

Dell PowerEdge Systems
Oracle Database on
Enterprise Linux x86_64

Storage and Network Guide

Version 1.5



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™, Optiplex™, PowerConnect™, PowerEdge™, and PowerVault™ are trademarks of Dell Inc. Broadcom™ is a trademark of Broadcom Corp. EMC®, Navisphere®, and PowerPath® are registered trademarks of EMC Corporation. Intel® is a registered trademark of Intel in the U.S. and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Red Hat® and Red Hat Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and 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 11g R2 Database	7
	Terminology Used in This Document	8
	Getting Help	8
	Dell Support.	8
	Oracle Support	8
2	Configuring Your Network	9
	Configuring the Public Network.	9
	Configuring the Private Network Using Bonding	10
	Setting Up User Equivalence	12
	Configuring ssh	12
	IP Address and Name Resolution Requirements	12
	For a Cluster Using GNS	12
	For a Cluster Using DNS	15
	Enabling IPMI	16
	Prerequisites for Enabling IPMI	17
	Configuring the Open IPMI Driver	17
	Configuring BMC Using IPMItool.	18

3	Setting Up a Fibre Channel Cluster	23
	Hardware Connections for a Fibre Channel Cluster	23
	Cabling Your Fibre Channel Storage System	25
	Configuring Direct-Attached Fibre Channel	25
	Configuring SAN-Attached Fibre Channel	26
4	Setting Up a SAS Cluster for the Dell PowerVault MD3000 and MD1000 Expansion Enclosures	29
	Setting Up the Hardware.	31
	Installing Host-Based Storage Software	32
	Verifying and Upgrading the Firmware	33
	Installing the SAS 5/E Adapter Driver	33
	Performing the Post Installation Tasks	33
	Setting Up an iSCSI Cluster for the Dell PowerVault MD32xx and MD12xx Storage Enclosures	34
	Installing Host-Based Storage Software	34
	Installing the MD Storage Software	34
	Post Installation Tasks	37
	Verifying Storage Array Discovery	37

5	Setting Up an iSCSI Cluster for the Dell PowerVault MD3000i and MD1000 Storage Enclosures	39
	Setting Up the Hardware	40
	Installing Host-based Software Needed for Storage	44
	Verifying and Upgrading the Firmware	44
	Post Installation Tasks.	44
	Setting Up an iSCSI Cluster for the Dell PowerVault MD32xxi and MD12xx Storage Enclosures.	45
	Installing Host-Based Storage Software.	45
6	Setting Up an iSCSI Cluster for the Dell EqualLogic PS Series Storage System	49
	EqualLogic Terminology	49
	Setting Up the EqualLogic iSCSI Storage System.	49
	Creating Volumes.	52
	Configuring the iSCSI Networks	54
	Configuring Host Access to Volumes.	55
7	Configuring Database Storage on the Host	65
	Verifying and Identifying Your Storage Disks.	65
	Adjusting Disk Partitions for Systems Running the Linux Operating System	68
	Example: fdisk Utility Arguments	68

Using the fdisk Utility to Adjust a Disk Partition	68
Configuring Shared Storage for Clusterware, Database, and Recovery Files in an RAC Environment.	69
Index	71

Overview

The document provides a generalized guide to configure the network and storage requirements for running the Dell|Oracle database on a system installed with the Red Hat Enterprise Linux or the Oracle Enterprise Linux operating system. This document applies to Oracle Database 11g R2 running on Red Hat Enterprise Linux 5.5 AS x86_64 or Oracle Enterprise Linux 5.5 AS x86_64. For the list of network and storage configurations supported by your system, go to dell.com/oracle and select **Dell Validated Components**.

Required Documentation for Deploying the Dell|Oracle 11g R2 Database

The Dell|Oracle Database Installation documentation set includes the following guides:

- *Dell PowerEdge Systems Oracle Database on Enterprise Linux x86_64-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.
- *Dell PowerEdge Systems Oracle Database on Enterprise Linux x86_64-Database Setup and Installation Guide*—Describes how to install and configure the Oracle database.
- *Dell PowerEdge Systems Oracle Database on Enterprise Linux x86_64-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 and iSCSI (Dell PowerVault MD3000 and Dell PowerVault MD3000i with Dell PowerVault MD1000 expansion) storage environment.
- This document uses the term Enterprise Linux that applies to both Red Hat Enterprise Linux and Oracle Enterprise Linux unless stated specifically.

Getting Help

Dell Support

- For detailed information about using your system, see the documentation that was shipped 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.
- For information on Dell enterprise training services, see dell.com/training.




NOTE: The 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 support.oracle.com.
- For information on installing and configuring Oracle, see the *Oracle Dell PowerEdge Systems Oracle Database on Enterprise Linux x86_64-Database Setup and Installation Guide* at support.dell.com/manuals.

Configuring Your Network

This section provides information about configuring the public and private cluster network.


 **NOTE:** Each node in a network requires a unique public and private Internet protocol (IP) address. An additional public IP address is required to serve as the virtual IP address for the client connections and the connection failover. Therefore, a total of three IP address are required for each node. The virtual IP address must belong to the same subnet as the public IP. All public IP addresses, including the virtual IP address, must be registered with the Domain Naming Service (DNS) and they must be routable. For more information, see "IP Address and Name Resolution Requirements" on page 12.


Depending on the number of NIC ports available, configure the interfaces as shown in Table 2-1.

Table 2-1. NIC Port Assignments

NIC Port	Ports Available
1	Public IP, virtual IP, and Single Client Access Name (SCAN) IP
2	Private IP (bonded)
3	Private IP (bonded)

Configuring the Public Network

 **NOTE:** Ensure that the public IP address is a valid and routable IP address.

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

To configure the public network 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:



NOTE: Ensure that the *Gateway* address is configured for the public network interface. If the *Gateway* address is not configured, the grid installation may fail.

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

- 3 Edit the `/etc/sysconfig/network` file, and, if necessary, replace `localhost.localdomain` with the qualified public node name. For example, the command for node 1:
`hostname=node1.domain.com`
- 4 Type `service network restart` to restart the network service.
- 5 Type `ifconfig` to verify that the IP addresses are set correctly.
- 6 To check your network configuration, ping each public IP address from a client on the LAN that is not a part of the cluster.
- 7 Connect to each node to verify that the public network is functioning. Type `ssh <public IP>` to verify that the secure shell (`ssh`) command is working.

Configuring the Private Network Using Bonding

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

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

- 1 Log in as `root`.
- 2 Add the following line to the `/etc/modprobe.conf` file:
`alias bond0 bonding`

- 3** In the `/etc/sysconfig/network-scripts/` directory, create or edit the `ifcfg-bond0` configuration file.
For example, using sample network parameters, the file appears as:

```
DEVICE=bond0
IPADDR=192.168.0.1
NETMASK=255.255.255.0
ONBOOT=yes
BONDING_OPTS="mode=6 miimon=100 max_bonds=2"
BOOTPROTO=none
```

DEVICE=bondn is the name required 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 the devices that are bonded as slaves.

- 4** For each device that is a bond member:

- a** In the directory `/etc/sysconfig/network-scripts/`, edit the `ifcfg-ethn` file:

```
DEVICE=ethn
HWADDR=<MAC ADDRESS>
ONBOOT=yes
TYPE=Ethernet
MASTER=bond0
SLAVE=yes
```

- b** Type `service network restart` and ignore all warnings.

- 5** On *each node*, type `ifconfig` to verify that the private interface is functioning.
The private IP address for the node must be assigned to the private interface `bond0`.
- 6** After all 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.
- 7** Connect to each node. Verify that the private network and the Secure Shell (`ssh`) are functioning correctly by typing:
`ssh <private IP>`

Setting Up User Equivalence

Configuring ssh

To configure ssh:

- 1 On the primary node, log in as **root**.
- 2 Go to the Grid binary folder→ **sshsetup** folder and run the following **sshUserSetup.sh** script:

```
sh sshUserSetup.sh -hosts "host1 host2" -user grid  
-advanced
```



```
sh sshUserSetup.sh -hosts "host1 host2" -user  
oracle -advanced
```

where *host1* and *host2* are the cluster node names.

After completing step 2, you can configure **ssh** on other nodes on the grid and log in as an **Oracle** user. During login you would not be prompted for a password.

IP Address and Name Resolution Requirements

For a Cluster Using GNS

To set up an Oracle 11g R2 RAC using Oracle Global Naming Server (GNS) you require:

- At least two interfaces configured on each node, one for the private IP address and one for the public IP address.
- A Domain Naming Server (DNS) running in the network for the address resolution of GNS virtual IP address.
- The DNS entry to establish DNS Lookup to direct DNS resolution of a sub-domain to the cluster.
- A DHCP service running on the public network for the cluster.
- Sufficient addresses on the DHCP to provide one IP address for each node's virtual IP, and three IP addresses for the cluster used by SCAN.

