for the OCR disk location (**/dev/raw/ocr.dbf**) and click **Next**.
- 12 In the **Voting Disk** window, enter a complete path for the partition to use for storing the voting disk (**/dev/raw/votingdisk**) and click **Next**.
- 13 In the **Summary** window, click **Install**.

When the installation is complete, a message appears indicating that you must run the root.sh script on all the nodes. The root.sh script automatically configures the cluster.

- 14 From the X Window System, as user `root`, run the root.sh script on each node, beginning with the local node.

Wait for root.sh to finish running on each node before you run it on the next node.

- 15 Click **OK** in the **Setup Privileges** window.
- 16 Click **Exit** in the **End of Installation** window and confirm by clicking **Yes**.
- 17 On *all nodes*, verify the CRS installation by typing the following command from the **/opt/oracle/product/10.1.0/crs_1/bin** directory:

```
olsnodes -n -v
```

A list of the public node names of all nodes in the cluster appears.

Installing the Oracle Database 10g Software

- 1 As user `root`, mount the *Oracle Database 10g* CD.
- 2 As user `oracle`, type `/mnt/cdrom/runInstaller`.

The Oracle Universal Installer starts.

- 3 In the **Welcome** window, click **Next**.
- 4 In the **Specify File Locations** window, verify that the complete Oracle home path is **/opt/oracle/product/10.1.0/db_1** and click **Next**.




NOTE: The Oracle home in this step must be different from the Oracle home name that you identified during the CRS installation. You cannot install the Oracle10g Enterprise Edition with RAC into the same home that you used for CRS.

- 5 In the **Specify Hardware Cluster Installation Mode** window, click **Select All** and click **Next**.
- 6 In the **Select Installation Type** window, click **Enterprise Edition** and click **Next**.

A window appears displaying the status of various checks being performed. When the checks are complete, click **Next**.

- 7 In the **Select Database Configuration** window, click **Do not create a starter database** and click **Next**.
- 8 Click **Install** in the **Summary** window.
- 9 When prompted, run `root.sh` on node 1.
 - a Press <Enter> to accept the default value for the local **bin** directory.
The Virtual Internet Protocol Configuration Assistant (VIPCA) starts.
 - b On the first VIPCA page, click **Next**.
 - c In the **List of Available Network Interfaces** window, select your public NIC or, if you have four NIC ports, the port reserved for the virtual IP address (see "Configuring the Public and Private Networks" on page 16), and click **Next**.

 **NOTE:** The public NIC designation that you select in this step must be available on all nodes. The NICs listed in the window are those found on node 1. If you select `eth0`, for example, each of the other nodes must also have a public NIC with device name `eth0`.
 - d In the **Virtual IPs for Cluster Nodes** window, enter an unused public virtual IP address and subnet mask for each node displayed and click **Next**.
The virtual IP address must be the same as you entered in the `/etc/hosts.equiv` file, and the subnet mask must be the same as the public mask.
 - e Click **Finish** in the summary window.
A progress window appears.
 - f When the configuration is complete, click **OK** and click **Exit** to exit the VIPCA.
 - g Run `root.sh` on each of the other nodes in your cluster.
Wait for `root.sh` to finish running on each node before you run it on the next node.
- 10 Click **OK** in the **Setup Privileges** window.
- 11 Click **Exit** in the **End of Installation** window and confirm by clicking **Yes**.

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, if it is not already running, by performing the following steps:
 - a Type `startx` at a command prompt.
 - b Open a terminal window and type `xhost +` at the prompt.
- 3 As user `oracle`, type `netca` at the next command prompt.
The **Net Configuration Assistant** window appears.

- 4 Select **Cluster Configuration** and click **Next**.
- 5 On the **TOPSNodes** page, click **Select All Nodes** and click **Next**.
- 6 On the **Welcome** page, select **Listener Configuration** and click **Next**.
- 7 On the **Listener Configuration, Listener** page, select **Add** and click **Next**.
- 8 On the **Listener Configuration, Listener Name** page, type `LISTENER` in the **Listener Name** field and click **Next**.
- 9 On the **Listener Configuration, Select Protocols** page, select **TCP** and click **Next**.
- 10 On the **Listener Configuration, TCP/IP Protocol** page, select **Use the standard port number of 1521** and click **Next**.
- 11 On the **Listener Configuration, More Listeners?** page, select **No** and click **Next**.
- 12 On the **Listener Configuration Done** page, click **Next**.
- 13 Click **Finish**.

Creating the Seed Database

This sections contains procedures for creating the seed database using either OCFS or ASM and for verifying the seed database.

Creating the Seed Database Using OCFS

- 1 On *node 1*, as user `oracle`, type `dbca -datafileDestination /u01` to start the Database Configuration Assistant (DBCA).
- 2 In the **Welcome** window, select **Oracle Real Application Cluster Database** and click **Next**.
- 3 In the **Operations** window, click **Create a Database** and click **Next**.
- 4 In the **Node Selection** window, click **Select All** and click **Next**.
- 5 In the **Database Templates** window, click **Custom Database** and click **Next**.
- 6 In the **Database Identification** window, enter a **Global Database Name** such as `racdb` and click **Next**.
- 7 In the **Management Options** window, click **Next**.
- 8 In the **Database Credentials** window, click **Use the Same Password for All Accounts**, complete password selections and entries, and click **Next**.
- 9 In the **Storage Options** window, select **Cluster File System** and click **Next**.
- 10 In the **Database File Locations** window, click **Next**.
- 11 In the **Recovery Configuration** window, click **Specify flash recovery area**, click **Browse** and select `/u02`, specify the flash recovery size, and then click **Next**.
- 12 In the **Database Content** window, click **Next**.
- 13 In the **Database Services** window, click **Next**.

