

Dell PowerEdge Systems
Oracle Database on Microsoft Windows
Server x64

Storage and Network Guide
Version 4.4

Notes, Cautions, and Warnings



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



CAUTION: A CAUTION indicates potential damage to hardware or loss of data if instructions are not followed.



WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Information in this publication is subject to change without notice.

© 2009–2010 Dell Inc. All rights reserved.

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

Trademarks used in this text: Dell™, the DELL logo, EqualLogic™, PowerEdge™, and PowerVault™ are trademarks of Dell Inc.; EMC® and PowerPath® are registered trademarks of EMC Corporation; Intel® is a registered trademark of Intel Corporation in the U.S. and other countries; Microsoft®, Windows®, and Windows Server® are either trademarks or registered trademarks of Microsoft Corporation in the US and/or other countries; Oracle® is a registered trademark of Oracle Corporation in the US and/or other countries.

Other trademarks and trade names may be used in this publication 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

1	Overview	7
	Required Documentation for Deploying the Dell Oracle Database	7
	Terminology Used in This Document	8
	Getting Help	8
	Dell Support.	8
	Oracle Support	9
2	Fibre Channel Cluster Setup	11
	Cabling Your Dell/EMC Fibre Channel Storage	13
3	SAS Cluster Setup for the Dell PowerVault MD3000	15
	Cabling Your SAS Storage System.	17
4	iSCSI Cluster Setup for the Dell PowerVault MD3000i and PowerVault MD1000 Expansion Enclosures	19
	Setting Up iSCSI Cluster With Dell PowerVault MD3000i Storage System and Dell PowerVault MD1000 Expansion Enclosures	21

	Cabling Your iSCSI Storage System.	21
5	iSCSI Cluster Setup for the Dell EqualLogic PS Series Storage Systems	25
	Cabling Dell EqualLogic iSCSI Storage System.	25
6	Configuring Network and Storage for Oracle RAC Database	29
	Configuring the Public and Private Networks.	29
	Configuring and Teaming the Private Network.	30
	Configuring NIC Teaming for Your Private Network Adapters	31
	Configuring the IP Addresses for Your Public and Private Network Adapters	33
	Installing the Host-Based Software Required for Storage.	36
	Installing Multi-Path Software for Storage	36
	Installing EMC PowerPath for Dell/EMC Systems.	36
	Installing Multi-Path Driver Software for PowerVault MD3000 or PowerVault MD3000i	36
	Installing Multi-Path Driver Software for EqualLogic iSCSI Storage Array	37
	Verifying Multi-Path Driver Functionality.	37
	Verifying and Upgrading the Firmware	37
	Configuring Your EqualLogic iSCSI Storage.	37
	Creating Volumes	37
	Configuring iSCSI Networks	39

Configuring Host Access to Volumes.	39
Configuring Microsoft iSCSI Initiator.	40
Verifying the Storage Assignment to the Nodes	41
Preparing the Disks for Oracle Clusterware, Database, and Backup.	42
Enabling the Automount Option for the Shared Disks	43
Preparing the OCR and Voting Disks for Clusterware on Windows Server 2003	43
Preparing the OCR and Votingdisk for Clusterware on Windows Server 2008	45
Preparing the Database Disk and Flash Recovery Area for Database Storage With OCFS.	46
Preparing the Database Disk and Flash Recovery Area for Database Storage With ASM.	47

Overview

The Storage and Networking Guide for Oracle Database on Microsoft Windows applies to:

- Oracle Database 10g R2 Enterprise Edition on Microsoft Windows Server 2003 R2 Standard or Enterprise x64 Edition or Windows Server 2008 SP2 Enterprise or Standard x64 Edition.
- Oracle Database 10g R2 Standard Edition on Windows Server 2003 R2 SP2 Standard x64 Edition or Windows Server 2008 SP2 Standard x64 Edition.

Required Documentation for Deploying the Dell Oracle Database

The documents required for installing the Dell | Oracle database are:

- *Oracle Database on Microsoft Windows Server x64 Operating System and Hardware Installation Guide*—Describes the required minimum hardware and software versions, how to install and configure the operating system, how to verify the hardware and software configurations, and how to obtain open source files.
- *Oracle Database on Microsoft Windows Server x64 Storage and Network Guide*—Describes how to install and configure the network and the storage solutions.
- *Oracle Database on Microsoft Windows Server x64 Database Setup and Installation Guide*—Describes how to install and configure the Oracle database.
- *Oracle Database on Microsoft Windows Server x64 Troubleshooting Guide*—Describes how to troubleshoot and resolve errors encountered during the installation procedures described in the previous modules.



NOTE: All modules provide information on how to receive technical assistance from Dell.

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 Dell PowerVault SAS or iSCSI (Dell PowerVault MD3000 and Dell PowerVault MD3000i with Dell PowerVault MD1000 expansion) storage environment.

Dell EqualLogic PS series storage arrays include storage virtualization technology. To better understand how these arrays operate, it is helpful to be familiar with some of the terminologies used to describe these arrays and their functions:

- **Member**—A single PS series array.
- **Group**—A set of one or more members that can be centrally managed. Host servers access the data through a single group IP address.
- **Pool**—A RAID that can consist of the disks from one or more members.
- **Volume**—A LUN or virtual disk that represents a subset of the capacity of a pool.

Getting Help

This section provides information on contacting Dell or Oracle for whitepapers, supported configurations, training, technical support, and general information.

Dell Support

- For detailed information about using your system, see the documentation that came with your system components.
- For whitepapers, Dell-supported configurations, and general information, see dell.com/oracle.
- For Dell technical support for your hardware and operating system software and to download the latest updates for your system, see support.dell.com. Information about contacting Dell is provided in the *Dell PowerEdge Systems Oracle Database on Microsoft Windows Server x64*

Operating System and Hardware Installation Guide and *Dell PowerEdge Systems Oracle Database on Microsoft Windows Server x64 Troubleshooting Guide* of your system.


- Dell Enterprise Training and Certification is now available; see 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, and for information about contacting Oracle, see oracle.com or see your Oracle documentation.
- Technical support, downloads, and other technical information is available at metalink.oracle.com.

2

Fibre Channel Cluster Setup

 **WARNING:** Before you begin any of the procedures in this section, read the safety information that shipped with your system. For additional best practices information, see dell.com/regulatory_compliance.

After a Dell Managed Services representative completes the setup of your Fibre Channel cluster, verify the hardware connections and the hardware and software configurations as described in this section. Figure 2-1 and Figure 2-2 show an overview of the connections required for the cluster and Table 2-1 summarizes the cluster connections.