Table 2-2 describes the different interfaces, IP address settings, and the resolutions in a cluster.

Table 2-2. IP Address Settings in a Cluster Using GNS

Interface	Type	Resolution
Public	Static	/etc/hosts
Private	Static	Not Required
Node virtual IP	DHCP	GNS
GNS virtual IP	Static	DNS
SCAN virtual IP	DHCP	GNS

Configuring a DNS Server

To configure a DNS server using GNS:

- 1 Configure GNS VIP address on DNS server—In the DNS, create a name resolution entry for the GNS virtual IP address in the forward Lookup file.

For example: `gns-server IN A 192.0.2.2`

where *gns-server* is the GNS virtual IP address given during grid installation. The address that you provide must be routable and should be in public range. For example, 192.0.2.2.

- 2 Configure the GNS sub-domain delegation—In the DNS, create an entry to establish DNS Lookup that directs the DNS resolution of a GNS sub-domain to the cluster. Add the following to the DNS Lookup file:
`clusterdomain.example.com. NS gns-server.example.com.`

where *clusterdomain.example.com* is the GNS sub domain (provided during grid installation) that you delegate and *gns-server.clustername.com* resolves to GNS virtual IP address.


Configuring a DNS Client

To configure a DNS client:

- 1 Add host entries within the `/etc/hosts` file domain name. On each node, modify lines in the `/etc/hosts` file by typing:

```
127.0.0.1 localhost.localdomain localhost  
<public IP node1> <fully qualified hostname node1>  
<public hostname node1>  
  
<public IP node2> <fully qualified hostname node2>  
<public hostname node2>
```
- 2 On all nodes in the cluster, edit the `resolv.conf` file to include name server entries that are resolvable to DNS servers.
For example, edit the `/etc/resolv.conf` file as:

```
options attempts: 2  
options timeout: 1  
search clusterdomain.example.com example.com  
nameserver 192.0.2.1  
nameserver 192.0.2.2
```

 **NOTE:** The total time-out period which is a combination of options attempted and options timed out should be less than 30 seconds.

Where `192.0.2.1` is the valid DNS server address and `192.0.2.2` is the GNS Virtual IP address in your network, and `clusterdomain.example.com` is the GNS sub domain.

- 3 Verify the order of configuration. `/etc/nsswitch.conf` controls name service lookup order.

In some system configurations, the Network Information System (NIS) can cause issues with Oracle SCAN address resolution. It is recommended that you place the NIS entry at the end of the search list.

For example, `hosts: files dns nis`

For a Cluster Using DNS

To set up an Oracle 11g R2 RAC using Oracle DNS (without GNS):

- 1 At least two interfaces configured on each node, one for the private IP address and one for the public IP address.
- 2 A SCAN NAME configured on the DNS for Round Robin resolution to three addresses (recommended) or at least one address. The SCAN addresses must be on the same subnet as virtual IP addresses and public IP addresses.



NOTE: For high availability and scalability, it is recommended that you configure the SCAN to use Round Robin resolution to three IP addresses. The name for the SCAN cannot begin with a numeral. For installation to succeed, the SCAN must resolve to at least one address.

Table 2-3 describes the different interfaces, IP address settings and the resolutions in a cluster.

Table 2-3. IP Address Settings in a Cluster Using DNS

Interface	Type	Resolution
Public	Static	/etc/hosts or DNS
Private	Static	Not required
Node virtual IP	Static	/etc/hosts or DNS
SCAN virtual IP	Static	DNS

Configuring a DNS Server

To configure changes on a DNS server for an Oracle 11g R2 cluster using a DNS (without GNS):

- 1 Configure SCAN NAME resolution on DNS server. In the DNS, create a name resolution entry for the GNS virtual IP address in the forward Lookup file. For example:

```
scancluster IN A 192.0.2.1
             IN A 192.0.2.2
             IN A 192.0.2.3
```

Where *scancluster* is the SCAN NAME provided during grid installation.



NOTE: The SCAN IP address must be routable and must be in public range.

Configuring a DNS Client

To configure the changes required on the cluster nodes for name resolution:

- 1 Add host entries in the `/etc/hosts` file. On each node, modify lines in the `/etc/hosts` file by typing:

```
127.0.0.1 localhost.localdomain localhost
<public IP node1> <fully qualified hostname
node1> <public hostname node1>
<public IP node2> <fully qualified hostname
node2> <public hostname node2>
<Virtual IP node1> <fully qualified Virtual
hostname node1> <Virtual hostname node1>
<Virtual IP node2> <fully qualified Virtual
hostname node2> <Virtual hostname node2>
```

- 2 You must configure the `resolv.conf` on the nodes in the cluster to contain name server entries that are resolvable to DNS server. Provide an entry similar to the following:

```
/etc/resolv.conf:
search ns1.domainserver.com
nameserver 192.0.2.100
```

Where `192.0.2.100` is a valid DNS server address in your network and `ns1.domainserver.com` is the domain server in your network.

- 3 Verify the order configuration. `/etc/nsswitch.conf` controls the name service order. In some configurations, the NIS can cause issues with Oracle SCAN address resolution. It is recommended that you place the NIS entry at the end of the search list. For example, `hosts: files dns nis`

Enabling IPMI

Intelligent Platform Management Interface (IPMI) provides a set of common interfaces to the hardware and firmware that system administrators can use to monitor system health and manage the system. With Oracle 11g R2, Oracle Clusterware can integrate IPMI to provide failure isolation support and to ensure cluster integrity.

Prerequisites for Enabling IPMI

Each cluster node requires a Baseboard Management Controller (BMC), running a firmware compatible with IPMI version 1.5 or later and configured for remote control using LAN.



NOTE: It is recommended that you use a dedicated management network (DRAC port) for IPMI.

The Linux rpm required for `ipmitool` is `OpenIPMI-tools-2.0.16-7.el5_4.1.x86_64.rpm`.

Configuring the Open IPMI Driver

- 1 Log in as `root`.
- 2 To enable the IPMI service, type:
`chkconfig ipmi on`
- 3 To start the IPMI service, type: `service ipmi start`.
- 4 Run the command `/sbin/lsmmod | grep ipmi` to confirm that the IPMI modules are loaded. For example:

```
# /sbin/lsmmod | grep ipmi
ipmi_devintf          44753  2
ipmi_si               77453  1
ipmi_msghandler      72985  2
ipmi_devintf,ipmi_si
```

- 5 Ensure that the system recognizes the IPMI device, using the following command: `ls -l /dev/ipmi0`

If the IPMI device has been dynamically loaded, then the output should be similar to the following:

```
# ls -l /dev/ipmi0
crw----- 1 root root 252, 0 Dec 23 16:41
/dev/ipmi0
```

Configuring BMC Using IPMITool

Use the following example to configure BMC using `ipmitool` version 2.0:

- 1 Log in as `root`.
- 2 Verify that `ipmitool` is communicating with the BMC using the IPMI driver. Use the following commands to check for the device ID in the output:

```
# ipmitool bmc info
Device ID                : 32
Device Revision          : 0
Firmware Revision        : 0.20
IPMI Version              : 2.0
Manufacturer ID          : 674
Manufacturer Name        : Unknown (0x2a2)
Product ID                : 256 (0x0100)
```

If `ipmitool` is not able to communicate with BMC, see "Configuring the Open IPMI Driver" on page 17 and ensure that the IPMI driver is running.

- 3 Enable IPMI over LAN:
 - a Determine the channel number for IPMI over LAN. Beginning with channel 1, run the following command until you find the channel that displays the required LAN attributes, such as the IP address:

```
# ipmitool lan print 1
. . .
IP Address Source : 0x01
IP Address : 140.87.155.89
. . .
```

- b Enable LAN access for the channel found. For example, where the channel is 1:

```
# ipmitool lan set 1 access on
```

4 Configure IP address settings for IPMI using one of the following procedures:

- Using dynamic IP addressing—Dynamic IP addressing is the default assumed by Oracle Universal Installer. It is recommended that you select this option so that nodes can be added or removed from the cluster more easily, as address settings can be assigned automatically. For example, if the channel is 1, then enter the following command to enable DHCP and set the channel:

```
# ipmitool lan set 1 ipsrc dhcp
```



NOTE: Use of DHCP requires a DHCP server on the subnet.