- 14 In the **Initialization Parameters** window, if your cluster has more than four nodes, change the **Shared Pool** value to 500 MB, and click **Next**.
- 15 In the **Database Storage** window, click **Next**.
- 16 In the **Creation Options** window, check **Create Database** and click **Finish**.
- 17 In the **Confirmation** window click **OK** to create the database.
The creation of the seed database may take more than an hour.
When the database creation is complete, the **Password Management** window appears.
- 18 Click **Exit**.
A message appears indicating that the cluster database is being started on all nodes.
- 19 On *each node*, first type `srvctl status database -d <database name>` at a command prompt to determine which database instance exists on that node; then 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.

Creating the Seed Database Using ASM

Perform the following steps to create the seed database using Oracle ASM:

- 1 On *node 1*, as user `oracle`, type `dbca &` to start the DBCA.
- 2 In the **Welcome** window, select **Oracle Real Application Cluster Database** and click **Next**.
- 3 In the **Operations** window, click **Create a Database** and click **Next**.
- 4 In the **Node Selection** window, click **Select All** and click **Next**.
- 5 In the **Database Templates** window, click **Custom Database** and click **Next**.
- 6 In the **Database Identification** window, enter a **Global Database Name** such as `racdb` and click **Next**.
- 7 In the **Management Options** window, click **Next**.
- 8 In the **Database Credentials** window, click **Use the Same Password for All Accounts**, complete password selections and entries, and click **Next**.
- 9 In the **Storage Options** window, click **ASM** and click **Next**.
- 10 In the **ASM Credentials** window, enter the password for user `SYS`, click **Create server parameter file**, change the location to `/dev/raw/spfile+ASM.ora`, and then click **Next**.
- 11 When a message appears indicating that DBCA is ready to create and start the ASM instance, click **OK**.

- 12 Under **Available Disk Groups**, click **Create New**.
- 13 Enter the information in the **Disk Group** window for the database files and click **OK**.
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 (for example, /dev/raw/ASM1).
A window appears indicating that disk group creation is in progress.
- 14 Under **Available Disk Groups**, click **Create New**.
- 15 Enter the information in the **Disk Group** window for the flashback recovery files and click **OK**.
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 (for example, /dev/raw/ASM2).
A window appears indicating that disk group creation is in progress.
- 16 In the **Select disk groups to be used as storage for database** window, check the disk group that you would like to use for database storage (for example, databaseDG) and click **Next**.
- 17 In **Select File Locations** window, check **Use Common Location for All Database Files**, and click **Next**.
- 18 In **Recovery Configuration** window, click **Browse**, select the flashback group that you created in step 15 (for example, flashbackDG), and click **Next**.
- 19 In **Database Content** window, click **Next**.
- 20 In **Database Services** window, click **Next**.
- 21 In the **Initialization Parameters** window, if your cluster has eight nodes, change the **Shared Pool** value to 500 MB, and click **Next**.
- 22 In the **Database Storage** window, click **Next**.
- 23 In the **Creation Options** window, select **Create Database** and click **Finish**.
- 24 In the **Confirmation** window click **OK** to create the database.
The creation of the seed database may take more than an hour.
When the database creation is complete, the **Password Management** window appears.
- 25 Click **Exit**.
A message appears indicating that the cluster database is being started on all nodes.
- 26 On *each node*, first type `srvctl status database -d <database name>` at a command prompt to determine which database instance exists on that node; then 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.

- 27 On *one node*, at the command prompt, 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.

Verifying the Private Interface

After you run DBCA, your cluster may use the public interface instead of the private interface. If this occurs, an unknown interface type warning appears in the Oracle alert log and performance may be degraded. To force cluster communications to the private interface, perform the following steps on *one node*:

- 1 Log in as `oracle`.
- 2 Type `sqlplus "/ as sysdba"` at the command prompt.

The `SQL>` prompt appears.

- 3 Enter the following lines at the `SQL>` prompt:

```
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>'
```

Continue entering lines for each node in the cluster.

- 4 Restart the database on all nodes by typing the following lines:

```
srvctl stop database -d <dbname>
srvctl start database -d <dbname>
```

- 5 Open the `/opt/oracle/admin/<dbname>/bdump/alert_<SID>.log` file and verify that the private IP addresses are being used for all instances.

Setting the oracle User Password

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` at the command prompt, and follow the onscreen instructions to create the oracle password.

Configuring and Deploying Oracle Database 10g (Single Node)

This section provides information on configuring your system initially or after reinstalling your Linux software following the procedures described in "Installing and Configuring Red Hat Enterprise Linux" on page 9.