Figure 2-1. Hardware Connections for a SAN-Attached Fibre Channel Cluster

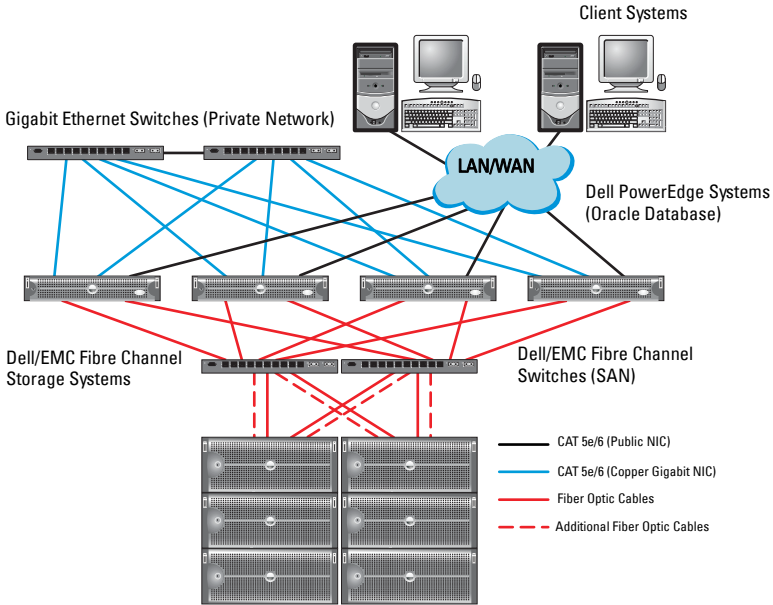


Table 2-1. Fibre Channel Hardware Interconnections

Cluster Component	Connections
Dell PowerEdge system node	<ul style="list-style-type: none"> • One Category 5 enhanced (CAT 5e) or CAT 6 cable from the public NIC to the LAN • One CAT 5e or CAT 6 cable from the private Gigabit NIC to the Gigabit Ethernet switch • One CAT 5e or CAT 6 cable from a redundant private Gigabit NIC to a redundant Gigabit Ethernet switch • One fiber optic cable from HBA 0 to Fibre Channel switch 0 • One fiber optic cable from HBA 1 to Fibre Channel switch 1

Table 2-1. Fibre Channel Hardware Interconnections (continued)

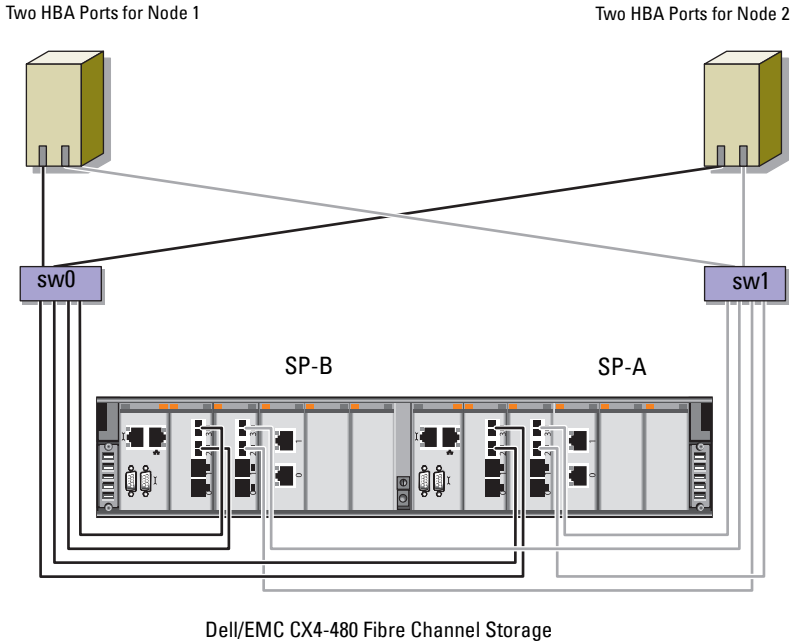
Cluster Component	Connections
Dell/EMC Fibre Channel storage system	<ul style="list-style-type: none">• Two CAT 5e or CAT 6 cables connected to the LAN• One to four fiber optic cable connections to each Fibre Channel switch. For example, for a four-port configuration:<ul style="list-style-type: none">– 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
Dell/EMC Fibre Channel switch	<ul style="list-style-type: none">• One to four fiber optic cable connections to the Dell/EMC Fibre Channel storage system• One fiber optic cable connection to each PowerEdge system HBA
Gigabit Ethernet switch	<ul style="list-style-type: none">• 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

Cabling Your Dell/EMC Fibre Channel Storage

You can configure your Oracle cluster storage system in a four-port storage area network (SAN) attached configuration, depending on your needs. See the following procedures for both configurations.

Figure 2-2 illustrates the hardware connections used in setting up a SAN-attached Fibre Channel cluster.


Figure 2-2. Cabling in a Dell/EMC SAN-Attached Fibre Channer Cluster



To configure your Oracle cluster storage system in a four-port, SAN-attached configuration (see Figure 2-2):

- 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-B port 0 to Fibre Channel switch 1.
- 4 Connect one optical cable from SP-B port 1 to Fibre Channel switch 0.
- 5 Connect one optical cable from HBA 0 on node 1 to Fibre Channel switch 0.
- 6 Connect one optical cable from HBA 1 on node 1 to Fibre Channel switch 1.
- 7 Connect one optical cable from HBA 0 of each additional node to Fibre Channel switch 0.
- 8 Connect one optical cable from HBA 1 of each additional node to Fibre Channel switch 1.

SAS Cluster Setup for the Dell PowerVault MD3000

 **WARNING:** Before you begin any of the procedures in this section, read the safety information that shipped with your system. For additional best practices information, see dell.com/regulatory_compliance.

This section provides information and procedures to configure your Dell PowerEdge systems and PowerVault MD3000 hardware and software to function in an Oracle Real Application Cluster (RAC) environment.

Verify the hardware connections, and the hardware and software configurations as described in this section using Figure 3-1, Table 3-1, and Table 3-2.

Figure 3-1. Cabling the Serial-Attached SCSI (SAS) Cluster and Dell PowerVault MD3000

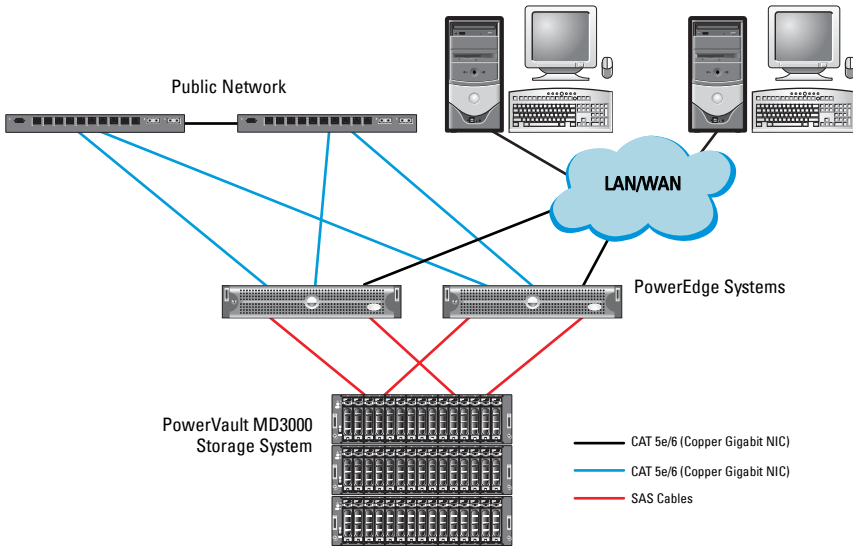


Table 3-1. SAS Cluster Hardware Interconnections

Cluster Component	Connections
PowerEdge system node	<ul style="list-style-type: none"> • 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 a PowerVault MD3000 storage system node using a SAS 5/E. See "Cabling Your SAS Storage System" on page 17.