- Using static IP addressing—If the BMC shares a network connection with the operating system, then the IP address must be on the same subnet. You must set not only the IP address, but also the proper values for netmask, and the default gateway. For example, assuming the channel is 1:

```
# ipmitool lan set 1 ipaddr 192.168.0.55
```

```
# ipmitool lan set 1 netmask 255.255.255.0
```

```
# ipmitool lan set 1 defgw ipaddr 192.168.0.1
```

The address specified in the example(192.168.0.55) is associated only with the BMC, and does not respond to normal pings.

5 Establish an administration account with a user name and password, using the following procedure (assuming the channel is 1):

- a** Configure BMC to authenticate password for **Admin** access over LAN. For example:

```
# ipmitool lan set 1 auth ADMIN MD5,PASSWORD
```

- b** List the account slots on the BMC, and identify an unused slot (a User ID with an empty user name field). For example:

```
# ipmitool channel getaccess 1
```

```
. . .
```

```
User ID : 4
```

```
User Name :
```

```
Fixed Name : No
```

```
Access Available : call-in / callback
Link Authentication : disabled
IPMI Messaging : disabled
Privilege Level : NO ACCESS
. . .
```

- c** Assign the desired administrator user name and password and enable messaging for the identified slot. Also set the privilege level for that slot when accessed over LAN (channel 1) to ADMIN (level 4). For example:

```
# ipmitool user set name 4 bmcuser
# ipmitool user set password 4 password
# ipmitool user enable 4
# ipmitool channel setaccess 1 4 privilege=4
# ipmitool channel setaccess 1 4 link=on
# ipmitool channel setaccess 1 4 ipmi=on
```

Where *bmcuser* is the administrative user name, and *password* is the password.

- d** Verify the setup using the command `lan print 1`. The following output must be displayed:

```
# ipmitool lan print 1
Set in Progress : Set Complete
Auth Type Support : NONE MD2 MD5 PASSWORD
Auth Type Enable : Callback : MD2 MD5
: User : MD2 MD5
: Operator : MD2 MD5
: Admin : MD5 PASSWORD
: OEM : MD2 MD5
```

```
IP Address Source : DHCP Address [or Static
Address]
IP Address : 192.168.0.55
Subnet Mask : 255.255.255.0
MAC Address : 00:14:22:23:fa:f9
SNMP Community String : public
IP Header : TTL=0x40 Flags=0x40 Precedence=...
Default Gateway IP : 192.168.0.1
Default Gateway MAC : 00:00:00:00:00:00
.
.
.
# ipmitool channel getaccess 1 4
Maximum User IDs : 10
Enabled User IDs : 2
User ID : 4
User Name : bmcuser
Fixed Name : No
Access Available : call-in / callback
Link Authentication : enabled
IPMI Messaging : enabled
Privilege Level : ADMINISTRATOR
```

- 6 Verify that BMC is accessible and controllable from a remote node in your cluster using the **bmc info** command.
For example, if *node2-ipmi* is the network host name assigned to BMC for node2, then to verify the BMC on node2 from node1, enter the following command on node1:

```
$ ipmitool -H node2-ipmi -U bmcuser -P password  
bmc info
```

Where *bmcuser* is the administrator account and *password* is the password.

If BMC is correctly configured, then the information about BMC on the remote node is displayed. If you see an error message, such as `Error: Unable to establish LAN session`, check the BMC configuration on the remote node.

Setting Up a Fibre Channel Cluster

⚠ 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 helps you to verify the hardware connections, and the hardware and software configurations of the the Fibre Channel cluster.

Hardware Connections for a Fibre Channel Cluster

Figure 3-1 illustrates the hardware connections of a Fibre Channel cluster.

Figure 3-1. Hardware Connections for a Fibre Channel Cluster




Table 3-1 lists the Fibre Channel hardware connections depicted in Figure 3-1 and summarises the cluster connections.

Table 3-1. Fibre Channel Hardware Interconnections

Cluster Component	Connections
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 fibre optic cable from HBA 0 to Fibre Channel switch 0 • One fibre optic cable from HBA 1 to Fibre Channel switch 1
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 fibre optic cable connections to each Fibre Channel switch. For example, for a four-port configuration: <ul style="list-style-type: none"> – One fibre optic cable from SPA port 0 to Fibre Channel switch 0 – One fibre optic cable from SPA port 1 to Fibre Channel switch 1 – One fibre optic cable from SPB port 0 to Fibre Channel switch 1 – One fibre 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 fibre optic cable connections to the Dell EMC Fibre Channel storage system • One fibre 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 Fibre Channel Storage System

Depending on your requirements, you can configure your Oracle database Fibre Channel cluster storage system in one of the following configurations:

- Direct-attached Fibre Channel (see Figure 3-2)
- Four-port SAN-attached Fibre Channel (see Figure 3-3)


Configuring Direct-Attached Fibre Channel

To configure the nodes in a direct-attached Fibre Channel configuration:

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

Figure 3-2 illustrates the cable connections in a direct-attached Fibre Channel cluster.

Figure 3-2. Cabling in a Direct-Attached Fibre Channel Cluster




Configuring SAN-Attached Fibre Channel

To configure your nodes in a four-port SAN-attached configuration:

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

Figure 3-3 illustrates the cable connections in a SAN-attached Fibre Channel cluster.

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



Setting Up a SAS Cluster for the Dell PowerVault MD3000 and 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.

To configure the Dell PowerEdge systems and the Dell PowerVault MD3000 and MD1000 storage enclosure to function in a Oracle Real Application Cluster (RAC) environment:

- 1 Verify the hardware and software configurations using Figure 4-1, Table 4-1, and Figure 4-2.
- 2 Follow the steps in "Setting Up the Hardware" on page 31.

⚠ CAUTION: Before you upgrade the firmware on your PowerVault MD3000 storage array, ensure that you backup your data.

Figure 4-1. Cabling the Serial-Attached SCSI (SAS) Cluster and the Dell PowerVault MD3000 Storage Enclosure




Table 4-1. 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 PowerVault MD3000 system node through the SAS 5/E

Table 4-1. SAS Cluster Hardware Interconnections (continued)

Cluster Component	Connections
PowerVault MD3000 storage enclosure	<ul style="list-style-type: none">• Two CAT 5e/6 cables connected to LAN (one from each storage processor module)• Two SAS connections to each PowerEdge system node through the SAS 5/E cables <p>NOTE: For more information on the PowerVault MD3000 storage enclosure interconnection, see "Setting Up the Hardware" on page 31.</p>
PowerVault MD1000 storage enclosure (optional)	Additional SAS cable connections as required for the PowerVault MD1000 expansion enclosure

Setting Up the Hardware

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:

- 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 storage enclosure.
- 2** Connect one SAS cable from the other port of the SAS controller of node 1 to the **In-0** port of RAID controller 1 in the PowerVault MD3000 storage 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 storage enclosure.
- 4** Connect one SAS cable from the other port of SAS controller of node 2 to the **In-1** port of RAID controller 1 in the PowerVault MD3000 storage enclosure.
- 5** If applicable, connect two SAS cables from the two PowerVault MD3000 storage enclosures **Out** ports to the two **In** ports of the first PowerVault MD1000 expansion enclosure.

- If applicable, connect two SAS cables from the two PowerVault MD1000 storage enclosures out ports to the **In-0** ports of the second PowerVault MD1000 expansion enclosure.



 **NOTE:** For information on configuring the PowerVault MD1000 expansion enclosure, see the PowerVault MD1000 storage system documentation at support.dell.com/manuals.

Figure 4-2. Cabling the Direct-Attached SAS Cluster



Installing Host-Based Storage Software

To install the necessary host-based storage software for the PowerVault MD3000 storage enclosure, use the Dell PowerVault Resource media that shipped with the system. To install the Modular Disk Storage Manager (MDSM) software on the master node and the Multi-Path Input/Output (MPIO) software on the remaining nodes, follow the procedures in the PowerVault MD3000 storage enclosure documentation.

Verifying and Upgrading the Firmware

To verify and upgrade the firmware:

- 1 Discover the direct-attached storage of the host system using the MDSM software that is installed on the host system.
- 2 Verify that the firmware for the following storage components is at the minimum required version:
 - RAID controller firmware
 - PowerVault MD3000 storage system firmware
 - PowerVault MD1000 expansion enclosure firmware



NOTE: For the minimum firmware version requirements, select the **Dell Validated Components** link at dell.com/oracle.

Installing the SAS 5/E Adapter Driver

Install the SAS 5/E driver from the *PowerVault MD3000 Resource* media.



NOTE: Ensure that the version of the SAS 5/E driver is the same or newer than the version listed in the Solution Deliverable List (SDL) that can be found for each Dell-validated component at dell.com/oracle.