Configuring the Public Network

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

Installing Oracle Database 10g

Perform the following steps to install Oracle 10g, version 10.1.0:

- 1 Log in as `root`.
- 2 For best installation performance, copy the *Oracle Database 10g* installation CD to a system hard drive.
Create a directory `/oracle_cd` and copy the CD to the directory.
- 3 Start the X Window System by typing `startx` at a command prompt.
- 4 Open a terminal window and type `xhost +` at the prompt.
- 5 Log in as `oracle`.
- 6 Type `/oracle_cd/runInstaller` at a command prompt.
The Oracle Universal Installer starts.
- 7 In the **Welcome** window, click **Next**.
- 8 In the **Specify File Locations** window, verify that the complete Oracle home path is `/opt/oracle/product/10.1.0/db_1` and click **Next**.
- 9 In the **Select a Product to Install** window, click **Oracle Database 10g 10.1.0.1.0** and click **Next**.
- 10 In the **Select Installation Type** window, click **Enterprise Edition** and click **Next**.
- 11 In the **Select Database Configuration** window, click **Do not create a starter database** and click **Next**.
- 12 Click **Install** in the **Summary** window.
- 13 When prompted, run `root.sh`.
A brief progress window appears, followed by the **End of Installation** window.
- 14 Click **Exit** and confirm by clicking **Yes**.

Configuring the Listener

- 1 Log in as `root`.
- 2 Start the X Window System, if it is not running, by typing `startx` at a command prompt.
- 3 Open a terminal window and type `xhost +` at the prompt.
- 4 Log in as `oracle`.
- 5 Type `netca`.
The Oracle Net Configuration Assistant opens.
- 6 Accept the default settings and click **Next** on all screens to complete the listener configuration.

Configuring Database Storage

If you have additional storage, perform the following steps:

- 1 Log in as `root`.
- 2 Type `cd /opt/oracle`.
- 3 Type `mkdir oradata`.
- 4 Using **fdisk**, create a partition where you want to store your database files (for example, **sdb1** if your storage device is **sdb**).
- 5 Verify the new partition by typing `cat /proc/partitions`.
If you do not observe the new partition, type `sfdisk -R /dev/sdb`.
- 6 Type `mke2fs -j /dev/sdb1`.
- 7 Modify the **/etc/fstab** file by adding an entry for the newly created file system.
- 8 Type `mount /dev/sdb1 /opt/oracle/oradata`.
- 9 Type `chown oracle.dba oradata`.

Creating the Seed Database

Perform the following steps to create a seed database with the Oracle Database Creation Assistant (DBCA):

- 1 Log in as `oracle`.
- 2 At a command prompt, type `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 **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.
The creation of the seed database may take more than an hour.
When the database creation is complete, the **Password Management** window appears.
- 17 Click **Exit**.
- 18 At the command prompt, 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"` at the command prompt.
The `SQL>` prompt appears.
 - b Enter 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.

Setting the oracle User Password

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` at the command prompt, and follow the onscreen instructions to create the oracle password.

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. Adding a node involves adding the node to the network layer and configuring the shared storage, and then adding the node to the clusterware, database, and database instance layers. To remove a node, you reverse the process and remove it from the database instance, the database, and finally the clusterware layers.

For further information about adding an additional node to an existing cluster, see the document titled *Oracle Real Application Clusters 10g Administration*.

Adding a New Node to the Network Layer

Prepare the new node by first performing the steps in "Installing and Configuring Red Hat Enterprise Linux" on page 9. Then, in "Configuring Networking and Storage for Oracle RAC 10g" on page 16, perform the steps in "Configuring the Public and Private Networks" on page 16 and "Verifying the Storage Configuration" on page 19.

Configuring the Shared Storage on the New Node

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. This section presents procedures for either ASM or OCFS.

Configuring Shared Storage Using ASM

If you are using ASM, ensure that the new nodes can access the ASM disks with the same permissions as the existing nodes.

Configure ASM disks by following this procedure:

- 1 Log in as `root`.
- 2 If your existing nodes were configured using raw devices for CRS, perform the following steps:

- a Type the following commands to change the names of the raw character devices to make them identifiable:

```
mv /dev/raw/raw1 /dev/raw/votingdisk
mv /dev/raw/raw2 /dev/raw/ocr.dbf
mv /dev/raw/raw3 /dev/raw/spfile+ASM.ora
```

- b Type the following commands to set user oracle ownership of the Cluster Repository and voting disk:

```
chown oracle.dba /dev/raw/votingdisk
chown oracle.dba /dev/raw/ocr.dbf
chown oracle.dba /dev/raw/spfile+ASM.ora
```

- 3 Type the following commands to change the names of the raw character devices to make them identifiable as configured on the existing nodes:

```
mv /dev/raw/raw4 /dev/raw/ASM1
mv /dev/raw/raw5 /dev/raw/ASM2
```

- 4 Type the following commands to set user oracle ownership of the Cluster Repository and voting disk:

```
chown oracle.dba /dev/raw/ASM1
chown oracle.dba /dev/raw/ASM2
```

- 5 Copy the **/etc/sysconfig/rawdevices** file from one of the existing nodes to the same location on the new node.
- 6 Type `chown oracle.dba /dev/raw/ASM*`.
- 7 Type `service rawdevices restart`.