Table 3-1. SAS Cluster Hardware Interconnections (continued)

Cluster Component	Connections
PowerVault MD3000	<ul style="list-style-type: none">• Two CAT 5e/6 cables connected to a LAN (one from each storage processor module).• Two SAS connections to each PowerEdge system node using a SAS 5/E controller. See "Cabling Your SAS Storage System" on page 17.
Gigabit Ethernet switch	<ul style="list-style-type: none">• 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.

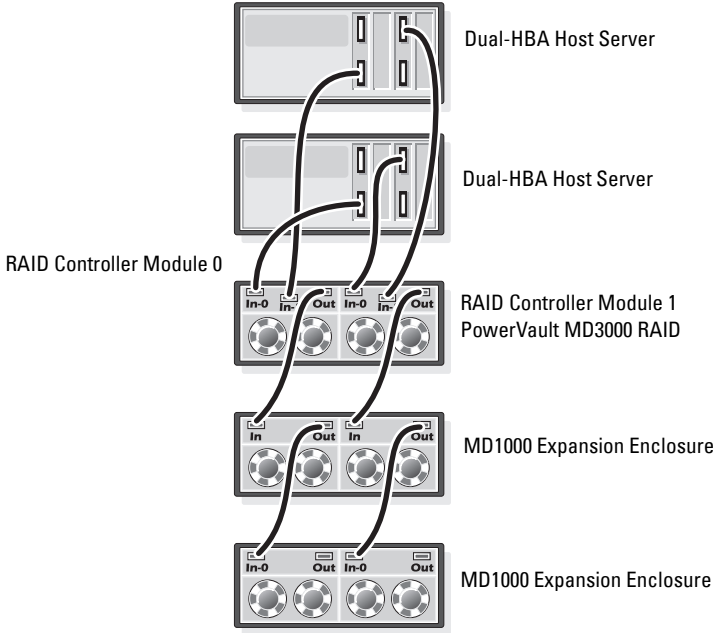
Cabling Your SAS Storage System

SAS clusters can only be installed in a direct-attached cluster and are limited to two nodes only.

To configure your nodes in a direct-attached configuration (see Figure 3-2):

- 1** Connect one SAS cable from a port of the SAS controller of node 1 to the In-0 port of RAID controller 0 in the PowerVault MD3000 RAID enclosure.
- 2** Connect one SAS cable from a port of the other SAS controller of node 1 to the In-0 port of RAID controller 1 in the PowerVault MD3000 RAID enclosure.
- 3** Connect one SAS cable from a port of the SAS controller of node 2 to the In-1 port of RAID controller 0 in the PowerVault MD3000 RAID enclosure.
- 4** Connect one SAS cable from a port of the other SAS controller of node 2 to the In-1 port of RAID controller 1 in the PowerVault MD3000 RAID enclosure.

Figure 3-2. Cabling in a Direct-Attached SAS Cluster



iSCSI Cluster Setup for the Dell PowerVault MD3000i and PowerVault MD1000 Expansion Enclosures

⚠ WARNING: Before you begin any of the procedures in this section, read the safety information that shipped with your system. For additional best practices information, see dell.com/regulatory_compliance.

This section provides information and procedures to configure your Dell PowerEdge systems and Dell PowerVault MD3000i hardware and software to function in an Oracle Real Application Cluster (RAC) environment.

Verify the hardware connections, and the hardware and software configurations, using the Supported Configuration figures contained in the *Dell PowerVault MD3000i Support Matrix* available at support.dell.com.

Table 4-1. iSCSI Hardware Interconnections

Cluster Component	Connections
One PowerEdge system node	<ul style="list-style-type: none"> • 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). <p>NOTE: For additional information on Dell PowerVault MD3000i system see your <i>Dell PowerVault MD3000i SetUp</i> documentation.</p>

Table 4-1. iSCSI Hardware Interconnections

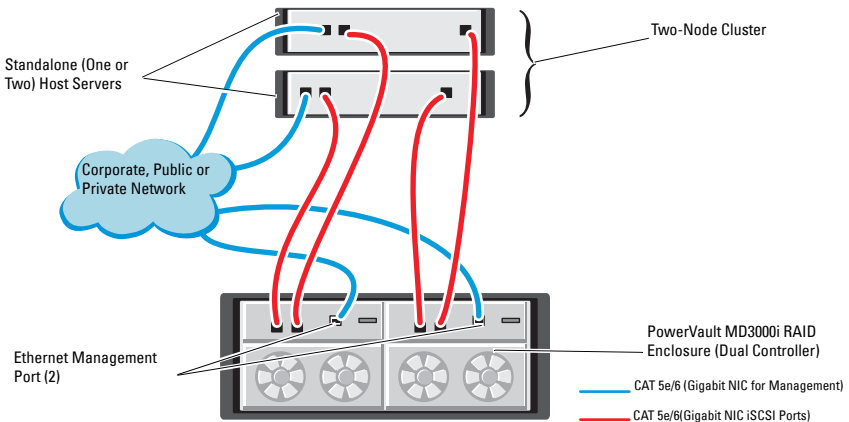
Cluster Component	Connections
One Dell PowerVault MD3000i storage system	<ul style="list-style-type: none">• Two CAT 5e/6 cables connected to LAN (one from each storage processor module) for the management interface.• Two CAT 5e/6 cables per storage processor for iSCSI interconnect. <p>NOTE: For additional information on Dell PowerVault MD3000i system see your <i>Dell PowerVault MD3000i Setup</i> documentation.</p>
One Dell PowerVault MD1000 storage expansion enclosure (optional)	Additional Serial-attached SCSI (SAS) cable connections as required for the MD1000 expansion enclosures.

Setting Up iSCSI Cluster With Dell PowerVault MD3000i Storage System and Dell PowerVault MD1000 Expansion Enclosures

Cabling Your iSCSI Storage System

Direct-attached iSCSI clusters are limited to two nodes only.

Figure 4-1. Cabling iSCSI Direct-Attached Clusters



To configure your nodes in a direct-attached configuration (see Figure 4-1):