For assistance in installing drivers on both nodes of the cluster, see the documentation that came with the PowerVault MD3000 storage enclosure and the SAS HBAs.

Performing the Post Installation Tasks

After installing the drivers and the software, perform the post installation tasks listed in the *PowerVault MD3000 Installation Guide*. Create the environment shown in the *Dell PowerEdge Systems Oracle Database on Enterprise Linux x86_64 Operating System and Hardware Installation*, and the *Linux Guide*. Both documents are available at support.dell.com/manuals.



NOTE: It is recommended that you configure the disks for the LUNs in a RAID 10 configuration.

Setting Up an iSCSI Cluster for the Dell PowerVault MD32xx and MD12xx Storage Enclosures

Setting Up the Hardware

For assistance in setting up your PowerVault MD32xx and PowerVault MD12xx expansion enclosure, please see the PowerVault documentation at support.dell.com/manuals.

Installing Host-Based Storage Software

Software Installation

The PowerVault MD3200 Series resource media contains software and drivers for both Linux and Microsoft Windows operating systems.

The root of the media contains a `readme.txt` file covering changes to the software, updates, fixes, patches, and other important data applicable to both Linux and Windows operating systems. The `readme.txt` file also specifies requirements for accessing documentation, information regarding versions of the software on the media, and system requirements for running the software.

For more information on supported hardware and software for Dell PowerVault systems, see the Support Matrix located at support.dell.com/manuals.



NOTE: It is recommended that you install all the latest updates available at support.dell.com.

Installing the MD Storage Software

The PowerVault MD3200 Series storage software installer provides features that include the core software and providers. The core software feature includes the host-based storage agent, multipath driver, and MD Storage Manager application used to configure, manage, and monitor the storage array solution.

The providers feature include providers for the Microsoft Virtual Disk Service (VDS) and Microsoft Volume Shadow-Copy Service (VSS) framework.



NOTE: For more information about the Microsoft VDS and Microsoft VSS providers, see the Owner's Manual.



NOTE: To install the software on a Windows or Linux system, you must have administrative or root privileges.

The PowerVault MD3200 Series resource media offers the following three installation methods:

- **Graphical Installation (Recommended)**—This is the recommended installation procedure for most users. The installer presents a graphical wizard-driven interface that allows customization of which components are installed.
- **Console Installation**—This installation procedure is useful for Linux users who do not desire to install an X-Window environment on their supported Linux platform.
- **Silent Installation**—This installation procedure is useful for users who prefer to create scripted installations.

Graphical Installation (Recommended)

To complete graphical installation:

- 1 Close all other programs before installing any new software.
- 2 Insert the resource media.
- 3 Select **Install MD32xx Storage Software**.



NOTE: If the host server is running Windows Server Core version or if the launcher does not automatically run, navigate to the root of the media and run **md_launcher.exe** for Windows or the autorun script for Linux.

- 4 Install the MD Storage software. Select one of the following installation options:
 - **Full**—This package installs core software, providers, and utilities. It includes the necessary host-based storage agent, multipath driver, MD Storage Manager, providers, and optional utilities.
 - **Host Only**—This package includes the host context agent, multipath drivers, and optional utilities required to configure the host.
 - **Management Station**—This package includes the MD Storage.
 - **Manager, providers, and optional utilities.**
 - **Custom**—This option allows you to select specific components.
- 5 Reboot each host server.

Console Installation



NOTE: Console installation only applies to Linux systems that are not running a graphical environment.

The autorun script in the root of the resource media detects when there is no graphical environment running and automatically starts the installer in a text-based mode. This mode provides the same options as graphical installation.

Silent Installation

This option allows you to install the software in an unattended mode. To run the silent installation in a Windows system:

- 1 To begin the installation with all the default options, navigate to the windows folder on the resource media and run the following command:

```
md32xx_install -i silent
```
- 2 To specify which features are to be installed silently, modify the **custom_silent.properties** file located in the same directory and copy it to a writable location on the host server. This file contains a list of supported features and their descriptions. To utilize the modified **custom_silent.properties** file, run the command:

```
md32xx_install -f  
C:\path\to\custom_silent.properties
```

To run the silent installation in a Linux system, perform the following steps:

- 1 To begin the installation with all the default options, navigate to the linux folder on the resource media and run the following command:

```
./md32xx_install.bin -i silent
```
- 2 To specify which features are to be installed silently, you must modify the **custom_silent.properties** file located in the same directory and copy it to a writable location on the host server. This file contains a list of supported features and their descriptions. To utilize the modified **custom_silent.properties** file, run the command:

```
./md32xx_install.bin -f  
/path/to/custom_silent.properties
```

Post Installation Tasks

Before using the PowerVault MD3200 Series storage array for the first time, complete a number of initial configuration tasks in the order shown. These tasks are performed using the MD Storage Manager (MDSM) software.

To complete the post installation tasks:

- 1 For out-of-band management, you must set the network configuration for each RAID controller module, including its Internet Protocol (IP) address, subnetwork mask (subnet mask), and gateway.



NOTE: You can set the network configuration using a DHCP server.

- 2 Start MDSM.

- On Microsoft Windows operating systems, click **Start**→ **Programs**→ **Dell**→ **MD Storage Manager**→ **Modular Disk Storage Manager Client**.
- On Linux operating systems, click the **MD Storage Manager** desktop icon. You can also launch MD Storage Manager in Linux by running the following commands at the command prompt:

```
cd/opt/dell/mdstoragesoftware/mdstoragemanager/client./SMclient
```

- 3 In the **Enterprise Management Window (EMW)**, go to **Setup**→ **Add Storage Array** and start the automatic discovery of systems.



NOTE: Only storage arrays configured with IP address on the same subnet as the management stations will be detected.



NOTE: It may take several minutes for the discovery process to complete. Closing the Automatic Discovery Status window before this stops the discovery process.

After discovery is complete, a confirmation screen appears.

- 4 Click **Close** to close the screen.

Verifying Storage Array Discovery

The Automatic Discovery option automatically discovers both in-band and out-of-band arrays and adds the storage arrays to the management domain. If the out-of-band or in-band connections shown on the summary page are not accurate, complete the following steps. For more information, see the *online help*.

To verify storage array discovery:

- 1 Check the hardware and connections for possible problems. For specific procedures on troubleshooting interface problems, see the *Owner's Manual*.
- 2 Verify that the array is on the local subnetwork. If it is not, click the **New** link to manually add it.
- 3 Verify that the status of each storage array is **Optimal**. If any array shows an Unresponsive status, complete one of the following steps:
 - a To remove the array from the management domain, highlight the array and then click **Remove** link.
 - b To add an array, click **New** link. In Linux, click the **Manual** button and click **OK**.
- 4 If the array still shows an Unresponsive status, contact Dell.

Initial Setup Tasks

To complete initial setup tasks:

- 1 The name of the first storage array found appears in the devices tab of the EMW. To see a list of all storage arrays found on the local network, expand the discovered storage arrays tab in the device pane of the EMW.
- 2 The default name for a newly installed MD3200 Series storage array is **Unnamed**. If another name appears in MDSM, click the down arrow next to the name and click **Unnamed** in the drop-down list.
- 3 Double-click on the discovered array to launch the array management window and make changes to the configuration of the array.
- 4 Click **Perform Initial Setup Tasks** to see an outline and links to the remaining post-installation tasks. The array is now ready to be configured.



NOTE: Before configuring the storage array, check the status icons on the Device pane to ensure that the enclosures in the storage array are in an Optimal status. For more information, see the *Owner's Manual* at support.dell.com/manuals.

Enabling Premium Features (Optional)

If you ordered premium features, follow the instructions on the Premium Features Card to install the additional features.

Setting Up an iSCSI Cluster for the Dell PowerVault MD3000i and MD1000 Storage 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 for configuring your Dell PowerEdge system and the PowerVault MD3000i and MD1000 storage enclosures to function in a Oracle Real Application Cluster (RAC) environment.

Verify the hardware connections, and the hardware and software configurations using the Supported Configuration section of the *Dell PowerVault MD3000i Support Matrix* at support.dell.com/manuals.