Configuring Shared Storage Using OCFS

If you are using Oracle Cluster File System for either CRS, quorum, or database files, ensure that the new nodes can access the cluster file systems in the same way that the other nodes access them.

- 1 Edit the **/etc/fstab** file on the new node and add OCFS volume information exactly as it appears on the existing nodes:

For example:

```
/dev/emcpowera1 /u01 ocfs _netdev 0 0
/dev/emcpowerb1 /u02 ocfs _netdev 0 0
/dev/emcpowerc1 /u03 ocfs _netdev 0 0
```

- 2 Create OCFS mount points on the new node as they exist on the existing nodes (for example, `/u01`, `/u02`, and `/u03`).
- 3 Run `ocfstool` to generate the OCFS configuration file **/etc/ocfs.conf** by performing the following steps:
 - a Start the X Window System, if it is not running, by typing `startx` at a command prompt.
 - b Type `ocfstool` at a command prompt.
 - c From the menu, click **Tools** and click **Generate Config**.
 - d Enter the private IP address and private host name of the node and click **OK**.
 - e Click **Exit**.
- 4 Type the following commands to load the OCFS module and mount all volumes listed in the **/etc/fstab** file:

```
/sbin/load_ocfs
mount -a -t ocfs
```

Adding a New Node to the Clusterware Layer

- 1 Log in as `oracle`.
- 2 From the `/opt/oracle/product/10.1.0/crs_1/oui/bin` directory on one of the existing nodes, type `addNode.sh`.
The Oracle Universal Installer starts.
- 3 In the **Welcome** window, click **Next**.
- 4 In the **Specify Cluster Nodes for Node Addition** 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 **Node Addition Summary** window appears.
- 5 Click **Next**.
The **Cluster Node Addition Progress** window displays the status of the cluster node addition process.
- 6 When prompted, run `rootaddnode.sh` on the local node.
When `rootaddnode.sh` finishes running, click **OK**.
- 7 When prompted, run `root.sh` on the new node.
When `root.sh` finishes running, click **OK**.
- 8 In the **End of Cluster Node Addition** window, click **Exit**.
- 9 From the `/opt/oracle/product/10.1.0/crs_1/oui/bin` directory on one of the existing nodes, type (for example) the following line:

```
racgons add_config node1-pub:4948 node2-pub:4948 node3-pub:4948
```


In this example, `node3` is being added to an existing two-node cluster.

Adding a New Node to the Database Layer

- 1 Log in as `oracle`.
- 2 From the `/opt/oracle/product/10.1.0/db_1/oui/bin` directory on one of the existing nodes, type `addNode.sh`.
The Oracle Universal Installer starts.
- 3 In the **Welcome** window, click **Next**.
- 4 In the **Specify Cluster Nodes for Node Addition** window, click the new node and click **Next**.
If all the verification checks pass, the **Node Addition Summary** window appears.

- 5 Click **Next**.

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

- 6 When prompted, run `root.sh` on the new node.

When `root.sh` finishes running, click **OK**.

- 7 In the **End of Cluster Node Addition** window, click **Exit**.

- 8 From the `/opt/oracle/product/10.1.0/db_1/oui/bin` directory on one of the existing nodes, type (for example) the following line:


```
vipca -nodelist node1-pub,node2-pub,node3-pub
```

In this example, `node3` is being added to an existing two-node cluster.

The Virtual Internet Protocol Configuration Assistant (VIPCA) starts.

- a On the first VIPCA page, click **Next**.

- b In the **List of Available Network Interfaces** window, select your public NIC and click **Next**.

 **NOTE:** The public NIC designation that you select in this step must be available on all nodes. The NICs listed in the window are those found on node 1. If you select `eth0`, for example, each of the other nodes must also have a public NIC with device name `eth0`.

- c In the **IP Address** window, enter an unused public virtual IP address and subnet mask for the new node and click **Next**.

- d Click **Finish** in the summary window.

A progress window appears.

- e When the configuration is complete, click **OK** and click **Exit** to exit the VIPCA.

Adding a New Node to the Database Instance Layer

- 1 On *node 1*, as user `oracle`, type `dbca &` to start the DBCA.

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

- 3 In the **Oracle Real Application Cluster Database** window, click **Next**.

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

- 5 In the **Instance Management** window, click **Add Instance** and click **Next**.

- 6 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.

- 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 Click **Next**.
- 9 In the **Adding an Instance** window, enter the instance name at the top of the page, select the new node name, and click **Next**.
- 10 In the **Services** window, click **Next**.
- 11 In the **Instance Storage** window, click **Finish**.
- 12 In the **Summary** window click **OK** to add the database instance.
A progress bar appears, followed by a message asking if you want to perform another operation.
- 13 Click **No** to exit DBCA.
- 14 On any node, type `srvctl status database -d <database name>` at a command prompt to determine if the instance has been successfully added.

