

Dell™ PowerEdge™ Cluster  
SE600L Systems

# Installation and Troubleshooting Guide

# Notes, Notices, and Cautions



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



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



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

---

**Information in this document is subject to change without notice.**

© 2008 Dell Inc. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, *PowerVault*, and *OpenManage* are trademarks of Dell Inc.; *Intel* is a registered trademark of Intel Corporation; *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 document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

# Contents

1	Overview . . . . .	7
	<b>Types of Clustering.</b> . . . . .	7
	<b>Supported Configuration.</b> . . . . .	8
	Required Components . . . . .	8
	Minimum Requirements . . . . .	10
	Cluster Nodes. . . . .	10
	Management Node. . . . .	11
	Cluster Storage. . . . .	11
	Fencing Mechanisms. . . . .	12
	Cluster Storage Management Software . . . . .	12
	<b>Other Documents You May Need</b> . . . . .	14
2	Cabling Your Cluster . . . . .	17
	<b>Cabling the Mouse, Keyboard, and Monitor</b> . . . . .	17
	<b>Cabling the Power Supplies.</b> . . . . .	17
	<b>Cabling Your Public and Private Networks.</b> . . . . .	20
	Cabling Your Public Network . . . . .	22
	Cabling Your Private Network . . . . .	22
	Using Dual-Port Network Adapters for Your Private Network . . . . .	22
	NIC Bonding . . . . .	23

	<b>Cabling the Storage Systems</b> . . . . .	<b>23</b>
	Cabling the Cluster in a Redundant Configuration With Single SAS 5/E HBAs . . . . .	23
	Cabling the Cluster in a Redundant Configuration With Dual SAS 5/E HBAs . . . . .	26
<b>3</b>	<b>Preparing Your Nodes for Clustering</b> . . . . .	<b>29</b>
	<b>Pre-Requisites</b> . . . . .	<b>29</b>
	Selecting a Management Infrastructure . . . . .	29
	<b>Configuring the Cluster Nodes</b> . . . . .	<b>29</b>
	Configuring SAS HBAs . . . . .	30
	Configuring Remote Access . . . . .	32
	Additional Configuration . . . . .	34
	<b>Configuring the Internal Storage Drives</b> . . . . .	<b>34</b>
	<b>Preparing the Operating System</b> . . . . .	<b>35</b>
	Accessing the Red Hat Network . . . . .	35
	Determine the Operating System Status . . . . .	35
	Installing the Red Hat Enterprise Linux 5.1 Operating System . . . . .	35
	Installing Cluster Groups . . . . .	40
	Dell Community Repositories . . . . .	40
	<b>Configuring Nodes for Storage Access</b> . . . . .	<b>42</b>
	Using the Dell PowerVault MD3000 Resource Media . . . . .	42

<b>Installing and Configuring the Shared Storage System</b> . . . . .	<b>47</b>
Setting Up Your Storage Array . . . . .	47
Configuring Host Access . . . . .	48
Creating a Host Group . . . . .	49
Creating Disk Groups and Virtual Disks . . . . .	50
Creating Host-to-Virtual Disk Mappings . . . . .	50
<b>4 Configuring Your Cluster</b> . . . . .	<b>51</b>
<b>Setting Up a High-Availability Cluster</b> . . . . .	<b>51</b>
Preparing Cluster Nodes for Conga . . . . .	51
Creating Your Cluster Using Conga. . . . .	52
Configuring Fencing Using Conga . . . . .	53
<b>Setting Up a Storage Cluster</b> . . . . .	<b>54</b>
Configuring a Storage Cluster With Conga. . . . .	55
Configuring a Storage Cluster With CLI Tools . . . . .	57
<b>Managing Services</b> . . . . .	<b>64</b>
Creating Services. . . . .	64
Example of Configuring NFS . . . . .	65
<b>5 Troubleshooting</b> . . . . .	<b>67</b>
<b>Physical Connectivity</b> . . . . .	<b>67</b>
<b>Verifying Node Connectivity</b> . . . . .	<b>67</b>
<b>Networking Issues</b> . . . . .	<b>68</b>
<b>Cluster Status</b> . . . . .	<b>68</b>

<b>Troubleshooting Conga</b> . . . . .	<b>69</b>
Running luci on a Cluster Node . . . . .	69
Issues While Creating a Cluster Initially . . . . .	69
<b>Configuration File Issues</b> . . . . .	<b>70</b>
<b>Determining Virtual Disk</b> . . . . .	<b>70</b>
<b>Logical Volume Issues</b> . . . . .	<b>71</b>
<b>Shared Storage Issues</b> . . . . .	<b>71</b>
<b>Testing Fencing Mechanisms</b> . . . . .	<b>72</b>
<b>A Checklist to Set Up Your Cluster</b> . . . . .	<b>73</b>

# Overview

This document provides instructions to use Red Hat® Enterprise Linux® 5.1 Advanced Platform, Red Hat Cluster Suite (RHCS), and Red Hat Global File System (GFS) to configure a Dell™ PowerEdge™ cluster SE600L system with two PowerEdge servers connected to a Dell PowerVault™ MD3000 serial-attached SCSI (SAS) storage array.

## Types of Clustering

Clustering is the process of configuring two or more systems to provide services as if they were a single system. These systems are known as cluster nodes. You can use clustering to enable load balancing between nodes, parallel processing, high availability, and storage clustering. The different types of clustering are discussed below:

- **Load Balancing** — Load balancing clusters accept requests from clients and use a defined algorithm to forward the requests to a pool of nodes. Red Hat Cluster Suite contains **linux virtual server** service that performs this function. Implementation of the linux virtual server function is beyond the scope of this document. For more information, see the Red Hat website at [www.redhat.com/cluster\\_suite](http://www.redhat.com/cluster_suite).
- **Parallel Processing or High Performance Compute Clustering (HPCC)** — Typical HPC clusters are controlled by a *master* or *front-end* server. The master server accepts job requests and divides them up among the compute nodes to take advantage of parallel processing. Red Hat Cluster Suite does not support HPCC architecture.
- **High Availability or Failover Clustering** — Includes a set of nodes capable of taking over the workload from another node in the same highly available cluster, if there is an unexpected failure. Red Hat cluster suite supports high availability cluster configuration. For more information, see the Red Hat cluster suite website at [www.redhat.com/cluster\\_suite](http://www.redhat.com/cluster_suite).
- **Storage Cluster** — A cluster comprised of systems that are accessing the same shared storage such as Red Hat Global File System. For more information, see the Red Hat GFS website at: [www.redhat.com/gfs](http://www.redhat.com/gfs).