Table 5-1. iSCSI 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) • One CAT 5e/6 cable from iSCSI Gigabit NIC to Gigabit Ethernet switch (iSCSI network)


Table 5-1. iSCSI Hardware Interconnections (continued)

Cluster Component	Connections
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 the PowerVault MD3000i storage enclosure, see the PowerVault MD3000i documentation at support.dell.com/manuals.</p>
PowerVault MD1000 storage expansion enclosure (optional)	Additional SAS cable connections, as required for the PowerVault MD1000 expansion enclosure

Setting Up the Hardware

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

Figure 5-1. Cabling iSCSI Direct-Attached Clusters



To configure your nodes in a direct-attached configuration, see Figure 5-1:


- 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 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 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 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 PowerVault MD3000i storage enclosure.
- 5** If applicable, connect two SAS cables from the two PowerVault MD3000 storage enclosures out ports to the two In ports of the first PowerVault MD1000 expansion enclosure.
- 6** If applicable, connect two SAS cables from the two PowerVault MD1000 storage enclosures out ports to the **In-0** ports of the second MD1000 expansion enclosure.



NOTE: For information on configuring the PowerVault MD1000 expansion enclosure, see the PowerVault MD3000i storage system documentation at support.dell.com/manuals.

Switched iSCSI clusters can support up to eight nodes.

Figure 5-2. Cabling iSCSI Switched Clusters



To configure your nodes in a switched configuration, see Figure 5-2:

- 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 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 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 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 PowerVault MD3000i storage enclosure.
- 9 If applicable, connect two SAS cables from the two PowerVault MD3000i storage enclosures out ports to the two In ports of the first PowerVault MD1000 expansion enclosure.
- 10 If applicable, connect two SAS cables from the two PowerVault MD3000 storage enclosures out ports to the **In-0** ports of the second PowerVault MD1000 expansion enclosure.



NOTE: For information on configuring the PowerVault MD1000 expansion enclosure, see the PowerVault MD3000i storage system documentation at support.dell.com/manuals. It is recommended to use a separate network for the iSCSI storage infrastructure. If a separate network cannot be dedicated for iSCSI, assign the storage function to a separate virtual local area network (VLAN); this action creates independent logical networks within a physical network.

Installing Host-based Software Needed for Storage

To install the necessary host-based storage software for the PowerVault MD3000i storage system, use the *Dell PowerVault Resource* media that came with your PowerVault MD3000i storage system. Follow the procedures in the PowerVault MD3000i storage enclosure documentation at support.dell.com/manuals, to install the Modular Disk Storage Manager Software (MDSM) on the Master node and the Multi-Path (MPIO) software on the remaining nodes.

Verifying and Upgrading the Firmware

To verify and update the firmware:

- 1 Discover the direct-attached storage of the host system using the MDSM software that is installed on the host system.
- 2 Verify that the firmware for the following storage components is at the minimum required version.
 - MD3000i storage system firmware
 - MD1000 expansion enclosure firmware



NOTE: For the minimum firmware version requirements, select the **Dell Validated Components** link at dell.com/oracle.

Post Installation Tasks

After installing the drivers and the software, perform the post installation tasks listed in the *PowerVault MD3000i Installation Guide* at support.dell.com/manuals, to create the environment shown in Table 5-1 on page 39.

Setting Up an iSCSI Cluster for the Dell PowerVault MD32xxi and MD12xx Storage Enclosures

Setting up the Hardware

For assistance in setting up or installing your PowerVault MD32xxi and PowerVault MD12xx expansion enclosure, please refer to the PowerVault documentation that can be found at support.dell.com/manuals.

Installing Host-Based Storage Software

The PowerVault MD3200i Series storage software installer provides features that include the core software, providers, and optional utilities. The core software feature includes the host-based storage agent, multipath driver, and MD Storage Manager (MDSM) application used to configure, manage, and monitor the storage array solution. The providers feature includes providers for the Microsoft Virtual Disk Service (VDS) and Microsoft Volume Shadow-Copy Service (VSS) framework. The Modular Disk Configuration Utility (MDCU) is an optional utility that provides a consolidated approach for configuring the management ports, iSCSI host ports, and creating sessions for the iSCSI Modular Disk storage arrays. It is recommended that you install and use MDCU to configure iSCSI on each host connected to the storage array.



NOTE: For more information about the Microsoft VDS and Microsoft VSS providers, see the Owner's Manual. To install the software on a Windows or Linux system, you must have administrative or root privileges.



NOTE: If Dynamic Host Configuration Protocol (DHCP) is not used, initial configuration of the management station must be performed on the same physical subnet as the storage array. Additionally, during initial configuration, at least one network adapter must be configured on the same IP subnet as the storage array's default management port (192.168.128.101 or 192.168.128.102). After initial configuration, the management ports are configured using MDSM and the management station's IP address can be changed back to the previous settings.

The PowerVault MD3200i Series resource media provided offers the following three installation methods:

- Graphical Installation (Recommended)—This is the recommended installation procedure for most users. The installer presents a graphical wizard-driven interface that allows customization of which components are installed.

- Console Installation—This installation procedure is useful for Linux users that do not desire to install an X-Window environment on their supported Linux platform.
- Silent Installation—This installation procedure is useful for users that prefer to create scripted installations.

Graphical Installation (Recommended)

To complete graphical installation:

- 1 Close all other programs before installing any new software.
- 2 Insert the resource media.
- 3 Select Install MD32xxi Storage Software.



NOTE: If the host server is running Windows Server Core version or if the launcher does not automatically run, navigate to the root of the media and run md_launcher.exe for Windows or the autorun script for Linux.

- 4 Install the MD Storage software. Select one of the following installation options:
 - Full—This package installs core software, providers, and utilities. It includes the necessary host-based storage agent, multipath driver, MD Storage Manager, providers, and optional utilities.
 - Host Only—This package includes the host context agent, multipath drivers, and optional utilities required to configure the host.
 - Management Station—This package includes the MD Storage Manager, providers, and optional utilities.
 - Custom—This option allows you to select specific components.
- 5 Reboot each host server.

Console Installation

NOTE: Console installation only applies to Linux systems that are not running a graphical environment.

The autorun script in the root of the resource media detects when there is no graphical environment running and automatically starts the installer in a text-based mode. This mode provides the same options as graphical installation with the exception of the MDCU specific options. The MDCU requires a graphical environment to operate.



NOTE: The console mode installer provides the option to install the MDCU. However a graphical environment is required to utilize the MDCU.

Silent Installation

This option allows you to install the software in an unattended mode. To run the silent installation in a Windows system, perform the following steps:

- 1 To begin the installation with all the default options, navigate to the windows folder on the resource media and run the following cmd:
`md32xxi_install -i silent`
- 2 To specify which features are to be installed silently, you must modify the `custom_silent.properties` file located in the same directory and copy it to a writable location on the host server. This file contains a list of supported features and their descriptions. To utilize the modified `custom_silent.properties` file, run the cmd:

```
md32xxi_install -f <path>/custom_silent.properties
```

Enabling Premium Features (Optional)

If you ordered premium features, follow the instructions on the Premium Features Card to install the additional features.

Setting Up an iSCSI Cluster for the Dell|EqualLogic PS Series Storage System

⚠ 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.

EqualLogic Terminology

The EqualLogic PS series storage array includes storage virtualization technology. To better understand how these arrays operate, it is helpful to be familiar with some of the terminology 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 systems 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

Setting Up the EqualLogic iSCSI Storage System

Host systems can be attached to the EqualLogic PS6000XV iSCSI array through an IP SAN industry-standard Gigabit Ethernet switch. Figure 6-1 shows the recommended network configuration for a dual control module EqualLogic PS6000XV array. This configuration includes two Dell PowerConnect 6200 series Gigabit Ethernet switches to provide the highest network availability and the maximum network bandwidth. If you are using the EqualLogic 10GbE iSCSI arrays, please refer to EqualLogic's documentation at support.equallogic.com.

NOTE: It is recommended to use two Gigabit Ethernet switches. 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. In such a configuration, there must be multiple ports with link aggregation providing the inter-switch, or trunk connection. From each of the control modules, it is recommended to connect one Gigabit interface to one Ethernet switch, and to connect the other two Gigabit interfaces to the other Ethernet switch.

Figure 6-1. Recommended Network Configuration





Figure 6-2 is an architectural overview of a sample Oracle Real Application Cluster (RAC) configuration with three Dell|EqualLogic PS6000XV arrays. The Dell|EqualLogic PS6000XV storage arrays provide the physical storage capacity for the RAC database. The group **oracle-group** includes three Dell|EqualLogic PS6000XV members: **oracle-member01**, **oracle-member02**, and **oracle-member03**. When a member is initialized, it can be configured with RAID 10, RAID 5, or RAID 50.

Figure 6-2. Configuring a Sample Oracle RAC With Three PS6000XV Arrays



NOTE: For more information on how to initialize an EqualLogic array, see the *Dell EqualLogic User's Guide* at support.dell.com/manuals.