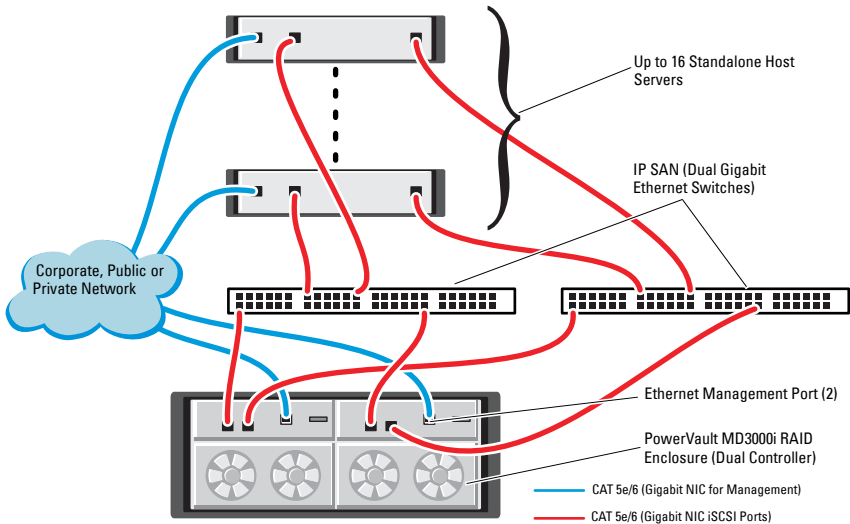
- 1 Connect one CAT 5e/6 cable from a port (iSCSI HBA or NIC) of node 1 to the In-0 port of RAID controller 0 in the Dell PowerVault MD3000i storage enclosure.
- 2 Connect one CAT 5e/6 cable from the other port (iSCSI HBA or NIC) of node 1 to the In-0 port of RAID controller 1 in the Dell PowerVault MD3000i storage enclosure.
- 3 Connect one CAT 5e/6 cable from a port (iSCSI HBA or NIC) of node 2 to the In-1 port of RAID controller 0 in the Dell PowerVault MD3000i storage enclosure.
- 4 Connect one CAT 5e/6 cable from the other port (iSCSI HBA or NIC) of node 2 to the In-1 port of RAID controller 1 in the Dell PowerVault MD3000i storage enclosure.

- 5 Connect two SAS cables from the two MD3000 out ports to the two In ports of the first Dell PowerVault MD1000 expansion enclosure (Optional).
- 6 Connect two SAS cables from the two MD1000 out ports to the In-0 ports of the second Dell PowerVault MD1000 expansion enclosure (Optional).

NOTE: For information on configuring the PowerVault MD1000 expansion enclosure, see the *Dell PowerVault MD3000 Storage System* documentation available at support.dell.com/manuals.

Switched iSCSI clusters can support up to eight nodes.

Figure 4-2. Cabling iSCSI Switched Clusters



To configure your nodes in a switched configuration (see Figure 4-2):

- 1 Connect one CAT 5e/6 cable from a port (iSCSI HBA or NIC) of node 1 to the port of network switch 1.
- 2 Connect one CAT 5e/6 cable from a port (iSCSI HBA or NIC) of node 1 to the port of network switch 2.
- 3 Connect one CAT 5e/6 cable from a port (iSCSI HBA or NIC) of node 2 to the port of network switch 1.

- 4 Connect one CAT 5e/6 cable from a port (iSCSI HBA or NIC) of node 2 to the port of network switch 2.
- 5 Connect one CAT 5e/6 cable from a port of switch 1 to the In-0 port of RAID controller 0 in the Dell PowerVault MD3000i storage enclosure.
- 6 Connect one CAT 5e/6 cable from the other port of switch 1 to the In-0 port of RAID controller 1 in the Dell PowerVault MD3000i storage enclosure.
- 7 Connect one CAT 5e/6 cable from a port of switch 2 to the In-1 port of RAID controller 0 in the Dell PowerVault MD3000i storage enclosure.
- 8 Connect one CAT 5e/6 cable from the other port of switch 2 to the In-1 port of RAID controller 1 in the Dell PowerVault MD3000i storage enclosure.
- 9 Connect two SAS cables from the two MD3000i out ports to the two In ports of the first Dell PowerVault MD1000 expansion enclosure (Optional).
- 10 Connect two SAS cables from the two MD1000 out ports to the In-0 ports of the second Dell PowerVault MD1000 expansion enclosure (Optional).




NOTE: For information on configuring the Dell PowerVault MD1000 expansion enclosure, see the *Dell PowerVault MD3000 Storage System* documentation available at support.dell.com/manuals.



NOTE: See the Solutions Deliverable List (SDL) found at dell.com/oracle and select the appropriate solution. After selecting the solution, verify that the firmware version for your storage is at or above the required firmware version found in the SDL.

iSCSI Cluster Setup for the Dell EqualLogic PS Series Storage Systems

 **WARNING:** Before you begin any of the procedures in this section, read the safety information that shipped with your system. For additional best practices information, see dell.com/regulatory_compliance.

Cabling Dell EqualLogic iSCSI Storage System

Host servers can be attached to the Dell EqualLogic iSCSI array through an IP storage area network (SAN) industry-standard Gigabit Ethernet switch.

Figure 5-1 shows the recommended network configuration for a dual control module PS5000XV array. This configuration includes two Dell PowerConnect 54xx Gigabit Ethernet switches, to provide highest network availability and maximum network bandwidth.

It is recommended that you use two Gigabit Ethernet switches because in the event of a switch failure in a single ethernet switch environment, all hosts lose access to the storage until the switch is physically replaced and the configuration restored. For such configurations there should be multiple ports with link aggregation providing the inter-switch, or trunk connection.

Additionally, it is recommended that one Gigabit interface connects to one Ethernet switch from each of the control modules, and the other two Gigabit interfaces connect to the other Ethernet switch.