Verifying the Private Interface

After you run DBCA, your cluster may use the public interface instead of the private interface. If this occurs, an unknown interface type warning appears in the Oracle alert log and performance may be degraded. To force cluster communications to the private interface, perform the following steps on *one node*:

- 1 Log in as `oracle`.
- 2 Type `sqlplus "/ as sysdba"` at the command prompt.
The `SQL>` prompt appears.
- 3 Enter the following lines at the `SQL>` prompt:


```
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>'
```

Continue entering lines for each node in the cluster.
- 4 Restart the database on all nodes by typing the following lines:


```
srvctl stop database -d <dbname>
srvctl start database -d <dbname>
```
- 5 Open the `/opt/oracle/admin/<dbname>/bdump/alert_<SID>.log` file and verify that the private IP addresses are being used for all instances.

Removing a Node From the Cluster

Deleting the Node From the Database Instance Layer

- 1 Log in as `oracle`.
- 2 From node 1, 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 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 **Finish**.

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 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 is complete, a dialog asks whether you want to perform another operation.

- 10 Click **No** to exit.
- 11 Type `srvctl config database -d <database name>` to verify that the node was removed.

Deleting the 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 (node3-pub for example):

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

- 3 On the node being deleted, log in as `root`.

- 4 Type the following command, using the public name of the node you are deleting (node3-pub for example):

```
/opt/oracle/product/10.1.0/db_1/install/rootdeletenode.sh
node3-pub
```

The CRS node applications are deleted. Ignore any warnings observed.

- 5 If you wish to remove the Oracle database software, type the following command:

```
rm -rf /opt/oracle/product/10.1.0/db_1/*.
```

Removing the Node from the Clusterware Layer

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

```
/opt/oracle/product/10.1.0/crs_1/install/rootdelete.sh remote
nosharedvar
```

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

```
/opt/oracle/product/10.1.0/crs_1/install/rootdeletenode.sh
<public nodename>, <node-number>
```

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

```
/opt/oracle/product/10.1.0/crs_1/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 /opt/oracle/product/10.1.0/crs_1/*
```

- 4 To verify that the node is no longer part of the cluster, type `olsnodes -n`.

Reinstalling the Software

- ➡ **NOTICE:** Reinstalling the software erases all information on the hard drives.
- ➡ **NOTICE:** You must disconnect all external storage from the system *before* you reinstall the software.
- ➡ **NOTICE:** Dell recommends that you perform regular backups of your database and individual nodes so that you do not lose valuable data. Reinstall the node software only as a last resort.

Installing the software using the *Deployment CD* created a redeployment partition on your hard drive that contains all of the software images that were installed on your system. The redeployment partition allows for quick redeployment of the Oracle software.

Reinstalling the software by using this method requires that you boot the system to the redeployment partition. When the system boots to this partition, it automatically reinstalls the Red Hat Linux operating system.

To reinstall software using this method, perform the following steps:

- 1 Disconnect the external storage.
- 2 Log in as `root` on the system on which you want to reinstall the software.
- 3 Type `fdisk /dev/sda` at the command prompt.
- 4 When prompted for a command within the `fdisk` program, type `a`.
- 5 When prompted for a partition number, type `2`.
- 6 When prompted for a command, again type `a`.
- 7 When prompted for another partition number, type `3`.
- 8 When prompted for a command, type `d`.
- 9 When prompted for a partition number, type `4`.
- 10 When prompted for a command, again type `d`.
- 11 When prompted for a partition number, type `3`.
- 12 When prompted for a command, type `p` to list the partitions.
Only the `/dev/sda2` partition is present; it is preceded by an asterisk.
- 13 When again prompted for a command, type `w`.
- 14 When the Linux command prompt appears, reboot your system.

Reinstallation completes automatically and the Linux operating system reboots and presents a login prompt. For information about configuring the system for use, see "Configuring Red Hat Enterprise Linux" on page 10 and continue through the remaining sections to reconfigure your system.

Additional Information

Supported Software Versions



NOTE: For this release of Dell Supported Configurations for Oracle, Emulex HBAs and PCI Express controllers are not supported.

Table 1-8 lists the supported software at the time of release. For the latest supported hardware and software, see www.dell.com/oracle and download the version 1.1 Solution Deliverable List.

Table 1-8. Supported Software Versions

| Software Component | Supported Versions |
|---|--|
| Red Hat Enterprise Linux AS (version 3) Quarterly Update 2 operating system | 2.4.21-15.EL, 2.4.21-15.ELsmp, and 2.4.21-15.ELhugemem |
| Oracle Patchset | None |

Table 1-8. Supported Software Versions (*continued*)

| Software Component | Supported Versions |
|---|--------------------|
| OCFS | ocfs-1.0.11 EL 3.0 |
| PowerPath for Linux | 3.0.6 |
| DKMS | 1.0.2-1 |
| Qlogic HBA 2340 driver | 6.07.02-RH2 |
| Qlogic HBA 2342 driver | 6.07.02-RH2 |
| PERC 3/DC Driver (megaraid2) | 2.10.1.1 |
| PERC 4/DC, PERC 4/Di, or PERC 4e/Di Driver (megaraid2) | 2.10.1.1 |
| Broadcom integrated NIC drivers (tg3) | 3.1 |
| Broadcom NIC drivers(tg3) | 3.1 |
| Intel PRO/100 S NIC drivers (e100) | 5.2.30.1 |
| Intel PRO/1000 XT/MT/MT DP NIC drivers (e1000) | 5.2.30.1 |