This document describes the procedure to configure a *high availability* and *storage cluster*.

# Supported Configuration

The following sections provide information about the supported hardware and software components of a PowerEdge Cluster SE600L system.

The PowerEdge Cluster SE600L system comprises of a high-availability cluster running the Red Hat Cluster Suite and a storage cluster running Red Hat Global File System. You can configure the high-availability functionality in a cluster without a storage cluster, but you cannot configure a storage cluster without the high-availability functionality.

## Required Components

The PowerEdge Cluster SE600L system provides a fully redundant high-availability solution. If any component including a SAS HBA, NIC, hard disk, storage controller, cable, or cluster node fails, the cluster service is operational and clients continue to access it.

The minimum required components for the cluster solution are:

- Two PowerEdge servers with at least two hard drives and two network interface cards each
- One PowerVault MD3000 storage array

The recommended configuration includes the following components:

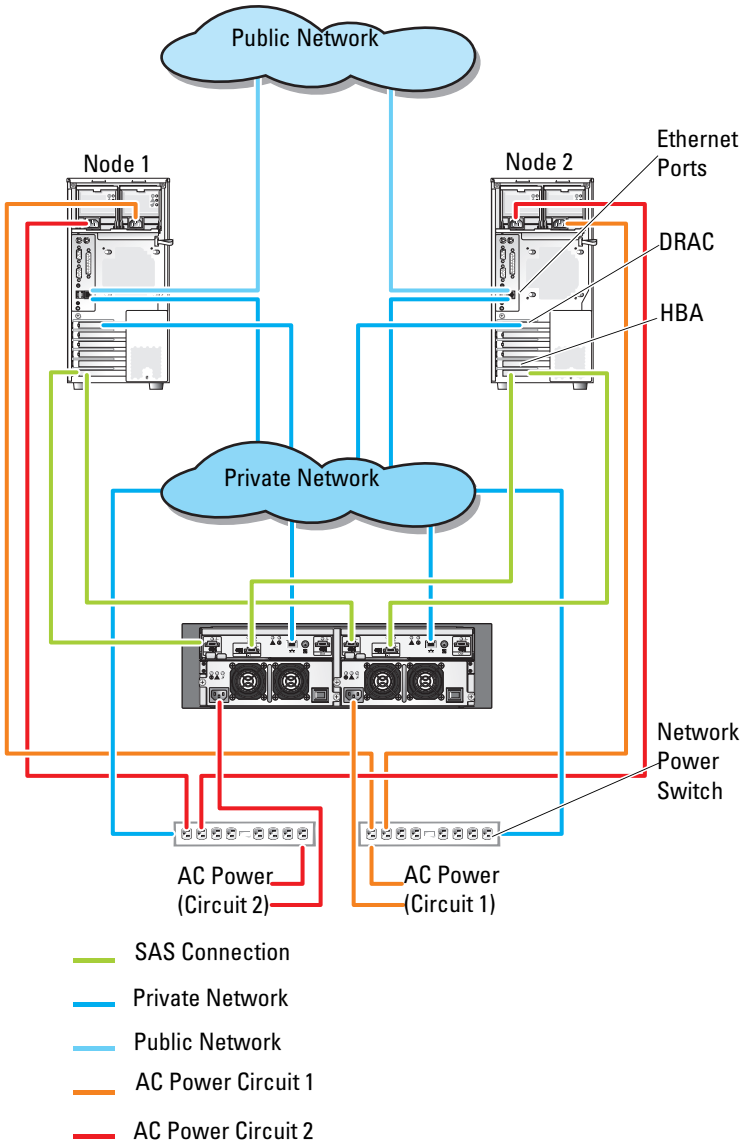
- Two PowerEdge servers with at least three hard drives and more than two network interface cards each
- One PowerVault MD3000 storage array
- Two PowerVault MD1000 expansion arrays
- Two network power switches
- Two network switches
- One management node installed with Red Hat Enterprise Linux 5 operating system



**NOTE:** You can use one of the cluster nodes as a management node, however it is recommended that you have another system dedicated as the management node.



**Figure 1-1. Example of a PowerEdge Cluster SE600L System Configuration**



## Minimum Requirements

The PowerEdge Cluster SE600L system supports the Red Hat Enterprise Linux 5.1 Advanced Platform operating system. The following list provides an overview of the minimum supported configuration:

- Minimum disk space 1 GB per node
- Minimum required RAM is 512 MB per node



**NOTE:** The amount of RAM supported per node also depends on your cluster node.



**NOTE:** Running different versions of Red Hat Enterprise Linux in a cluster is supported only during a rolling upgrade.

## Cluster Nodes

The PowerEdge Cluster SE600L system requires two PowerEdge systems in a homogeneous pair. Each PowerEdge system must run the Red Hat Enterprise Linux 5.1 AP operating system. Table 1-1 provides the requirements for each PowerEdge system in the cluster.

**Table 1-1. Requirements for Each Cluster Node**

Component	Minimum Requirement
Processor	At least one processor for each cluster node
RAM	At least 512 MB RAM installed on each cluster node for Red Hat Enterprise Linux 5.1 Advanced Platform
Host Bus Adapter (HBA)	Two SAS 5/E HBAs for each cluster node
Network Interface Cards (NICs)	At least two NICs per cluster node—one NIC for the public network and another NIC for the private network <b>NOTE:</b> It is recommended that the NICs on each public network are identical and that the NICs on each private network are identical.
Internal Disk Controller	One controller connected to internal disks for each node. <b>NOTE:</b> It is strongly recommended that you use hardware or software RAID with the internal drives.
Power Supplies	At least two power supplies per node supplied by different power sources.

## Management Node

It is recommended that you configure a dedicated management node running Red Hat Enterprise Linux operating system. The management node performs the following functions in your PowerEdge Cluster SE600L system configuration:

- Manages and monitors the cluster and cluster storage
- Manages and monitors the PowerVault MD3000 storage array
- Performs package management for nodes

Additionally a dedicated management node can also provide the following functions that are required for your cluster environment. The detailed description of the following functions are outside the scope of this document:

- Kickstart network installation server
- Pre-execution environment (PXE) server
- Dynamic host control protocol (DHCP) server
- Domain Naming Service (DNS)

## Cluster Storage

The PowerEdge Cluster SE600L system supports the following storage enclosures:

- One PowerVault MD3000 RAID enclosure
- Up to two PowerVault MD1000 storage expansion enclosures


Table 1-2 provides the configuration requirements for the shared storage system.