Figure 5-1. Recommended Network Configuration

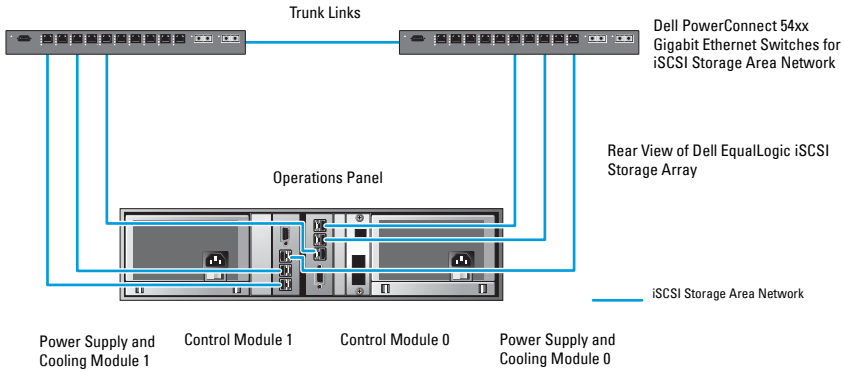


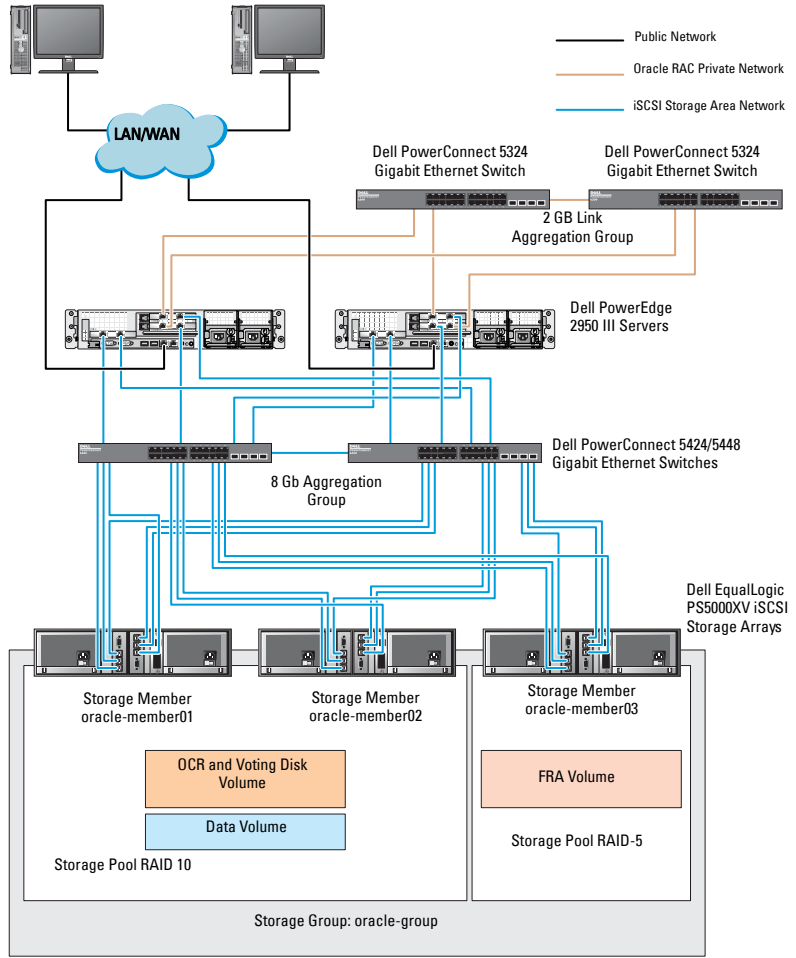
Figure 5-2 is an architecture overview of a sample Oracle RAC configuration with three PS5000XV arrays.

Table 5-1. Cable Color Legend

Cable Color	Denotes
Blue	iSCSI storage area network (SAN)
Brown	Oracle RAC private interconnect network
Black	public network

Table 5-1 describes the cable colors and their significance. The PS5000XV storage arrays provide the physical storage capacity for the Oracle RAC database.

Figure 5-2. Sample Oracle RAC Configuration With Three PS5000XV Arrays



As illustrated in Figure 5-2, the group named **oracle-group** includes three PS5000XV members:

- oracle-member01
- oracle-member02
- oraclemember03

When a member is initialized, it can be configured with RAID 10, RAID 5, or RAID 50. For more information on how to initialize an EqualLogic array, see the *Dell EqualLogic User's Guide*.

A PS Series storage group can be segregated into multiple tiers or pools. Tiered storage provides administrators with greater control over how disk resources are allocated. At any one time, a member can be assigned to only one pool. It is easy to assign a member to a pool and also to move a member between pools with no impact to data availability. Pools can be organized according to different criteria, such as disk types or speeds, RAID levels, and application types.

In Figure 5-2 pools are organized by member RAID levels:

- One pool with the name RAID-10 consists of RAID 10 members.
- One pool with the name RAID-5 consists of RAID 5 members.

Configuring Network and Storage for Oracle RAC Database

This section provides information about:

- Configuring the public and private networks.
- Verifying the storage configuration.
- Configuring the shared storage for Oracle Clusterware and the Oracle Database.



NOTE: Oracle RAC requires an ordered list of procedures. To configure networking and storage in a minimal amount of time, perform the procedures listed in this chapter in order.

Configuring the Public and Private 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 address. All public IP addresses, including the virtual IP address should be registered with the domain naming service (DNS). If a DNS server is not available, IP addresses have to be registered in the hosts file on all cluster nodes.

Depending on the number of NIC ports available, configure the public and private interfaces as shown in Table 6-1.

Table 6-1. NIC Port Assignments

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

Configuring and Teaming the Private Network

Before you deploy the cluster, assign a private IP address and host name to each cluster node. This procedure ensures that the nodes can communicate with each other through the private interface.

Table 6-2 provides an example of a network configuration for a two-node cluster.



NOTE: This example assumes all the IP addresses are registered in the hosts file of all cluster nodes.



NOTE: The two bonded NIC ports for a private network should be on separate PCI buses. For example, a bonded pair can consist of one on-board NIC and one add-on NIC card.

Table 6-2. Network Configuration Example for a Two-Node Cluster

Host Name	Type	IP Address	Registered In
rac1	Public	155.16.170.1	%SystemRoot%\system32\drivers\etc\hosts
rac2	Public	155.16.170.2	%SystemRoot%\system32\drivers\etc\hosts
rac1-vip	Virtual	155.16.170.201	%SystemRoot%\system32\drivers\etc\hosts
rac2-vip	Virtual	155.16.170.202	%SystemRoot%\system32\drivers\etc\hosts
rac1-priv	Private	10.10.10.1	%SystemRoot%\system32\drivers\etc\hosts
rac2-priv	Private	10.10.10.2	%SystemRoot%\system32\drivers\etc\hosts

Configuring NIC Teaming for Your Private Network Adapters



NOTE: TCP Offload Engine (TOE) functionality of a TOE-capable NIC is not supported in this solution.


To configure NIC teaming for your private network adapters:

- 1** On node 1, identify two network adapters that are used for NIC teaming.
- 2** Connect an ethernet cable from each selected network adapter to the private network switch.
- 3** If node 1 is configured with Broadcom NICs, go to step 4. If node 1 is configured with Intel NICs, configure NIC teaming by performing the following steps:
 - a** Right-click **My Computer** and select **Manage**.
 - b** On the **Computer Management** window, select **Device Manager**.
 - c** Expand the **Network Adapters** tab.
 - d** Right-click the **Intel NIC**, which is identified for NIC teaming and select **Properties**.
 - e** Click the **Teaming** tab.
 - f** Select **Team with other Adapters** and then select **New Team**.
 - g** Specify a name for NIC team and click **Next**.
 - h** On the **Select the adapters to include in this team** box, select the remaining network adapters that you identified for NIC teaming and click **Next**.
 - i** On the **Select a team mode** list box, select **Adaptive Load Balancing**.
 - j** Click **Finish** to complete the teaming.
 - k** On the **Team Properties** window, click **OK**.
 - l** On the Intel NIC's **Properties** window, click **OK**.
 - m** Close the **Computer Management** window.


- 4 If node 1 is configured with Broadcom NICs, configure NIC teaming by performing the following steps. If not go to step 5.
 - a Click **Start**→**Programs**→**Broadcom**→**Broadcom Advanced Control Suite 3**.

The **Broadcom Advanced Control Suite 3** window is displayed.
 - b Highlight **Team Management**, and click **Teams** and select **Create a Team**.

The **Broadcom Teaming Wizard** window is displayed.
 - c Click **Next**.
 - d In the **Enter the name for the team** field, type `Private` and click **Next**.
 - e On the **Team Type** field, select **Smart Load Balancing (TM) and Failover (SLB)** and click **Next**.
 - f On the **Assigning Team Members** window, in the **Available Adapters** box, select the network adapters you identified for NIC teaming and add them to the **Team Members** box.


 **NOTE:** When configuring integrated Broadcom NICs into a team, the following warning message is displayed on Dell PowerEdge 6850 systems.
Including this adapter in a team will disrupt the system management features. Click **Yes** to proceed.
 - g Click **Next**.
 - h On the **Designating a Standby Member** window, either select a NIC or select **Do not configure a standby member**.
 - i On the **Configuring Live Link** window, select **No** and click **Next**.
 - j On the **VLAN Support** window, if there is a VLAN you require to configure, select **Add VLAN**. If not, select **Skip manage VLAN**.
 - k Click **Preview** to verify the NIC team and the adapters. After verification, select **Commit changes to system** and exit the wizard. Click **Finish**.
 - l On the **Information Message** window, click **Yes** to proceed.
 - m On the **Broadcom Advanced Control Suite 3** window, click **File** then **Exit**.
- 5 Repeat step 1 to step 4 on the remaining nodes.