An EqualLogic 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 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 6-2, pools are organized by member RAID levels:

- One pool with the name RAID 10 consists of RAID10 members.
- One pool with the name RAID 5 consists of RAID5 members.

Creating Volumes

Before storing data, configure the physical disks of the EqualLogic PS5000XV storage array 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. The volume 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 to data availability. In addition, based on the overall workload of the storage hardware resources within the pool, automatic data placement and automatic load balancing occurs within a pool.

Table 6-1 shows a sample volume configuration.

Table 6-1. Configuring Volumes for Oracle RAC

Number of Volumes	Minimum Size	RAID	Number of Partitions	Used For	Operating System Mapping
Five	1024 MB	10	One	Voting Disk, Oracle Cluster Registry (OCR)	ASM disk group for Cluster Registry (OCR) and Voting disk
One	Larger than the size of your database	10	One	Data	ASM disk group DATABASEDG
One	Minimum twice the size of your second area volume(s)	5	One	Flash Recovery Area	ASM disk group FLASHBACKDG
One (Optional)	10 GB	10	One	Shared DB Home	ASM disk group for Shared DB Home

Create volumes in the Dell |EqualLogic PS5000XV array and create access list to allow all host iSCSI network interfaces to access the volumes. For example, the following volumes are created:

- mdi-ocr
- mdi-data1
- mdi-fra1

Configuring the iSCSI Networks

It is recommended to configure the host network interfaces for iSCSI traffic to use **Flow Control** and **Jumbo Frame** for optimal performance. Use the `ethtool` utility to configure **Flow Control**.

Use the following command to check for **Flow Control (RX/TX Pause)** on the interfaces: `# ethtool -a <interface>`

For example:

```
# ethtool -a eth2
Pause parameters for eth2:
Autonegotiate:      on
RX:                 on
TX:                 on
```

This example shows that **Flow Control** is already turned on. If not, use the following command to turn on **Flow Control**:

```
# ethtool -A <interface> rx on tx on
```

Jumbo Frame is configured in the `/etc/sysconfig/network-scripts/ifcfg-<interface>` scripts, by adding the `MTU="<mtu-value>"` parameter.

The example below shows the **MTU** set to **9000**.

```
# cat /etc/sysconfig/network-scripts/ifcfg-eth2
DEVICE=eth2
HWADDR=00:15:17:80:43:50
ONBOOT=yes
TYPE=Ethernet
BOOTPROTO=none
IPADDR=10.16.7.125
NETMASK=255.255.255.0
USERCTL=no
MTU=9000
```

Verify the **Jumbo Frame** setting using the `ifconfig` command:

```
$ ifconfig eth2
eth2      Link encap:Ethernet  HWaddr
00:15:17:80:43:50
          inet addr:10.16.7.125  Bcast:10.16.7.255
Mask:255.255.255.0
          inet6 addr: fe80::215:17ff:fe80:4350/64
Scope:Link
```

```
UP BROADCAST RUNNING MULTICAST  MTU:9000
Metric:1
RX packets:3348411 errors:0 dropped:0
overruns:0 frame:0
TX packets:2703578 errors:0 dropped:0
overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:10647052076 (9.9 GiB) TX
bytes:11209177325 (10.4 GiB)
Memory:d5ee0000-d5f00000
```

Configuring Host Access to Volumes

This section provides information about configuring the host access to iSCSI volumes using the `iscsiadm` tool. The `iscsiadm` tool is the open-iSCSI administration utility.

- 1 Log in to the system as `root`. Verify the open-iSCSI initiator software is installed on all host systems:

```
rpm -qa|grep -i iscsi-initiator
```

If the open-iSCSI initiator RPM is installed, the following output is displayed:

```
iscsi-initiator-utils-6.2.0.871-0.16.el5_4.1
```

If the output is not displayed, install the open-iSCSI initiator RPM.

- 2 Start up the iSCSI service.

```
service iscsi start
```
- 3 Enable the iSCSI service start-up at boot time.

```
chkconfig --add iscsi
chkconfig iscsi on
chkconfig --list iscsi
```
- 4 Obtain the hardware address of each network interface on the host used for iSCSI traffic.

```
grep -i hwaddr /etc/sysconfig/network-
scripts/ifcfg-ethn,
```

where `n` is the network interface number.

- 5 Create an interface for each network interface on the host used for iSCSI traffic.

```
iscsiadm -m iface -I iface_name --op=new,
```

where, **iface_name** is the name assigned to the interface.

```
iscsiadm -m iface -I iface_name --op=update -n  
iface.hwaddress -v hardware_address
```

where, **hardware_address** is the hardware address of the interface obtained in step 4.

For example, the following commands create an interface named **eth0-iface** for the **eth0** interface whose hardware address is **00:18:8B:4E:E6:CC**:

```
# iscsiadm -m iface -I eth0-iface --op=new
```

The new interface **eth0-iface** is added.

```
# iscsiadm -m iface -I eth0-iface --op=update -n  
iface.hwaddress -v 00:18:8B:4E:E6:CC  
eth0-iface updated
```

- 6 Verify the interfaces are created and associated properly:

```
iscsiadm -m iface
```

- 7 Modify the CHAP information in **/etc/iscsi/iscsid.conf** on the host.

```
node.session.auth.authmethod = CHAP
```

```
node.session.auth.username = username
```

```
node.session.auth.password = password
```

```
discovery.sendtargets.auth.authmethod = CHAP
```

```
discovery.sendtargets.auth.username = username
```

```
discovery.sendtargets.auth.password = password
```

where, **username** is the CHAP username defined in the EqualLogic storage, and **password** is the CHAP password defined in the EqualLogic storage.

- 8 Modify the replacement timeout value within the **/etc/iscsi/iscsid.conf** file. Change the value from the default 120 to 30

```
node.session.timeo.replacement_timeout = 30
```

- 9 Restart the iSCSI service for the new configuration to take effect.

```
service iscsi stop
```

```
service iscsi start
```

- 10 Discover the targets from all **ifaces** created in step 5.

```
iscsiadm -m discovery -t st -p group_ip_address --  
interface=iface_name1 --interface=iface_name2 --
```


interface=iface_name3 --interface=iface_name4, where, **group_ip_address** is the IP address of the EqualLogic storage group, **iface_name1**, **iface_name2**, **iface_name3**, **iface_name4**, etc, are the network interfaces (as defined in step 5) on the host that is used for iSCSI traffic.

For example, the following command discovers four volumes at group IP address **10.16.7.100**, from a host with two interfaces named as **eth0-iface** and **eth1-iface**:

```
# iscsiadm -m discovery -t st -p 10.16.7.100 --
interface=eth0-iface --interface=eth1-iface
10.16.7.100:3260,1 iqn.2001-05.com.equallogic:0-
8a0906-90ee59d02-e26f999767b4942e-mdi-ocr
10.16.7.100:3260,1 iqn.2001-05.com.equallogic:0-
8a0906-90ee59d02-e26f999767b4942e-mdi-ocr
10.16.7.100:3260,1 iqn.2001-05.com.equallogic:0-
8a0906-93ee59d02-674f999767d4942e-mdi-data1
10.16.7.100:3260,1 iqn.2001-05.com.equallogic:0-
8a0906-93ee59d02-674f999767d4942e-mdi-data1
10.16.7.100:3260,1 iqn.2001-05.com.equallogic:0-
8a0906-95ce59d02-2e0f999767f4942e-mdi-data2
10.16.7.100:3260,1 iqn.2001-05.com.equallogic:0-
8a0906-95ce59d02-2e0f999767f4942e-mdi-data2
10.16.7.100:3260,1 iqn.2001-05.com.equallogic:0-
8a0906-97be59d02-d7ef99976814942e-mdi-fra1
10.16.7.100:3260,1 iqn.2001-05.com.equallogic:0-
8a0906-97be59d02-d7ef99976814942e-mdi-fra1
```

11 Verify all volumes are discovered from all the **ifaces** on the host:

```
iscsiadm -m discovery --print=1
```

For example:

```
# iscsiadm -m discovery --print=1
```

SENDTARGETS:

```
DiscoveryAddress: 10.16.7.100,3260
Target: iqn.2001-05.com.equallogic:0-8a0906-
90ee59d02-e26f999767b4942e-mdi-ocr
Portal: 10.16.7.100:3260,1
Iface Name: eth0-iface
```

```
Iface Name: eth1-iface
```

```
Target: iqn.2001-05.com.equallogic:0-8a0906-93ee59d02-674f999767d4942e-mdi-data1
```

```
Portal: 10.16.7.100:3260,1
```

```
Iface Name: eth0-iface
```

```
Iface Name: eth1-iface
```

```
Target: iqn.2001-05.com.equallogic:0-8a0906-95ce59d02-2e0f999767f4942e-mdi-data2
```

```
Portal: 10.16.7.100:3260,1
```

```
Iface Name: eth0-iface
```

```
Iface Name: eth1-iface
```

```
Target: iqn.2001-05.com.equallogic:0-8a0906-97be59d02-d7ef99976814942e-mdi-fra1
```

```
Portal: 10.16.7.100:3260,1
```

```
Iface Name: eth0-iface
```

```
Iface Name: eth1-iface
```

```
iSNS:
```

```
No targets found.
```

```
STATIC:
```

```
No targets found.
```

12 Log in to all targets (volumes) from each interface you created in step 5:

```
iscsiadm -m node -p group_ip_address --interface iface_name --login,
```

where, **group_ip_address** is the IP address of the EqualLogic storage group, and **iface_name** is the network interface (as defined in step 5) on the host that is used for iSCSI traffic.