**Table 1-2. Requirements for the Storage System**

<b>Hardware Components</b>	<b>Minimum Requirement</b>
Power and cooling requirements	Two integrated hot-pluggable power supply/cooling fan modules
Physical disks	At least two physical disks in the PowerVault MD3000 RAID enclosure

**Table 1-2. Requirements for the Storage System (continued)**

Hardware Components	Minimum Requirement
Cables	<ul style="list-style-type: none"><li>• Four 1 metre, 2 metre, or 4 metre SAS cables for the redundant configuration.</li><li>• Two 1 metre, 2 metre, or 4 metre SAS cables for each additional PowerVault MD1000 storage expansion enclosure</li></ul>

 **NOTE:** RAID 0 and independent disks are possible but are not recommended for a high-availability system because they do not offer data redundancy if a disk failure occurs.


 **NOTE:** The PowerVault MD3000 storage array does not support more than two connected systems, therefore only two nodes are possible in a cluster.

## Fencing Mechanisms

Fencing is used to ensure data integrity on the shared storage. If a node does not respond to the cluster services in a timely manner, it is removed from the cluster. There are several possible methods to use for fencing. For the complete list of fencing devices, see the Red Hat website at [www.redhat.com/cluster\\_suite/hardware](http://www.redhat.com/cluster_suite/hardware).

The PowerEdge Cluster SE600L system is tested with the following devices:

- Dell Remote Access Controller (DRAC) 5
- PowerEdge Intelligent Platform Management Interface (IPMI)
- APC® network power switch (AP7867A)

 **NOTE:** It is recommended that you use two fencing methods, such as a network power switch and a remote access controller (DRAC or IPMI).

## Cluster Storage Management Software

The following sections describe various storage management software components.

### **Dell PowerVault Modular Disk Storage Manager Client**

The PowerVault Modular Disk Storage Manager Client runs on the management station to centrally manage the PowerVault MD3000 RAID enclosures. You can use the PowerVault Modular Disk Storage Manager to perform tasks such as creating or managing RAID arrays, binding virtual disks, and downloading firmware.

### **PowerVault Modular Disk Storage Manager Agent**

The PowerVault Modular Disk Storage Manager Agent resides on each cluster node and collects server-based topology data that can be managed by the Modular Disk Storage Manager Client.

### **Multi-Path Proxy Driver**

The Multi-Path Proxy driver resides on each cluster node and manages the redundant data path between the server and the RAID enclosure. For proper functioning of the Multi-Path Proxy driver, ensure that the configuration includes redundant HBAs and cabling.

The Multi-Path Proxy driver identifies the existence of multiple paths to a virtual disk and establishes a preferred path to that disk. If any component in the preferred path fails, the driver automatically re-routes I/O requests to the alternate path so that the storage array continues to operate without interruption.

In a redundant cluster configuration, the automatic failback feature is disabled. Therefore, when a failed component is repaired or replaced, the virtual disk(s) do not automatically transfer to the preferred controller. You can manually initiate a failback using the Modular Disk Storage Manager Client or Command Line Interface (CLI).

## Other Documents You May Need



**CAUTION:** The *Product Information Guide* provides important safety and regulatory information. Warranty information may be included within this document or as a separate document.

- The *Rack Installation Guide* included with your rack solution describes how to install your system into a rack.
- The *Setting Up Your System* document provides an overview of initially setting up your system.
- The *Dell PowerVault MD3000 Hardware Owners Manual* and the *Dell PowerEdge System Hardware Owners Manual* provide information about the host servers and storage array.
- The *Dell PowerVault MD3000 Systems Installation Guide* provides an overview of setting up and cabling your storage array.
- *The PowerVault Modular Disk Storage Manager CLI Guide* provides information about using the command line interface (CLI).
- The *Dell PowerVault MD3000 Resource Media* provides documentation for configuration and management tools, as well as the full documentation set included in this section.
- The *Dell PowerVault Modular Disk Storage Manager User's Guide* provides instructions for using the array management software to configure the storage array.
- The *Dell PowerVault MD Systems Support Matrix* provides information on supported software and hardware for PowerVault MD systems. This document can be located on the Dell Support website at [support.dell.com](http://support.dell.com).
- The *System Administrator's Guide* provides system operation and management information.
- Documentation for any components you purchased separately provides information to configure and install these options.
- Updates are sometimes included with the system to describe changes to the system, software, and/or documentation.



**NOTE:** Always read the updates first because they often supersede information in other documents.

- For Red Hat documentation, see the following links:
  - Red Hat Enterprise Linux manuals  
[www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise)
  - Red Hat Hardware Certification - [hardware.redhat.com](http://hardware.redhat.com)
  - Red Hat Cluster Suite - [www.redhat.com/cluster\\_suite/](http://www.redhat.com/cluster_suite/)
  - Red Hat Global File System - [www.redhat.com/gfs/](http://www.redhat.com/gfs/)
  - Red Hat Network - [www.redhat.com/docs/manuals/RHNetwork](http://www.redhat.com/docs/manuals/RHNetwork)
  - Additional information is also available from the Red Hat training class RH436—*Red Hat Enterprise Clustering and Storage Management*

Release notes or readme files may be included to provide last-minute updates to the system documentation or advanced technical reference material intended for experienced users or technicians.





# Cabling Your Cluster

The following sections provide information on how to cable various components of your cluster.

## Cabling the Mouse, Keyboard, and Monitor

When you are installing the cluster configuration in a rack, it is recommended but not required that you connect a mouse, keyboard, and a monitor. For instructions about cabling the node connections, see the documentation included with your rack.

## Cabling the Power Supplies

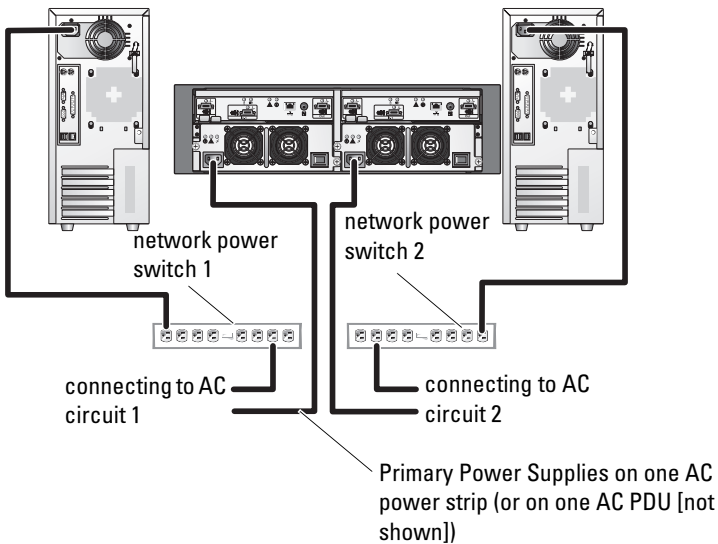
To ensure that the specific power requirements are satisfied, see the documentation for each component in your cluster solution.