Configuring the IP Addresses for Your Public and Private Network Adapters


 **NOTE:** The TOE functionality of TOE-capable NIC is not supported in this solution.

To configure the IP addresses for your public and private network adapters:

- 1 Update the adapter's network interface name, if required. Otherwise, go to step 3.
 - a On node 1, click **Start** and navigate to **Settings**→**Control Panel**→**Network Connections**.
 - b On the **Network Connections** window, right-click the public network adapter you want to rename, and select **Rename**.

 **NOTE:** When you configure your network adapters on the cluster nodes, use identical names for the public adapters and the private adapters on all cluster nodes. Otherwise, the Oracle database installer generates an error and prevents you from completing the installation procedure.

- c Rename the public adapter name to **Public** and press <Enter>.
 - d Right-click the Private NIC team you want to rename and select **Rename**.
 - e Rename the Private NIC team to **Private** and press <Enter>.
- 2 Configure the IP addresses.

 **NOTE:** You must set a default gateway for your public interface, otherwise, the Clusterware installation may fail.


- a Click **Start**→**Settings**→**Control Panel**→**Network Connections**→**Public**→**Properties**.
 - b Double-click **Internet Protocol (TCP/IP)**.
 - c Click **Use the following IP address**, enter the required IP address, default gateway address and the DNS server IP address, and click **OK**.
 - d In the **Public Properties** window, select **Show icon in notification area when connected**.

The network adapter status appears in the system tray after you perform the procedures in this section.

- e Click **OK**.

f On the **Properties** window, click **Close**.

g Repeat step a through step f on the Private NIC team.

 **NOTE:** Private NIC team does not require a default gateway address and DNS server entry.

3 Ensure that the public and private network adapters appear in the appropriate order for access by network services.

a On the Windows desktop, click **Start**→**Settings**→**Control Panel**→**Network Connections**.

b On the **Network Connections** window, click **Advanced** and select **Advanced Settings**.


c Select the **Adapter and Bindings** tab.

d Ensure that the network adapters are listed in the following order:

i. Public

ii. Private


iii. <Any other network adapter>

 **NOTE:** Click the up-and down-arrow keys to change the adapter order.

e Click **OK**.

f Close the **Network Connections** window.

4 On all nodes, add the public, private, and virtual IP addresses and host name to the `%SystemRoot%\system32\drivers\etc\hosts` file.


 **NOTE:** Add the public and virtual IP addresses to the hosts file only if they are not registered with the DNS server.

For example, the following entries use the adapter IP and host name as shown in Table 6-3.

Table 6-3. IP Address and Node Names

IP Address	Node Name
155.16.170.1	rac1
155.16.170.2	rac2
10.10.10.1	rac1-priv
10.10.10.2	rac2-priv

IP Address	Node Name
155.16.170.201	rac1-vip
155.16.170.202	rac2-vip

 **NOTE:** Registering the private IP addresses with the DNS server is not required because the private network IP addresses are not accessible from the public network.

- 5 Repeat step 1 to step 4 on the remaining nodes.
- 6 Ensure that the cluster nodes can communicate with the public and private networks.

a On node 1, open a command prompt window.


b At the command prompt, type:

```
ping <public_host_name>
```

```
ping <private_host_name>
```

where *<public_host_name>* and *<private_host_name>* are the host names for the public and private network adapters in the remaining node.

If the node's network adapters do not respond to ping commands, check your network configuration and then repeat this step.

 **NOTE:** Virtual internet protocol (VIP) is configured later and cannot be pinged at this point.

c Repeat step a through step b on the remaining nodes.

Installing the Host-Based Software Required for Storage

If you are installing Dell/EMC Fibre Channel Storage, see the Dell/EMC documentation that came with your system to install the EMC Naviagent software.

If you are installing a Dell PowerVault storage, see the Dell PowerVault documentation that came with your system to install the Modular Disk Storage Manager (MDSM) software from the *Dell PowerVault Resource* media.



NOTE: The Microsoft iSCSI Software Initiator is required for a Dell PowerVault MD3000i storage system. The Microsoft iSCSI Software Initiator can be downloaded from microsoft.com.

Installing Multi-Path Software for Storage

Follow the documentation accompanying the storage device to install the appropriate multi-path software.

Installing EMC PowerPath for Dell/EMC Systems

To install EMC PowerPath for Dell/EMC Systems:

- 1 On node 1, install EMC PowerPath.



NOTE: For more information, see the *EMC PowerPath* documentation that came with your Dell/EMC storage system.

- 2 On completion of the installation procedure, restart your system.
- 3 Repeat step 1 and step 2 on the remaining nodes.

Installing Multi-Path Driver Software for PowerVault MD3000 or PowerVault MD3000i

In the case of a SAS or iSCSI (PowerVault MD3000 or PowerVault MD3000i storage system) cluster, the Multi-Path software should already be installed on your hosts as instructed in the section "Installing the Host-Based Software Required for Storage" on page 36.

Installing Multi-Path Driver Software for EqualLogic iSCSI Storage Array

For more information see "Installing and Configuring Dell EqualLogic Host Integration Tool (HIT) Kit" on page 39.

Verifying Multi-Path Driver Functionality

To verify the multi-path driver functionality:

- 1 Right-click **My Computer** and select **Manage**.
- 2 Expand **Storage** and click **Disk Management**.
One disk is displayed for each LUN assigned in the storage.
- 3 Ensure that each LUN is configured as a **Basic** disk.
- 4 Repeat step 1 to step 3 on the remaining nodes.

Verifying and Upgrading the Firmware

- Discover the host server's direct-attached storage using the MDSM software that is installed on the host server.
- Verify that the firmware for the following storage components is at the minimum required version. Refer to the Solutions Deliverable List (SDL) for the firmware version requirements.
 - PowerVault MD3000i storage system firmware
 - MD1000 expansion enclosure firmware

Configuring Your EqualLogic iSCSI Storage

Creating Volumes

Before data can be stored, the PS5000XV physical disks must be configured into usable components, known as volumes. A volume represents a portion of the storage pool, with a specific size, access controls, and other attributes. A volume can be spread across multiple disks and group members and is seen on the network as an iSCSI target. Volumes are assigned to a pool and can be easily moved between pools, with no impact on data availability. In addition,

automatic data placement and automatic load balancing occurs within a pool, based on the overall workload of the storage hardware resources within the pool.

Table 6-4. Volumes for Oracle RAC Configuration

Volume	Minimum Size	RAID	Number of Partitions	Used For	Operating System Mapping
First Area Volume	3 GB	10	One extended partition with five logical drives - 3 x 250 MB for Voting Disk and 2 x 300 MB for OCR	Voting disk and Oracle Cluster Registry (OCR)	One extended partition with five logical drives: 3 x Voting Disk, 2 x OCR
Second Area Volume(s)	Larger than the size of your database	10	One	Data	ASM disk group DATABASEDG
Third Area Volume(s)	Minimum twice the size of your second area volume(s)	5	One	Flash Recovery Area	ASM disk group FLASHBACKDG


 **NOTE:** It is recommended that you keep the OCR mirror and the voting disk mirrors on a different volume to avoid warning messages during OCR and voting disk configuration.