Configuring Automatic Reboot for a Hung Operating System

Install managed system software for Red Hat Enterprise Linux by performing the following steps:

- 1 Log on with administrator privileges to the system where you want to install the managed system components.
- 2 Exit any open application programs and disable any virus-scanning software.
- 3 Start the X Window System, if it is not running, by typing `startx` at a command prompt.
- 4 Open a terminal window and type `xhost +` at the prompt.
- 5 Insert the *Dell OpenManage Systems Management* CD into the CD drive on the system.
- 6 Type `mount /mnt/cdrom` to mount the CD.
- 7 Click **start.sh** located in the root directory of the CD to start the setup program.
- 8 Click **Next** on the **Welcome to Dell OpenManage Systems Management Installation** screen.
- 9 Read and accept the software license agreement to continue.

The setup program provides both an **Express Setup** option and a **Custom Setup** option. The **Express Setup** option (recommended) automatically installs all of the software components necessary to manage your system. The **Custom Setup** option allows you to select which software components you want to install.

The rest of this procedure is based on the **Express Setup** option. See the *Server Administrator User's Guide* for information about the **Custom Setup** option.

- 10 Click **Express Setup**.
- 11 Read the information on the **Installation Summary** screen, and then click **Next**.
The setup program automatically installs all of the managed system software for your hardware configuration.
- 12 When the installation is complete, click **Finish**.

See the *Server Administrator User's Guide* for information about the uninstallation of the managed system software.

To configure the automatic reboot option, perform the following steps:

- 1 At a command prompt, type:

```
omconfig system recovery action=reboot
```

This command sets the timer to a default setting of 480 seconds, which is the delay before automatic reboot of an unresponsive system.

- 2 To change the timer setting to a different value, type:

```
omconfig system recovery timer=<seconds>
```

- 3 To verify the system reboot timer settings, type:

```
omreport system recovery
```

Determining the Private Network Interface

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

- 1 Determine which types of NICs are in your system.

Refer to table Table 1-9 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 determine which you have.

Table 1-9. Integrated NICs

| System | Integrated NICs |
|----------------|--------------------------------|
| PowerEdge 1750 | Broadcom NetXtreme Gigabit (2) |
| PowerEdge 1850 | Intel PRO/1000 (2) |
| PowerEdge 2600 | Intel PRO/1000 |
| PowerEdge 2650 | Broadcom NetXtreme Gigabit (2) |
| PowerEdge 2850 | Intel PRO/1000 (2) |
| PowerEdge 4600 | Broadcom NetXtreme Gigabit (2) |
| PowerEdge 6600 | Broadcom NetXtreme Gigabit (2) |
| PowerEdge 6650 | Broadcom NetXtreme Gigabit (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. This is your private NIC.
- Determine which driver module your private NIC uses.
The Broadcom NetXtreme Gigabit uses **tg3**, and the Intel PRO/1000 family uses **e1000**.
- View the **/etc/modules.conf** file by typing `more /etc/modules.conf` at a command prompt.
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.
- Observe which Ethernet interfaces (eth*X*) have been assigned to the type of Gigabit NIC that is connected to the Gigabit switch.
If there is only one entry in **modules.conf** for your driver module type, then you have successfully identified the private network interface.
- 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" for the correct driver module until you have identified the correct Ethernet interface.

Troubleshooting

Table 1-10 provides recommended actions for problems that you may encounter while deploying and using your Red Hat Enterprise Linux and Oracle software.

Table 1-10. Troubleshooting

| Category | Problem / Symptom | Cause | Recommended Corrective Action |
|---------------------------|---|--|---|
| 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. | <p>Ensure that SGA size does not exceed 65% of total system RAM.</p> <p>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.</p> |
| Performance and stability | Unknown interface type warning appears in Oracle alert file. Poor performance may be observed. | The public interface is being used for the cluster communications. | <p>Force cluster communications to the private interface by performing the following steps on <i>one node</i>:</p> <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> <p>Continue entering lines for each node in the cluster.</p> 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. |

Table 1-10. Troubleshooting (*continued*)

| Category | Problem / Symptom | Cause | Recommended Corrective Action |
|----------|---|---|---|
| OCFS | System hangs or displays a kernel panic message. | Red Hat Linux system by default has a daily cron job routine set up to perform updatedb on certain file systems in /etc/cron.daily and /etc/updatedb.conf . OCFS does not support updatedb on its file systems. | <p>1 Edit the /etc/cron.d/slocate.conf file and add OCFS to the excluded file system list. The file should resemble the following:</p> <pre>#!/bin/sh renice +19 -p \$\$ >/dev/null 2>&1 /usr/bin/updatedb -f "ocfs nfs,smbfs,ncpfs,proc,devpts" -e "/tmp,/var/tmp,/usr/tmp,/afs,/net"</pre> <p>2 Edit the /etc/updatedb.conf file and add OCFS to the PRUNEFS list. The file should resemble the following:</p> <pre>PRUNEFS="ocfs devpts NFS nfs afs proc smbfs autofs auto iso9660" PRUNEPATHS="/tmp /usr/tmp /var/tmp /afs /net" export PRUNEFS export PRUNEPATHS</pre> |
| OCFS | OCFS does not recognize newly replaced NIC. | If you change the NIC whose name is listed in /etc/ocfs.conf , you must run <code>ocfs_uid_gen -r</code> to update the MAC address in the ocfs.conf file before you can mount OCFS volumes. | If you replace the NIC whose IP address is listed in /etc/ocfs.conf , type <code>ocfs_uid_gen -r</code> before you load the OCFS driver or mount OCFS partitions. |
| OCFS | Large file copy, move, dd and other operations on OCFS files cause system to hang occasionally. | The default fileutils package does not support <code>o_direct</code> file operations, which are required by OCFS files. | Download the latest OCFS supported fileutils package from Oracle Technology Network. |
| NETCA | NETCA fails, resulting in database creation errors. | The public network, hostname, or virtual IP is not listed in the /etc/hosts.equiv file. | Before launching netca, make sure that a hostname is assigned and that the public IP address and the virtual IP address are listed in the /etc/hosts.equiv file. |