It is recommended that you adhere to the following guidelines to protect your cluster solution from power-related failures:

- Plug each power supply into a separate AC circuit.
- Plug each power supply into separate optional network power switches.
- Use uninterruptible power supplies (UPS).
- For some environments, consider having backup generators and power from separate electrical substations.

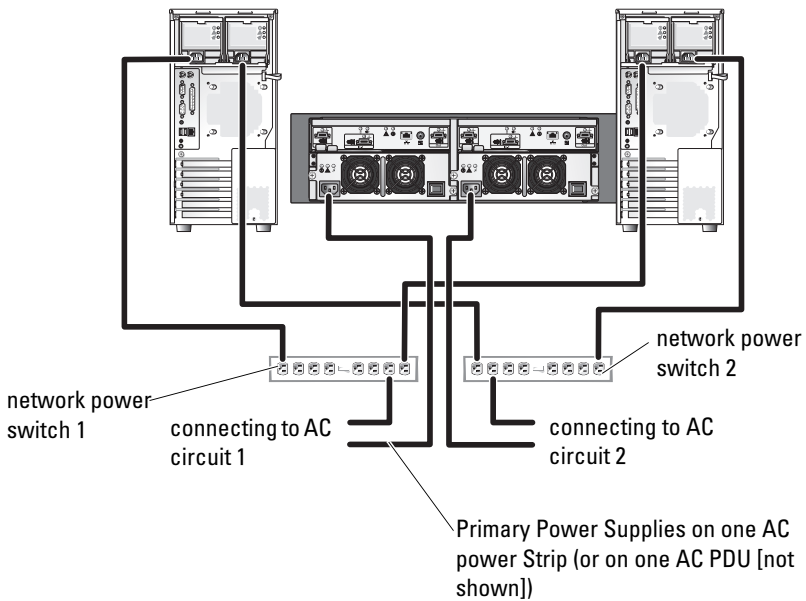
Figure 2-1 and Figure 2-2 illustrate recommended methods for power cabling of a cluster solution consisting of two Dell™ PowerEdge™ systems and one storage system. To ensure redundancy, the primary power supplies of all the components are grouped into one or two circuits and the redundant power supplies are grouped into a different circuit.

**Figure 2-1. Power Cabling Example With One Power Supply in the PowerEdge Systems**



**NOTE:** This illustration is intended only to demonstrate the power distribution of the components.

**Figure 2-2. Power Cabling Example With Two Power Supplies in the PowerEdge Systems**



**NOTE:** This illustration is intended only to demonstrate the power distribution of the components.

# Cabling Your Public and Private Networks


The network adapters in the cluster nodes provide at least two network connections for each node. These connections are described in Table 2-1.

**Table 2-1. Network Connections**

<b>Network Connection</b>	<b>Description</b>
Public Network	<p>All connections to the client local area network (LAN).</p> <p>At least one public network must be configured for mixed mode (public mode and private mode) for private network failover.</p>
Private Network	<p>A dedicated connection for sharing cluster status information between the cluster nodes.</p> <p>Network adapters connected to the LAN can also provide redundancy at the communications level in case the cluster interconnect fails.</p> <p>If you have optional DRAC cards, cable them to the private network. In this configuration, do not use point-to-point topology, use a network switch.</p> <p>If you have optional network power switches, cable them to the private network. Cable the PowerVault MD3000 Storage Controller Ethernet ports to the private network.</p>

Figure 2-3 shows an example of network adapter cabling in which dedicated network adapters in each node are connected to the public network and the remaining network adapters are connected to each other (for the private network).

**Figure 2-3. Example of Network Cabling Connection**



## Cabling Your Public Network

Any network adapter supported by a system running TCP/IP may be used to connect to the network segments. You can install additional network adapters to support additional public network segments or to provide redundancy in the event of a faulty network adapter or switch port.

## Cabling Your Private Network

The private network connection to the cluster nodes is provided by a second or subsequent network adapter that is installed in each node. This network is used for intra-cluster communications. Cable all cluster network components including cluster node private network interface cards, network power switches, remote access controllers (DRAC/IPMI), and MD3000 Storage Controller Ethernet ports. Table 2-2 lists the required hardware components and connection method for two possible private network configurations.

**Table 2-2. Private Network Hardware Components and Connections**

<b>Method</b>	<b>Hardware Components</b>	<b>Connection</b>
Network switch	Fast Ethernet or Gigabit Ethernet network adapters and switches.	Connect standard Ethernet cables from the network adapters in both cluster nodes to a Fast Ethernet or Gigabit Ethernet switch.
Point-to-Point	Copper Gigabit Ethernet network adapters.	Connect a standard Ethernet cable between the Gigabit Ethernet network adapters in both cluster nodes.

## Using Dual-Port Network Adapters for Your Private Network

You can configure your cluster to use the public network as a failover for private network communications. However, if dual-port network adapters are used, do not use two ports on the same adapter to connect to both the public and private networks.

## NIC Bonding

Network Interface Card (NIC) bonding combines two or more NICs to provide load balancing and/or fault tolerance. Your cluster supports NIC bonding. Use the same NICs in a bond for consistent performance. For information on configuring bonding, see the Red Hat Deployment Guide section "Channel Bonding Interfaces" at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

## Cabling the Storage Systems

This section provides information for connecting your cluster to a storage system.

Storage management can be either in-band through the SAS connection, or out-of-band using an Ethernet connection. For out-of-band storage management, cable the Ethernet ports on the storage array to the private network.



**NOTE:** It is recommended that you configure your Dell PowerVault™ MD3000 storage enclosure to use out-of-band management with a dedicated management node.

### Cabling the Cluster in a Redundant Configuration With Single SAS 5/E HBAs

Each cluster node attaches to the storage system using one SAS 5/E dual-port HBA and two SAS cables. In this configuration, there are redundant storage paths from the cluster node to the storage system. If a component fails in the storage path such as the port, the cable, or the storage controller, the Multi-Path Proxy driver automatically reroutes the I/O requests to the alternate path so that the storage array continues to operate without interruption.