The following example logs into three volumes from each of the two **ifaces** (**eth0-iface** and **eth1-iface**) on a host.

```
# iscsiadm -m node -p 10.16.7.100 --interface eth0-iface --login
```

```
Logging in to [iface: eth0-iface, target: iqn.2001-05.com.equallogic:0-8a0906-90ee59d02-e26f999767b4942e-mdi-ocr, portal: 10.16.7.100,3260]
```

```
Logging in to [iface: eth0-iface, target:
iqn.2001-05.com.equallogic:0-8a0906-95ce59d02-
2e0f999767f4942e-mdi-data2, portal:
10.16.7.100,3260]
```

```
Logging in to [iface: eth0-iface, target:
iqn.2001-05.com.equallogic:0-8a0906-93ee59d02-
674f999767d4942e-mdi-data1, portal:
10.16.7.100,3260]
```

```
Logging in to [iface: eth0-iface, target:
iqn.2001-05.com.equallogic:0-8a0906-97be59d02-
d7ef99976814942e-mdi-fral, portal:
10.16.7.100,3260]
```

```
Login to [iface: eth0-iface, target: iqn.2001-
05.com.equallogic:0-8a0906-90ee59d02-
e26f999767b4942e-mdi-ocr, portal:
10.16.7.100,3260]: successful
```

```
Login to [iface: eth0-iface, target: iqn.2001-
05.com.equallogic:0-8a0906-95ce59d02-
2e0f999767f4942e-mdi-data2, portal:
10.16.7.100,3260]: successful
```

```
Login to [iface: eth0-iface, target: iqn.2001-
05.com.equallogic:0-8a0906-93ee59d02-
674f999767d4942e-mdi-data1, portal:
10.16.7.100,3260]: successful
```

```
Login to [iface: eth0-iface, target: iqn.2001-
05.com.equallogic:0-8a0906-97be59d02-
d7ef99976814942e-mdi-fral, portal:
10.16.7.100,3260]: successful
```

```
# iscsiadm -m node -p 10.16.7.100 --interface
eth1-iface --login
```

```
Logging in to [iface: eth1-iface, target:
iqn.2001-05.com.equallogic:0-8a0906-90ee59d02-
e26f999767b4942e-mdi-ocr, portal:
10.16.7.100,3260]
```

```
Logging in to [iface: eth1-iface, target:
iqn.2001-05.com.equallogic:0-8a0906-95ce59d02-
2e0f999767f4942e-mdi-data2, portal:
10.16.7.100,3260]
```

```
Logging in to [iface: eth1-iface, target:
iqn.2001-05.com.equallogic:0-8a0906-93ee59d02-
674f999767d4942e-mdi-datal, portal:
10.16.7.100,3260]
```

```
Logging in to [iface: eth1-iface, target:
iqn.2001-05.com.equallogic:0-8a0906-97be59d02-
d7ef99976814942e-mdi-fral, portal:
10.16.7.100,3260]
```

```
Login to [iface: eth1-iface, target: iqn.2001-
05.com.equallogic:0-8a0906-90ee59d02-
e26f999767b4942e-mdi-ocr, portal:
10.16.7.100,3260]: successful
```

```
Login to [iface: eth1-iface, target: iqn.2001-
05.com.equallogic:0-8a0906-95ce59d02-
2e0f999767f4942e-mdi-data2, portal:
10.16.7.100,3260]: successful
```

```
Login to [iface: eth1-iface, target: iqn.2001-
05.com.equallogic:0-8a0906-93ee59d02-
674f999767d4942e-mdi-datal, portal:
10.16.7.100,3260]: successful
```

```
Login to [iface: eth1-iface, target: iqn.2001-
05.com.equallogic:0-8a0906-97be59d02-
d7ef99976814942e-mdi-fral, portal:
10.16.7.100,3260]: successful
```

- 13** Display and verify all active connections and sessions:
`iscsiadm -m session -i`
- 14** Verify that the partitions are visible in operating system:
`cat /proc/partitions`
- 15** Repeat step 1 to step 14 on all the other hosts in the cluster.

Configuring Device Mapper Multipath to Volumes

- 1 Run the `/sbin/scsi_id` command against the devices created for Oracle to obtain their unique device identifiers:

```
/sbin/scsi_id -gus /block/<device>
```

For example:

```
# scsi_id -gus /block/sda
```

- 2 Uncomment the following section in `/etc/multipath.conf`.

```
blacklist {
    wwid <insert local disk WWID here>
    devnode "^(ram|raw|loop|fd|md|dm-
|sr|scd|st)[0-9]*"
    devnode "^hd[a-z]"
}
```

- 3 Uncomment the following section in `/etc/multipath.conf`.

```
defaults {
    udev_dir                /dev
    polling_interval        10
    selector                "round-robin 0"
    path_grouping_policy    multibus
    getuid_callout          "/sbin/scsi_id -g -u -s
/block/%n"
    prio_callout            /bin/true
    path_checker            readsector0
    rr_min_io               100
    max_fds                 8192
    rr_weight               priorities
    failback                immediate
    no_path_retry           fail
    user_friendly_names     yes
}
```

- 4 Add the following section in `/etc/multipath.conf`. The **WWID** is obtained from step 1. Ensure the alias names are consistent on all hosts in the cluster.

```
multipaths {
    multipath {
        wwid    WWID_of_volume1
        alias   alias_of_volume1
    }
    multipath {
        wwid    WWID_of_volume2
        alias   alias_of_volume2
    }
}
```

(Add a multipath subsection for each additional volume.)

```
}
```

The following sample includes configurations of four volumes.

```
multipaths {
    multipath {
        wwid    36090a028d059ee902e94b4
6797996fe2
        alias   ocr
    }
    multipath {
        wwid    36090a028d059ee932e94d4
6797994f67
        alias   data1
    }
    multipath {
        wwid    36090a028d059ce952e94f4
6797990f2e
        alias   data2
    }
    multipath {
        wwid    36090a028d059be972e9414
689799efd7
        alias   fra1
    }
}
```

- 5 Restart the multipath daemon and verify the alias names are displayed in the `multipath -ll` output.

```
service multipathd restart
chkconfig multipathd on
multipath -ll
```

For example,

```
fra1 (36090a028d059be972e9414689799efd7) dm-13
EQLOGIC,100E-00
[size=5.0G][features=1
queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][enabled]
\_ 96:0:0:0 sds 65:32 [active][ready]
\_ 92:0:0:0 sdab 65:176 [active][ready]
ocr (36090a028d059ee902e94b46797996fe2) dm-11
EQLOGIC,100E-00
[size=2.0G][features=1
queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][enabled]
\_ 93:0:0:0 sdf 8:80 [active][ready]
\_ round-robin 0 [prio=0][enabled]
\_ 86:0:0:0 sdad 65:208 [active][ready]
data2 (36090a028d059ce952e94f46797990f2e) dm-8
EQLOGIC,100E-00
[size=20G][features=1
queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][enabled]
\_ 97:0:0:0 sdc 8:32 [active][ready]
\_ 98:0:0:0 sdd 8:48 [active][ready]
data1 (36090a028d059ee932e94d46797994f67) dm-18
EQLOGIC,100E-00
[size=20G][features=1
queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=0][enabled]
\_ 95:0:0:0 sdq 65:0 [active][ready]
\_ 89:0:0:0 sdac 65:192 [active][ready]
```


- 6 Verify the `/dev/mapper/*` devices are created. These devices names must be used to access and interact with multipath devices in the subsequent sections.