Table 1-10. Troubleshooting (*continued*)

| Category | Problem / Symptom | Cause | Recommended Corrective Action |
|----------|--|---|---|
| NETCA | NETCA cannot configure remote nodes or you get a raw device validation error while running DBCA. | The /etc/hosts.equiv file either does not exist or does not include the assigned public or virtual IP addresses. | Verify that the /etc/hosts.equiv file on each node contains the correct public and virtual IP address. Try to rsh to other public names and VIP addresses as <code>oracle user</code> . |
| CRS | CRS fails to start when you reboot the nodes or type <code>/etc/init.d/init.crs start</code> . | 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 /opt/oracle/product/10.1.0/crs_1/. • Verify that each node has access to the quorum disk and the disk is writeable by the root user. • Check the last line in the file \$ORA_CRS_HOME/css/log/ocssd.log • If you see <code>clssnmvWriteBlocks: Failed to flush writes to (votingdisk)</code>, verify that the /etc/hosts file on each node contains correct IP addresses for all node hostnames, including the virtual IP addresses. Verify that you can ping the public and private hostnames. Also verify that quorum disk is writeable. |
| 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 /opt/oracle/product/10.1.0/crs_1/ 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 I/O problem and attempt to start the service again by rebooting the node or by running <code>root.sh</code> from /opt/oracle/product/10.1.0/crs_1/ . |

Table 1-10. Troubleshooting (*continued*)

| Category | Problem / Symptom | Cause | Recommended Corrective Action |
|----------|--|---|---|
| CRS | When you run root.sh 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 bs=8192 count=12800 dd if=/dev/zero of= /dev/raw/votingdisk bs= 8192 count=2560</pre> 2 Attempt to start the service again by rebooting the node or by running root.sh from /opt/oracle/product/10.1.0/crs_1/. |
| CRS | When you run root.sh, CRS fails to start. | The oracle user does not have permissions on /var/tmp (specifically /var/tmp/.oracle). | <ol style="list-style-type: none"> 1 Make oracle user the owner of /var/tmp/.oracle by typing <code>chown oracle.oinstall /var/tmp/.oracle</code> 2 Attempt to start the service again by rebooting the node or by running root.sh from /opt/oracle/product/10.1.0/crs_1/. |
| CRS | When you run root.sh, CRS fails to start. | Other CRS troubleshooting steps have been attempted without success. | <ol style="list-style-type: none"> 1 Enable debugging by adding the following line to root.sh: <code>set -x</code> 2 Attempt to start the service again by running root.sh from /opt/oracle/product/10.1.0/crs_1/. 3 Check log files in the following directories to diagnose the issue: \$ORA_CRS_HOME/crs/log \$ORA_CRS_HOME/crs/init \$ORA_CRS_HOME/css/log \$ORA_CRS_HOME/css/init \$ORA_CRS_HOME/evm/log \$ORA_CRS_HOME/evm/init \$ORA_CRS_HOME/srvm/log 4 Check /var/log/messages for any error messages regarding CRS init scripts. 5 Capture all log files for support diagnosis. |

Table 1-10. Troubleshooting (*continued*)

| Category | Problem / Symptom | Cause | Recommended Corrective Action |
|-----------------------|---|---|---|
| CRS | Node continually reboots. | The node does not have access to the quorum disk on the shared storage. | <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 quorum disk is available for read and write. If it is not available, check hardware connections and ensure that OCFS volumes are mounted. 4 Reboot and type <code>/etc/inet.d/init.crs enable</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. |
| DBCA | While creating the seed database using DBCA on OCFS volumes, you get error <code>ORA-60</code> , <code>ORA-06512</code> , or <code>ORA-34740</code> . | Known intermittent issue. | Click Ignore ; the seed database is created normally. |
| Software installation | You receive <code>dd</code> failure error messages while installing the software using <i>Deployment CD 1</i> . | Using copies, rather than the original Red Hat CDs. | Use the original Red Hat CDs that you received with your system. |
| Software installation | Configuration script 005-oraclesetup fails to create <code>/opt</code> directory. | While installing with <i>Deployment CD 1</i> , the mount point <code>/opt</code> was not specified. | Either reinstall the operating system using <i>Deployment CD 1</i> or create a mount point <code>/opt</code> , make user <code>oracle</code> the owner, and rerun 005-oraclesetup from <i>Deployment CD 2</i> . |