To cable the cluster:

- 1** Connect cluster node 1 to the storage system.
  - a** Install a SAS cable from the cluster node 1 HBA port 0 to the RAID controller module 0 port In-0.
  - b** Install a SAS cable from the cluster node 1 HBA port 1 to the RAID controller module 1 port In-0.
- 2** Connect cluster node 2 to the storage system.
  - a** Install a SAS cable from the cluster node 2 HBA port 0 to the RAID controller module 0 port In-1.
  - b** Install a SAS cable from the cluster node 2 HBA port 1 to the RAID controller module 1 port In-1.

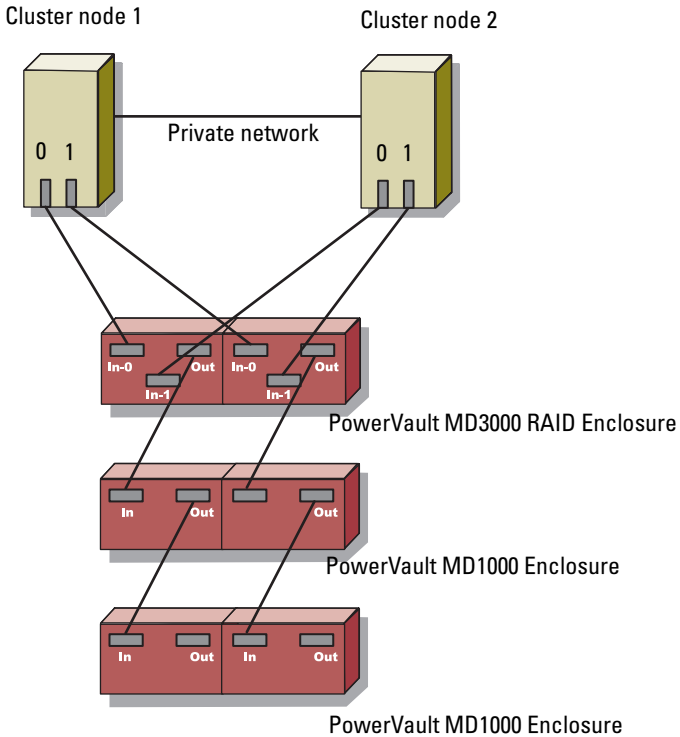


**NOTE:** If the HBA on the primary server fails, the node is removed from the cluster and all services using the node are relocated.



Figure 2-4 shows a redundant cluster configuration with a single SAS 5/E HBA.

**Figure 2-4. Redundant Cluster Configuration With Single SAS 5/E HBA**



**NOTE:** The SAS out port provides SAS connection for cabling to MD1000 expansion enclosure(s).

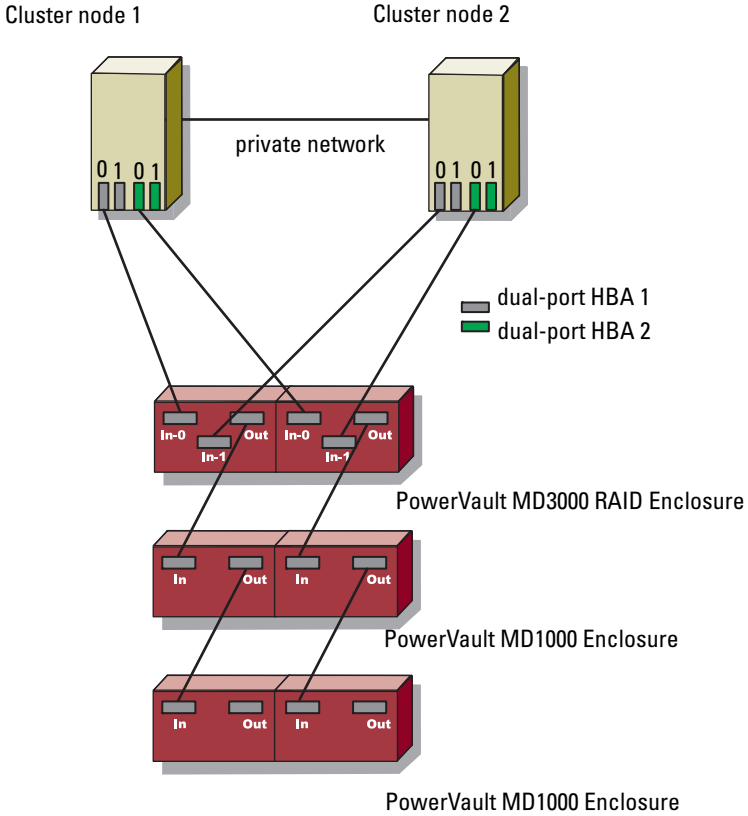
## **Cabling the Cluster in a Redundant Configuration With Dual SAS 5/E HBAs**

Each cluster node attaches to the storage system using two SAS 5/E HBA and two SAS cables. In this configuration, there are redundant storage paths from the cluster nodes to the storage system. If a component fails in the storage path such as the port, the cable, or the storage controller, the Multi-Path Proxy driver automatically re-routes the I/O requests to the alternate path so that the storage array continues to operate.

To cable the cluster:

- 1** Connect cluster node 1 to the storage system.
  - a** Install a SAS cable from the cluster node 1 HBA port 0 to the RAID controller module 0 port In-0.
  - b** Install a SAS cable from the cluster node 1 HBA port 1 to the RAID controller module 1 port In-0.
- 2** Connect cluster node 2 to the storage system.
  - a** Install a SAS cable from the cluster node 2 HBA port 0 to the RAID controller module 0 port In-1.
  - b** Install a SAS cable from the cluster node 2 HBA port 1 to the RAID controller module 1 port In-1.

**Figure 2-5. Redundant Cluster Configuration With Dual SAS 5/E HBAs**





# Preparing Your Nodes for Clustering

## Pre-Requisites

To setup the Dell™ PowerEdge™ Cluster SE600L system, ensure you have the following:

- The components listed in "Required Components" on page 8
- Red Hat® Enterprise Linux® 5.1 Advanced Platform installation media
- *Dell PowerVault™ MD3000 Resource Media* or the latest ISO image of the media from the Dell Support website at [support.dell.com](http://support.dell.com)

## Selecting a Management Infrastructure

Ensure you have completed all cabling and hardware setup. Before continuing, select the management infrastructure. You can manage a storage array in two ways:

Out-of-band management—For out-of-band management, data is separate from commands and events. Data travels through the host-to-controller SAS interface cables, while commands and events travel through the Ethernet cables.

In-band management—For in-band management, commands, events, and data travel through the host-to-controller SAS interface cables. Unlike out-of-band management, commands and events are on the same network as data.



**NOTE:** It is recommended that you use out-of-band management for the PowerEdge Cluster SE600L system.

## Configuring the Cluster Nodes

To prepare the cluster nodes for setup, execute the instructions in the following sections.