Table 6-4 shows a sample volume configuration. Create volumes in PS5000XV array and create an access list to allow all host iSCSI network interfaces to access the volumes.

Configuring iSCSI Networks

It is recommended that the host network interfaces for iSCSI traffic are configured to use **Flow Control** and **Jumbo** frame for optimal performance.

To set **Flow Control** and **Jumbo** frame:

- 1 Select **Start**→**Settings**→**Network Connections**.
- 2 Highlight the iSCSI network interface, and right click **Properties**.
- 3 Click **Configure**.
- 4 Click **Advanced**.
- 5 Highlight **Jumbo Packet**, and set its value to 9014 bytes.
- 6 Highlight **Performance Options**, then click **Properties**.
The **Performance Options** window is displayed.
- 7 In the **Performance Options** window, highlight **Flow Control**, and set its value to **Rx & Tx Enabled**.
- 8 Repeat step 2 to step 7 for all other network interfaces used for iSCSI traffic.

Configuring Host Access to Volumes

Installing and Configuring Dell EqualLogic Host Integration Tool (HIT) Kit

To install and configure the Dell EqualLogic Host Integration Tool (HIT) kit:

- 1 Download the latest HIT tool kit on your database server.
- 2 Double-click **Setup64.exe** to launch the installation window.
- 3 In the **Welcome to the Host Integration Tools**, update program window, click **Next**.
- 4 In the **License Agreements** window, select **I accept the terms of all the license agreements**, and click **Next**.
- 5 On the **Installation Type** window, click **Next**.
- 6 On the **Ready to install the components** window, click **Install**.
The **Installation Status** window and the **Software Update Installation Wizard** window is displayed.
- 7 On the **Installation Complete** window, click **Finish**.

- 8** On the **System Restart Required** window, select **Yes, I want to restart my computer now**, and click **OK**.
- 9** When the server restarts, a **Remote Setup Wizard** window is displayed.
- 10** Select **Configure MPIO** settings for this computer, then click **Next**.
- 11** Move the iSCSI network subnets under **Subnets included for MPIO**. Move all other network subnets under **Subnets excluded from MPIO**. Select Default load balancing policy (Least Queue Depth). Click **Finish**.
- 12** Repeat step 1 to step 11 on all other hosts in the cluster.

Configuring Microsoft iSCSI Initiator

To configure Microsoft iSCSI initiator:

- 1** Double-click the **Microsoft iSCSI Initiator** icon on desk top.
- 2** On the **iSCSI Initiator Properties** window **Discovery** tab, the IP address of the EqualLogic group address should already be populated by the HIT tool kit.
- 3** On the **iSCSI Initiator Properties** window **Targets** tab, click on one volume that was created for the Oracle database. Click **Log On**.
- 4** On the **Log On to Target** window, check the two boxes by **Automatically restore this connection when the system boots, and Enable multi-path**.
- 5** Click **Advanced**.
- 6** On the **Advanced Settings** window **General** tab, perform the following configurations:
 - a** Choose **Microsoft iSCSI Initiator by the Local adapter** drop-down menu.
 - b** Choose an iSCSI initiator IP address by the **Source IP** drop-down menu.
 - c** Choose **EqualLogic group IP address** by the **Target Portal** drop-down menu.
 - d** Select the **CHAP logon information** check box.
 - e** Enter the CHAP user name defined in EqualLogic storage, by the **User name** box.

- f** Enter the CHAP password defined in EqualLogic storage, by the **Target** secret box.
- g** Click **OK**.
- 7** On the **Log On to Target** window, click **OK**.
- 8** On the **iSCSI Initiator Properties** window **Targets** tab, the status of the logged on volume should be **Connected**.
- 9** Repeat step 3 to step 8 to log on to the same volume for every other iSCSI initiator IP addresses.
- 10** Repeat step 3 to step 9 to log on to all other volumes created for the database.
- 11** Repeat step 1 to step 10 on all other hosts in the cluster.


Verifying the Storage Assignment to the Nodes

To verify the storage assignment to the nodes:

- 1** On the Windows desktop, right-click **My Computer** and select **Manage**.
- 2** On the **Computer Management** window, click **Device Manager**.
- 3** Expand **Disk drives**.
- 4** Under **Disk drives**, ensure that four small computer system interface (SCSI) disk devices appear for each LUN or virtual disk assigned in the storage.
- 5** Expand **Storage** and click **Disk Management**.

If the **Welcome to the Initialize and Convert Disk Wizard** is displayed, perform step a through step d. Otherwise, go to step 6.

- a** On the **Welcome to the Initialize and Convert Disk Wizard** window, click **Next**.
- b** On the **Select Disks to Initialize** window, in the **Disks** window, select the disks that are associated with your storage LUNs/virtual disks and click **Next**.
- c** In the **Select Disks to Convert** window, clear the disk(s) that you selected in step b and click **Next**.

 **NOTE:** This procedure ensures that your disks are configured as basic disks.

- d** Click **Finish**.

- 6 On the **Disk Management** window, verify if four disks appear. The disks should be similar in size to each other and to the LUNs/virtual disks that are assigned to the nodes in the storage system.
- 7 Repeat step 1 to step 6 on the remaining nodes.

Preparing the Disks for Oracle Clusterware, Database, and Backup

This section provides information on creating the logical drives for the following disks:

- **Oracle Cluster Registry disk (OCR)**—Contains the cluster configuration information
- **Voting disk**—Provides arbitration between the cluster nodes when the private network or attached storage is unavailable to one or more nodes
- **Data and backup disks**—Provide storage areas for creating the database (data disk) and saving the backup and log data (backup disk)

During the cluster configuration described in this document, you create partitions on your shared storage. When you create the partitions, ensure that the cluster nodes can detect the LUNs or logical disks that are created in the attached storage system.

To prepare the disks for Oracle Clusterware, identify the OCR, voting, data, and flash recovery area disks. After you identify the appropriate disks, perform the following steps on node 1:

- 1 Enable the Automount option for the shared disks.
- 2 Prepare the OCR and voting disks for Clusterware on Windows Server 2003.
- 3 Prepare the OCR and voting disk for Clusterware on Windows Server 2008.
- 4 Prepare the database disk and flash recovery area for database storage with OCFS.
- 5 Prepare the database disk and flash recovery area for database storage with ASM.

Enabling the Automount Option for the Shared Disks

To enable the Automount option for the shared disks:

- 1 On node 1, click **Start** and select **Run**.
- 2 In the **Run** field, type `cmd` and click **OK**.
- 3 At the command prompt, type `diskpart` and press <Enter>.
- 4 At the `DISKPART` command prompt, type `automount enable` and press <Enter>.
The following message is displayed: **Automatic mounting of new volumes enabled.**
- 5 At the `DISKPART` command prompt, type `exit` and press <Enter>.
- 6 Close the command prompt.
- 7 Repeat step 1 to step 6 on each of the remaining nodes.