For example,

```
# ls -lt /dev/mapper/*
crw----- 1 root root  10, 63 Dec 15 11:22
/dev/mapper/control
brw-rw---- 1 root disk 253, 18 Dec 15 11:51
/dev/mapper/data1
brw-rw---- 1 root disk 253,  8 Dec 15 13:47
/dev/mapper/data2
brw-rw---- 1 root disk 253, 13 Dec 15 11:51
/dev/mapper/fral
brw-rw---- 1 root disk 253, 11 Dec 15 11:51
/dev/mapper/ocr
brw-rw---- 1 root disk 253,  6 Dec 15 11:22
/dev/mapper/osvg-crs
brw-rw---- 1 root disk 253,  3 Dec 15 11:22
/dev/mapper/osvg-home
brw-rw---- 1 root disk 253,  4 Dec 15 11:22
/dev/mapper/osvg-opt
brw-rw---- 1 root disk 253,  0 Dec 15 11:22
/dev/mapper/osvg-root
brw-rw---- 1 root disk 253,  7 Dec 15 11:22
/dev/mapper/osvg-swap
brw-rw---- 1 root disk 253,  1 Dec 15 11:22
/dev/mapper/osvg-tmp
brw-rw---- 1 root disk 253,  2 Dec 15 11:22
/dev/mapper/osvg-usr
brw-rw---- 1 root disk 253,  5 Dec 15 11:22
/dev/mapper/osvg-var
```

- 7 Repeat step 1 to step 6 on all other hosts in the cluster.

Configuring Database Storage on the Host

 **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.

Oracle Real Application Clusters (RAC) requires an ordered list of procedures. The following sections describe how to verify, adjust, and configure the storage disks for the Fibre Channel, direct-attached SAS, or iSCSI storage on your host.

Verifying and Identifying Your Storage Disks

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


 **NOTE:** The procedure in this section describes how to deploy the Oracle database for direct-attached SAS storage as well as for Fibre Channel storage. In the following procedures, Fibre Channel storage nomenclature is used. If direct-attached SAS or iSCSI storage (Dell PowerVault MD3000, MD3000i, or Dell|EqualLogic storage array) is used, see Table 7-1 to translate Fibre Channel nomenclature to direct-attached SAS or iSCSI storage nomenclature.

Table 7-1. Fibre Channel, Direct-Attached SAS, and EqualLogic Array Nomenclature

Fibre Channel Storage	Direct-Attached SAS or iSCSI (MD3000/MD3000i)	EqualLogic Array Volumes
LUNs	Virtual disks	Volumes
/dev/emcpower(X)	/dev/sd(X)	/dev/sd(X)
PowerPath	Multi-Path (MPIO)	Device Mapper

To verify that each node can detect each storage LUN or logical disk:

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



NOTE: If you are installing the cluster or reinstalling the software on a node, you must perform step 1.

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

```
more /proc/partitions
```

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



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

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

In the case of a direct-attached SAS or iSCSI configurations, the virtual disks appears as `/dev/sdb` and `/dev/sdc`.

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

- All PowerPath pseudo devices appear in the file with similar device names across all nodes.
For example: `/dev/emcpowera`, `/dev/emcpowerb`, and `/dev/emcpowerc`.
- In the case of the PowerVault MD3000, MD3000i, or the EqualLogic storage array, all the virtual disks or volumes appear in the file with similar device names across all nodes.
For example: `/dev/sdb`, `/dev/sdc`, and `/dev/sdd`

- The external storage logical volumes appear as SCSI devices and each node is configured with the same number of LUNs, virtual disks, or volumes. For example, if the node is configured with a SCSI drive or RAID container attached:

On a Fibre Channel storage device with three logical disks:


- `sda` identifies the node's RAID container or internal drive
- `emcpowera`, `emcpowerb`, and `emcpowerc` identifies the LUNs (or PowerPath pseudo devices)

On a direct-attached SAS or iSCSI storage device with three virtual disks:

- `sda` identifies the RAID container or internal drive of the node
- `sdb`, `sdc`, and `sdd` identifies the external storage logical volumes


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

Adjusting Disk Partitions for Systems Running the Linux Operating System

 **CAUTION:** In a system running the Linux operating system, align the partition table before data is written to the LUN/virtual disk. The partition map is rewritten and all data on the LUN/virtual disk is destroyed.

Example: fdisk Utility Arguments

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

 **NOTE:** In this example, the disk `/dev/emcpowera` already contains a primary partition `/dev/emcpowera1`. In the case of the PowerVault MD3000, MD3000i, or EqualLogic storage array, the partitions are created on `/dev/sdb1`.

```
fdisk /dev/emcpowera
```

 **NOTE:** Create a partition on `/dev/emcpowera` before performing the following steps.

```
x # expert mode
b # adjust starting block number
1 # choose partition 1
128 # set it to 128, (This is the default stripe element size on the
Dell|EMC CX series Fibre Channel storage)
w # write the new partition
```

For LUNs that have a snapshot, clone, or MirrorView image, it is recommended that you use the `fdisk` method than the LUN alignment offset method. It is also recommended that you make a SAN copy of the sources and targets.

Using the fdisk Utility to Adjust a Disk Partition

To use the `fdisk` utility to adjust a disk partition:

- 1 At the command prompt, type the following:
`fdisk <PartitionName>`,
where `<PartitionName>` is the name of the partition that you are adjusting.
For example, if the partition name is `/dev/emcpowera`, type:
`fdisk /dev/emcpowera`

The system displays the following message:

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

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

1) software that runs at boot time (e.g., old versions of LILO)

2) booting and partitioning software from other OSs

(e.g., DOS FDISK, OS/2 FDISK)

- 2 At the command prompt, type the following **fdisk** utility argument: **x**
- 3 At the command prompt, type the following **fdisk** utility argument: **b**
- 4 When prompted for the partition number, type the partition number at the command prompt. For example: **1**
- 5 Specify the new location on the disk partition for the beginning of data. For example: **128**
- 6 At the command prompt, type the following **fdisk** utility argument: **w**
The system displays the following message:
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
- 7 Repeat step 1 through step 6 for all Oracle data LUNs.

Configuring Shared Storage for Clusterware, Database, and Recovery Files in an RAC Environment

Use the ORACLEASM library driver to configure disks for Automatic Storage Management (ASM). To configure the ASM disks:

- 1 Log in as **root**.
- 2 Open a terminal window and perform the following steps on all nodes:
 - a Type `service oracleasm configure`
 - b Enter the following inputs for all the nodes:

Default user to own the driver interface []:grid


Default group to own the driver interface []:asmadmin

Start Oracle ASM library driver on boot (y/n) [n]:y

Fix permissions of Oracle ASM disks on boot (y/n) [y]:y

- 3 Perform this step only if the RAC configuration uses shared storage and a Linux Device Mapper Multipath driver.

- a Set the ORACLEASM_SCANORDER parameter in `/etc/sysconfig/oracleasm` to `dm`.

 **NOTE:** If you are using EMC storage, change the **ORACLEASM_SCANORDER** parameter in `/etc/sysconfig/oracleasm` to **emcpower**.

- b Reboot the system to apply the changes.

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

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

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

- 5 Repeat step 4 for any additional ASM disks that need to be created.

- 6 Verify that the ASM disks are created and marked for ASM usage. In the terminal window, type the following and press <Enter>:

```
service oracleasm listdisks
```

The disks that you created in step 4 are displayed. For example:

```
ASM1
```

```
ASM2
```

- 7 Ensure that the remaining nodes are able to access the ASM disks that you created in step 4. On each remaining node, open a terminal window, type the following, and press <Enter>:

```
service oracleasm scandisks
```

Index

A

alias names, 62

B

bonded pair, 9

C

cable

 CAT 5e, 24

 CAT 6, 24

cabling

 Fibre Channel storage, 25

 iSCSI cluster, 40

D

Dell/EMC Fibre Channel, 24

DellEMC Fibre Channel

 storage, 24

 switch, 24

Domain Naming Service, 9

E

EMC

 Navisphere, 66

 PowerPath, 66

 ethernet switch, 24

F

Fibre Channel

 direct-attached configuration, 25

 SAN-attached configuration, 26

H

help, 8

 Dell support, 8

 Oracle support, 8

I

ifconfig, 10

IP address

 private IP address, 9

 public IP address, 9

 virtual IP address, 9

J

Jumbo Frame, 54

M

Modular Disk Storage
 Manager, 32

MPIO software, 32

MTU, 54

N

network bonding, 10

network parameters, 11

NIC

 port assignments, 9

P

PowerConnect switch, 49

PowerPath pseudo devices, 66

R

Resource Media

 PowerVault MD3000, 33

S

SAS

 cluster setup, 29, 39, 49, 65

T

term

 Group, 49

 Member, 49

 Pool, 49

 Volume, 49

V

volumes, 52