## Configuring SAS HBAs

- 1 Turn on each node.
- 2 During the BIOS post, use the option to invoke the serial-attached SCSI (SAS) 5/E configuration (which is typically <Ctrl><C> key combination). The following screen appears.

```
Dell SAS 5 Host Bus Adapter BIOS
MPTBIOS-6.12.05.00 (2007.09.29)
Copyright 2000-2007 LSI Corporation

Adapter(s) disabled by user
Press Ctrl-C to run SAS Configuration Utility...
```

The following message appears:

```
Dell SAS 5 Host Bus Adapter BIOS
MPTBIOS-6.12.05.00 (2007.09.29)
Copyright 2000-2007 LSI Corporation

Adapter(s) disabled by user
Please wait, invoking SAS Configuration Utility...
```

- 3 In the SAS configuration utility menu, select an adapter from the list as shown in the figure below:

```
LSI Corporation Config Utility For Dell SAS 5
Adapter List Global Properties
Adapter      PCI   PCI   PCI   FW Revision
              Bus  Dev  Fnc
SAS1068     0B   08   00   0.10.51.00-IT
SAS1068     0F   08   00   0.10.51.00-IT
```

After you select an adapter, the adapter properties are displayed as shown below:

```
LSI Corporation Config Utility For Dell SAS 5 v6.12.05.00
Adapter Properties -- SAS1068

Adapter                SAS1068
PCI Address (Bus/Dev/Func) 0B:08:00
MPT Firmware Revision    0.10.51.00-IT
SAS Address              500188B1:4D924100
NVDATA Version           25.10
Status                   Disabled
Boot Order               0
Boot Support              [Disabled]

SAS Topology

Advanced Adapter Properties
```

- 4 Record the SAS Address. The SAS address is used to assign virtual disks later.
- 5 Toggle the option for **Boot Support** and select **Disabled**. This field enables you to boot from the PowerVault MD3000 storage array and is not supported with the PowerEdge Cluster SE600L systems.
- 6 Save the configuration and exit.
- 7 Repeat the procedure on all other SAS 5/E adapters.
- 8 Exit the configuration utility and reboot the system.

## Configuring Remote Access

- 1 In the BIOS post, use the <Ctrl><E> key combination to go to the Remote Access Configuration Utility menu when the following message appears:

```
Remote Access Configuration Utility 1.05
Copyright 2006 Dell Inc. All Rights Reserved

Baseboard Management Controller Revision 1.33
Remote Access Controller Revision (Build 06.05.12) 1.0
Primary Backplane Firmware Revision 1.05

IP Address: 172. 16. 0 .102
Netmask: 255.255.255. 0
Pressa<Ctrl-E>1for0Remote Access Setup within 5 sec.....
```

- 2 The Remote Access Configuration Utility menu appears.

```
Remote Access Configuration Utility
Copyright 2006 Dell Inc. All Rights Reserved 1.05


Baseboard Management Controller Revision 1.33
Remote Access Controller Revision (Build 06.05.12) 1.0
Primary Backplane Firmware Revision 1.05

IPMI Over LAN ..... Off
NIC Selection ..... Dedicated
LAN Parameters ..... <ENTER>
Advanced LAN Parameters ..... <ENTER>
Virtual Media Configuration ..... <ENTER>
LAN User Configuration ..... <ENTER>
Reset To Default ..... <ENTER>
System Event Log Menu ..... <ENTER>


Up,Down Arrow to select   SPACE,+,- to change   ESC to exit   F1=Help
```




- If you have purchased the PowerEdge systems with Dell Remote Access Controller (DRAC) cards, ensure the option **IPMI Over LAN** is set to **off**. If IPMI Over LAN is enabled, an additional network interface card (NIC) may be required for a fully-redundant setup. This issue occurs as IPMI uses one of on-board NICs, thereby preventing redundant network configuration without additional NICs.

 **NOTE:** If you are using IPMI instead of a DRAC card for fencing on a PowerEdge 1950 system, a fully-redundant setup is not possible due to the limited number of slots for additional cards.

- In the **Remote Access Configuration Utility** menu, select **LAN Parameters** and to assign an IP address or select **DHCP**.

 **NOTE:** If you are using the DHCP option, you must assign the IP address from a DHCP server. Record the MAC Address from this menu to assign a static IP through DHCP later. For instructions on configuring a DHCP server, see the "Configuring a DHCP Server" section in the *Red Hat Deployment Guide* at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

 **NOTE:** Ensure that the IP address assigned is on the same subnet as the cluster private network because the nodes communicate with the remote access controllers as part of cluster fencing operations.

```

Remote Access Configuration Utility
Copyright 2006 Dell Inc. All Rights Reserved 1.05

Baseboard Management Controller Revision      1.33
Remote Access Controller Revision (Build 06.05.12)  1.0
Primary Backplane Firmware Revision          1.05

IPMI Over LAN Selection:
NIC Selection:
LAN Parameters:
Advanced Virtual LAN User:
Reset To System Defaults:
  RACP+ Encryption Key ..... <ENTER>
  IP Address Source ..... DHCP
  Ethernet IP Address ..... 172.16.0.102
  MAC Address ..... 00:18:8B:38:5E:F5
  Subnet Mask ..... 255.255.255.0
  Default Gateway ..... 172.16.0.254
  VLAN Enable ..... Off
  VLAN ID ..... 0001
  icated
  TER>
  TER>
  TER>
  TER>
  TER>
  TER>

Up,Down Arrow to select   SPACE,+,- to change   ESC to exit   F1=Help

```

- Save changes and Exit.

## Additional Configuration

You will need to configure your DRAC cards for telnet communication with the fencing agents. DRAC cards are shipped with secure shell (ssh) enabled and telnet disabled by default. Support for ssh in fencing agents is planned for a future release.

- 1 Connect to each DRAC with the following command:  

```
[root]# ssh {ip address of DRAC}
```

For example:  

```
[root]# ssh 192.168.120.100
```
- 2 Repeat this process on all nodes with DRACs.
- 3 Enable telnet with the following command:  

```
[root@drac]# racadm config -g cfgSerial -o  
cfgSerialTelnetEnable 1
```



**NOTE:** Perform the procedure in this section after your nodes are configured, if you cannot connect to DRAC using `ssh` at this time. You can also use the web user interface (UI) or Dell OpenManage to configure your DRACs. For more information on configuring your DRAC, see the documentation the Dell Support website at [support.dell.com](http://support.dell.com).

## Configuring the Internal Storage Drives

If you have added new physical disks to your system or are setting up the internal drives in a Redundant Array of Independent Discs (RAID) configuration, you must configure the RAID using the RAID controller's BIOS configuration utility before you can install the operating system.