Preparing the OCR and Voting Disks for Clusterware on Windows Server 2003


To prepare the OCR and voting disks for Clusterware on Windows Server 2003:

- 1 On the Windows desktop, right-click **My Computer** and select **Manage**.
- 2 Expand **Storage** and click **Disk Management**.
The storage disk that you initialized in the "Verifying Multi-Path Driver Functionality" on page 37 is displayed as **Unallocated**.
- 3 Right-click the partition area of the first shared disks assigned to the cluster nodes and select **New Partition**.
The **Welcome to the New Partition** wizard is displayed.
- 4 Click **Next**.
- 5 On the **Select Partition Type** window, select **Extended partition** and click **Next**.
- 6 On the **Specify Partition Size** window, accept the default partition size and click **Next**.
- 7 Click **Finish**.


The disk partition area you selected in step 3 is configured as an extended partition.

- 8** Repeat step 3 to step 7 on all shared disks that are assigned to the cluster nodes.
- 9** Create a logical drive for the OCR disk.
 - a** On the partition area of the disk identified for OCR and voting disk (2 GB LUN/virtual disk), right-click the free space and select **New Logical Drive**.
The **Welcome to the New Partition** wizard is displayed.
 - b** Click **Next**.
 - c** On the **Select Partition Type** window, select **Logical drive** and click **Next**.
 - d** On the **Specify Partition Size** window, type 120 in the **Partition size in MB** field and click **Next**.
 - e** On the **Assign Drive Letter or Path** window, select **Do not assign a drive letter or drive path** and click **Next**.
 - f** On the **Format Partition** window, select **Do not format this partition** and click **Next**.
 - g** Click **Finish**.
 - h** Repeat step a through step g to create an additional OCR disk.
- 10** Create a logical drive for the Voting Disk.
 - a** On the partition area of the disk identified for the OCR and voting disk (2 GB LUN/virtual disk), right-click the free space and select **New Logical Drive**.
The **Welcome to the New Partition** wizard is displayed.
 - b** Click **Next**.
 - c** In the **Select Partition Type** window, select **Logical drive** and click **Next**.
 - d** In the **Specify Partition Size** window, in the **Partition size in MB** field, type 50 and click **Next**.
 - e** In the **Assign Drive Letter or Path** window, select **Do not assign a drive letter or drive path** and click **Next**.

- f In the **Format Partition** window, select **Do not format this partition** and click **Next**.
- g Click **Finish**.
- h Repeat step a to step g to create two additional voting disk partitions.

 **NOTE:** If you are using Redundant Voting Disk and OCR, repeat the steps outlined in step 9 and step 10 for the redundant Voting Disk and OCR.

Preparing the OCR and Votingdisk for Clusterware on Windows Server 2008

 **NOTE:** The steps to create extended partitions and to create logical drives are identical to the steps of Windows Server 2003, however, it is no longer possible to perform these steps in Windows Server 2008 from the Disk Management GUI.


To create extended partitions and logical drives:

- 1 **Start**→**Run** and type `cmd` and click **Enter** to open a command prompt window.
- 2 Type `diskpart` in the command prompt to open the diskpart interface.
- 3 Enter `list disk` and a table is displayed as shown:

```
DISKPART> list disk
```


Disk ###	Status	Size	Free
Disk 0	Online	37 GB	0 MB
Disk 1	Online	2048 MB	2014 MB
Disk 2	Online	100 GB	100 GB
Disk 3	Online	200 GB	200 GB


- 4 Enter `select disk [disk #]` of the target disk.

 **NOTE:** The following steps explain how to create partitions and logical disks for MBR disk.

- 5 Type `create partition extended` and press **Enter**.
- 6 Type `create partition logical size=<size of volume in megabytes>` and press **Enter**.

Preparing the Database Disk and Flash Recovery Area for Database Storage With OCFS

 **NOTE:** When using Automatic Storage Management (ASM), the ASM data disk group should be larger than your database (multiple LUNs) and the ASM Flash Recovery Area disk group should be at least twice the size of your data disk group.

 **NOTE:** If you are creating the logical drives that are used to create the ASM storage disk, ignore the following steps and see "Preparing the Database Disk and Flash Recovery Area for Database Storage With OCFS" on page 46.

To create logical drives that are used to create the Oracle's Clustered File System (OCFS) storage disk:

- 1** On node 1, create one logical drive for the Database.
 - a** Locate the disk that is assigned for the Oracle Database.
 - b** On the disk partition area, right-click the free space and select **New Logical Drive**.
The **Welcome to the New Partition** wizard is displayed.
 - c** Click **Next**.
 - d** In the **Select Partition Type** window, select **Logical drive** and click **Next**.
 - e** In the **Specify Partition Size** window in the **Partition** size in **MB** field, type the appropriate size and click **Next**.
 - f** In the **Assign Drive Letter or Path** window, select **Do not assign a drive letter or drive path** and click **Next**.
 - g** In the **Format Partition** window, select **Do not format this partition** and click **Next**.
 - h** Click **Finish**.
- 2** On node 1, create one logical drive for the Flash Recovery Area.
 - a** Locate the disk that is assigned for the Flash Recovery Area.
 - b** Perform step b to step h in step 1.
- 3** Restart all other nodes and login as administrator.

Preparing the Database Disk and Flash Recovery Area for Database Storage With ASM



NOTE: If you are creating the logical drives that are used to create the OCFS storage disk, ignore the following steps and follow the procedures in "Preparing the Database Disk and Flash Recovery Area for Database Storage With OCFS" on page 46.

To create logical drives that are used to create ASM disk storage:

- 1** Create one logical drive for the Database.
 - a** Locate the disk that is assigned for the Oracle database.
 - b** On the disk partition area, right-click the free space and select **New Logical Drive**.
The **Welcome to the New Partition** wizard is displayed.
 - c** Click **Next**.
 - d** In the **Select Partition Type** window, select **Logical drive** and click **Next**.
 - e** In the **Specify Partition Size** window, type the appropriate size in the **Partition size in MB** field and click **Next**.
 - f** In the **Assign Drive Letter or Path** window, select **Do not assign a drive letter or drive path** and click **Next**.
 - g** In the **Format Partition** window, select **Do not format this partition** and click **Next**.
 - h** Click **Finish**.
- 2** Create one logical drive for the Flash Recovery Area.
 - a** Locate the disk that is assigned for the Flash Recovery Area.
 - b** Perform step b to step h in step 1.
- 3** Restart all other nodes and log in as the administrator.

Removing the Assigned Drive Letters

To remove the assigned drive letters:

- 1** On the Windows desktop for each node, right-click **My Computer** and select **Manage**.
- 2** On the **Computer Management** window, expand **Storage** and click **Disk Management**.
- 3** If you find any drive letters assigned to the drives that you created in "Preparing the OCR and Voting Disks for Clusterware on Windows Server 2003" on page 43, perform the following steps:
 - a** Right-click the logical drive and select **Change Drive Letter and Paths**.
 - b** In the **Change Drive Letter and Paths** window, select the drive letter and click **Remove**.
 - c** In the **Confirm** window, click **Yes**.
 - d** Repeat step a through step c for the remaining logical drives on the storage partition.