Table 1-10. Troubleshooting (*continued*)

| Category | Problem / Symptom | Cause | Recommended Corrective Action |
|-----------------------------|---|---|---|
| Software installation | When connecting to the database as a user other than <code>oracle</code> , you receive the error messages ORA01034: ORACLE not available and Linux Error 13: Permission denied. | 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> . |
| Fibre Channel storage | 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 www.dell.com/oracle for the supported versions. Update as required the driver, BIOS, and firmware for the Fibre Channel HBAs. |
| 8-node OCFS setup with DBCA | You receive the error message ORA-04031 unable to allocate 4180 bytes of shared memory. | The default memory allocation for an 8-node cluster is too small. | In the Initialization Parameters Window , change the value of the Shared Pool to 500 MB from the default value of 95 MB and click Next . |
| PowerPath | One node does not enumerate PowerPath device names in the same order as the other nodes in a cluster, causing a cluster configuration failure. | PowerPath device scan issue. | Type <code>cat /proc/partitions</code> . Verify that each device labeled emcpower_x , where <i>x</i> is a letter assigned by PowerPath, corresponds to a device with the same label and size on each of the other nodes. If this is not the case, copy the file <code>/opt/emcpower/emcpower.conf</code> from the first node to the other nodes and restart PowerPath. |

Getting Help

For detailed information on the use of your system, see the documentation that came with your system components. For training information for your Oracle software and application clusterware, go to www.oracle.com or see your Oracle documentation for information on contacting Oracle. Technical support, downloads, and other technical information are available at the Oracle MetaLink website at metalink.oracle.com. For white papers, Dell Supported Configurations, and general information, visit www.dell.com/oracle. 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 *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.

Obtaining and Using Open Source Files

The software contained on the *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 refer to SKU 420-4534 when making such request. There may be a nominal fee charged to you for the physical act of transferring a copy.

Index

A

- adding and removing nodes, 34
- additional configuration options
 - adding and removing nodes, 34
- additional information, 41
 - configuring automatic reboot, 42
 - determining the private network interface, 43
- ASM configuration, 23

B

- bonding, 17

C

- cluster
 - Fibre Channel hardware connections, example, 12
 - SCSI hardware connections, example, 14
- cluster setup
 - Fibre Channel, 11
 - SCSI, 14

- configuring
 - ASM, 23
 - database storage (single node), 32
 - OCFS, 22
 - Oracle Database 10g (single node), 31
 - Oracle RAC 10g, 16
 - Red Hat Enterprise Linux, 10
 - shared storage, 22
 - shared storage for CRS, 21
 - shared storage using ASM, 23
 - shared storage using OCFS, 22

- configuring automatic reboot, 42
- configuring Oracle 10g, 11
 - verifying hardware and software configurations, 11
- configuring Oracle Database 10g (single node), 31-32
 - creating the seed database, 32
- configuring Oracle RAC 10g, 16
 - creating the seed database, 27
- configuring shared storage, 22
 - ASM, 23
 - OCFS, 22
- configuring shared storage for CRS, 21
- configuring the private and public networks, 16

- configuring the private network, 17
- configuring the public network, 17
- creating the seed database, 27, 32
 - ASM, 28
 - OCFS, 27
- CRS
 - installing, 24
- CRS configuration, 21

D

- deploying Oracle RAC 10g, 16
- determining the private network interface, 43
- documentation, 9

E

- examples
 - Fibre Channel cluster hardware connections, 12
 - SCSI cluster hardware connections, 14

F

- Fibre Channel cluster setup, 11

G

getting help, 50

H

hardware

Fibre Channel cluster
minimum requirements, 8

Fibre Channel
interconnections, 12

SCSI cluster minimum
requirements, 8

SCSI interconnections, 14

single-node minimum
requirements, 9

hardware and software
configurations

Fibre Channel, 13

SCSI, 16

I

installing

CRS, 24

Oracle Database 10g, 25

Oracle Database 10g (single
node), 31

Oracle RAC 10g, 24

Red Hat Enterprise Linux, 9

using Dell Deployment CD, 9

integrated NICs, 44

L

license agreements, 9

listener configuration, 26, 32

N

node

adding and removing, 34

removing, 39

O

OCFS configuration, 22

Oracle Database 10g

installing, 25

installing (single node), 31

single node configuration, 31

Oracle RAC 10g

ASM configuration, 23

configuration, 16

CRS configuration, 21

installing, 24

OCFS configuration, 22

shared storage

configuration, 22

P

passwords

setting, 30, 33

private network

configuring, 16-17

determining the interface, 43

public network

configuring, 16-17

R

Red Hat Enterprise Linux

installing, 9

reinstalling

software, 40

removing a node, 39

S

SCSI cluster setup, 14

setting up nodes, 15

seed database

creating, 27, 32

verifying, 30, 33

setting up SCSI nodes, 15

software

reinstalling, 40

requirements, 7, 41

software and hardware

requirements, 7

supported storage devices, 41

T

troubleshooting, 44

V

verifying

hardware configuration, 11

seed database, 30, 33

software configuration, 11

storage configuration, 19