For the best balance of fault tolerance and performance, it is recommended that you use RAID 5 for the internal disks. For more information on RAID configurations, see the documentation for your specific RAID controller.



**NOTE:** If you are not using the Dell PowerEdge Expandable RAID Controller (PERC) RAID solution and want to configure fault tolerance, use the software RAID included with the Red Hat Enterprise Linux operating system. For more instructions, see the "Configuring Software RAID" section in the *Red Hat Deployment Guide* on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

# Preparing the Operating System



**NOTICE:** Disconnect all serial-attached SCSI (SAS) storage cables during any operating system installation. Failure to do so can result in lost data, boot sector installation issues, and multipath drive ordering problems.

## Accessing the Red Hat Network

If you do not have an Red Hat Network (RHN) account, log on to [www.redhat.com/register](http://www.redhat.com/register) and create an account. The cluster nodes are shipped with RHN subscription information. If you are unable to locate the registration information, contact Dell Support.

## Determine the Operating System Status

If your PowerEdge system shipped with Red Hat Enterprise Linux 5 operating system installed from the factory, you must upgrade to Red Hat Enterprise Linux 5.1 operating system. To upgrade the operating system, see "Updating Your Operating System" on page 39. To install the operating system on your PowerEdge systems, see "Installing the Red Hat Enterprise Linux 5.1 Operating System" on page 35.

## Installing the Red Hat Enterprise Linux 5.1 Operating System

If the operating system media is not included with your PowerEdge system, download the latest Red Hat Enterprise Linux (version 5 for 64-bit x86\_64) ISO image(s) from the Red Hat website at [rhn.redhat.com](http://rhn.redhat.com). If you require physical media, contact your Dell sales representative.

After obtaining the correct ISO image(s) or physical media, choose one of the following installation methods.

### Installing the Operating System Using Physical Media

Create physical media for installation. For more information, see the section "Can You Install Using the CD-ROM or DVD?" in the *Red Hat Installation Guide* located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).



**NOTE:** You can also use the *Dell Systems Build and Update Utility* media to install the operating system. You can download the latest media from the Dell Support website at [support.dell.com](http://support.dell.com).

## Installing the Operating System Through the Network

Create a network installation server. For more information, see the section "Preparing for a Network Installation" in the *Red Hat Installation Guide* located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

## Using a Kickstart Installation

Your PowerEdge Cluster SE600L system can be easily installed with the SE600L kickstart script located at <http://linux.dell.com/files/se600l/ks.cfg>. However, your nodes must be connected to the Internet to use this script. If your nodes do not have internet access, you will need to create your own kickstart file. See "Creating Your Own Kickstart File" on page 36.

For complete details on kickstart, see section *Kickstart Installations* in the *Red Hat Installation Guide* on [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

## Creating Your Own Kickstart File

You can download the SE600L kickstart script and modify it to meet your needs, or if your cluster nodes do not have direct internet access.

Download the SE600L kickstart script from <http://linux.dell.com/files/se600l/ks.cfg>

For more information on creating a kickstart installation, see the section "Creating a Kickstart Installation" in the *Red Hat Installation Guide* located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise/](http://www.redhat.com/docs/manuals/enterprise/).

For information on using the kickstart GUI, see the section *Kickstart Configurator* of the *Red Hat Installation Guide* on [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise) for information.

## Using the Network Installation

You can use one of the network installation methods described in the following sections to install your Red Hat Enterprise Linux operating system.



**NOTE:** Use the Pre-Boot Execution Environment (PXE) for easy installation of multiple nodes. For more information on setting up PXE, see the section "PXE Network Installations" in the *Red Hat Installation Guide* located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

### ***Using PXE Boot to Install the Operating System From the Network***

Configure an entry that points to the network installation server you configured in "Installing the Operating System Through the Network" on page 36 using the kickstart that you created in "Creating Your Own Kickstart File" on page 36. Example of a PXE entry:

```
label node1

    menu label node1 RHEL 5.1 (registered with RHN,
        running SE600L HA Cluster)

    kernel
    images/os/linux/redhat/rhel/5.1/x86_64/vmlinuz

    append initrd=
    images/os/linux/redhat/rhel/5.1/x86_64/initrd.img
```

### ***Using Media to Install the Operating System From the Network***

You can boot from any Red Hat Enterprise Linux 5.1 CD1 or DVD and specify your network installation server using the `askmethod` parameter. For example:

```
boot: linux askmethod
```

You can also use the kickstart file that you created in "Using a Kickstart Installation" on page 36. For example:

```
boot: linux ks=
http://linux.dell.com/files/se600l/ks.cfg
```



**NOTE:** If you have modified the SE600L kickstart script or provided your own, specify that kickstart file instead.


### **Synchronize Node System Clocks**

Synchronize the clocks on all nodes. For more information on configuring system time, see the *Red Hat Deployment Guide* located on the Red Hat website at [www.redhat.com](http://www.redhat.com).

## Registering Your Nodes With RHN

Register all nodes with RHN at [rhn.redhat.com/register](http://rhn.redhat.com/register) by executing the following command:

```
[root]# rhn_register
```

 **NOTE:** If you encounter a message that indicates the node is already registered with RHN, you do not need to register it again.


For more information, see the section "RHN Registration" in the *Red Hat Installation Guide* located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

Your installation number provides access to all required RHN Channels. Verify that all nodes are configured on the following RHN Channels:

- Red Hat Enterprise Linux (v. 5 for 64-bit x86\_64)
- Red Hat Enterprise Linux Clustering (v. 5 for 64-bit x86\_64)
- Red Hat Enterprise Linux Cluster-Storage (v. 5 for 64-bit x86\_64)


Run the following command on each node to verify that they are registered:

```
[root]# echo "repo list" | yum shell
```

 **NOTE:** A future version of `yum` will support the command `yum repolist`.

If all three channels are not registered:

- 1 Log in to [rhn.redhat.com](http://rhn.redhat.com).
- 2 Click **Systems**.
- 3 If your cluster nodes are not listed, use a filter such as **Recently Registered** on the left-pane.
- 4 Select your cluster node from the list.
- 5 In the **Subscribed channels** section, click **Alter Channel Subscriptions**.
- 6 Select all channels and click **Change Subscriptions**.

 **NOTE:** For more information on registering with RHN and subscribing to channels, see Red Hat network documentation website at [www.redhat.com/docs/manuals/RHNetwork](http://www.redhat.com/docs/manuals/RHNetwork).

## Using the Dell-se600l Post Installation Script

If you did not use the kickstart file to install your operating system, you can use the `dell-se600l` utility with the following command:

```
[root]# yum install dell-se600l
```

Ensure that you have configured the Dell community repository using the procedure in "Installing Dell Community Repositories" on page 41.

## Updating Your Operating System

After you register the cluster nodes with RHN, you can use the `yum` utility to manage software updates. Red Hat Enterprise Linux operating systems also include a GUI-based tool called `system-config-packages`. To access `system-config-packages`, go to **Applications**→**Add/Remove Software** that provides a front-end for the `yum` utility.

If needed, install `system-config-packages` with the command:

```
[root]# yum install pirut
```

Update to the latest software with the following command:

```
[root]# yum update
```

This command updates all packages for your operating system to the latest version. The PowerEdge Cluster SE600L system has been tested and qualified with Red Hat Enterprise Linux 5.1 operating system.

## Configuring the Firewall

To configure your firewall:

- 1 Ensure that all nodes can communicate with each other by host name and IP address.

See the Troubleshooting section under "Verifying Node Connectivity" on page 67 to verify node connectivity. For more information, see the section "Before Configuring a Red Hat Cluster" in the *Red Hat Deployment Guide* located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).



**NOTE:** If you are not using DNS to resolve names, edit the local host file at `/etc/hosts` on each node and create entries for each node. Ensure that the local host file of each node contains entries for all nodes.

- 2 Configure necessary IP traffic between the nodes. Run the following commands on each node to allow traffic between the nodes:

```
[root]# iptables -I INPUT -s {node1 fully
qualified hostname or ip address} -j ACCEPT

[root]# iptables -I INPUT -s {node2 fully
qualified hostname or ip address} -j ACCEPT

[root]# iptables -I INPUT -s {management node
fully qualified hostname or ip address} -j ACCEPT

[root]# service iptables save
```

This allows all traffic from each node. However, if you require more control over the security between the nodes, you can allow only specific ports that are needed. For more information on firewalls, see the section "Firewalls" in the *Red Hat Enterprise Linux Deployment Guide* located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

For information about IP port numbers, protocols, and components used by Red Hat Cluster Suite, see the section *Enabling IP Ports on Cluster Nodes* in the *Red Hat Enterprise Linux 5.1 Cluster Administration Guide* located on the Red Hat website at [www.redhat.com](http://www.redhat.com).

## Installing Cluster Groups

If your systems were installed with the kickstart file, then they should already have the correct cluster groups. Run the following command to verify or install the correct groups:

```
[root]# yum groupinstall "Clustering" "Cluster
Storage"
```

## Dell Community Repositories

Your Dell PowerEdge systems can be managed using the **yum** utility that already manages packages for Red Hat Enterprise Linux 5.1 on your nodes. This provides easy access to any Dell provided Linux software, Dell PowerEdge firmware updates, some **dkms** drivers, and other open source software. For more information, visit [linux.dell.com/repo/software/](http://linux.dell.com/repo/software/)



If your nodes do not have access to the internet, skip this section and perform manual firmware updates to your PowerEdge systems as needed. Visit [support.dell.com](http://support.dell.com) for information.



**NOTE:** The repositories are community supported, and not officially supported by Dell. For support, use the linux-poweredge mailing list on [lists.us.dell.com](http://lists.us.dell.com). You may also obtain the latest MD3000 and OpenManage Server Assistant (OMSA) ISO images from [support.dell.com](http://support.dell.com).

## Installing Dell Community Repositories



**NOTE:** Ensure that you enter the CLI commands mentioned below in one line.

To install the Dell Community Software Repository, use the following command:

```
[root]# wget -q -O -  
http://linux.dell.com/repo/software/bootstrap.cgi |  
bash
```

To install the Dell Community Hardware Repository, use the following command:

```
[root]# wget -q -O -  
http://linux.dell.com/repo/hardware/bootstrap.cgi |  
bash
```

(Optional) To leverage the repositories to install OpenManage DRAC components, use the following command:

```
[root]# yum install srvadmin-rac5
```

OpenManage gives you the ability to use the **racadm** command to manage your DRAC components. This installs the required components. If you need the complete OpenManage functionality, use this command instead:

```
[root]# yum install srvadmin-all
```

Install the SAS HBA driver:

```
[root]# yum install mptlinux
```

Install the Multi-Path Proxy driver:

```
[root]# yum install linuxrdac
```

## Configuring Nodes for Storage Access

To ensure proper operation of Red Hat Cluster Suite in the PowerEdge Cluster SE600L system, configure each node to access the same PowerVault MD3000 storage array. If your nodes were installed with the SE600L script or you completed these steps from the "Dell Community Repositories" on page 40, skip to "Installing the PowerVault Modular Disk Storage Manager" on page 45. To configure each node to access the same storage array:

- 1 Install and configure HBA driver.
- 2 Install and configure Multi-Path Proxy driver.
- 3 Configure the PowerVault MD3000 storage array.

After you complete these steps, each node accesses to the same shared storage. For more information, see the **Readme.txt** on the *Dell PowerVault MD3000 Resource Media*.

### Using the Dell PowerVault MD3000 Resource Media

The SAS HBA driver and Multi-Path Proxy driver are available on the *Dell PowerVault MD3000 Resource Media*. If you are not using the Dell Community repositories, or do not have at least version 14 media, download the latest ISO image from the Dell Support website at [support.dell.com](http://support.dell.com).

#### Pre-requisites

Before running the *Dell PowerVault MD3000 Resource Media* installer, ensure that the packages libXp, libXtst, and kernel development packages installed on each node that are needed for the installer.

- 1 Run the following commands to ensure that the packages are installed:  

```
[root]# yum install libXp libXtst
```
- 2 The *Dell PowerVault MD3000 Resource Media* requires development tools in order to install the Dell HBA drivers and Multi-Path Proxy driver. Before running the installer, ensure your node has these packages by entering the following commands:  

```
[root]# yum install kernel-devel gcc
```

- 3 To use the PowerVault Modular Disk Storage Manager ensure that X is running. If X is not running, execute the following commands:

```
[root]# yum groupinstall base-x
```

If X does not start up properly, run the following configuration utility:

```
[root]# system-config-display
```

Enter runlevel 5 with the following command:

```
[root]# init 5
```

Or start X with the following command:

```
[root]# startx
```

- 4 To view the various documentation located on the *Dell PowerVault MD3000 Resource Media* install a web browser on your system:

```
[root]# yum install firefox
```

### Launching the Main Menu

To install the required software:

- 1 Access the *Dell PowerVault MD3000 Resource Media* in one of the two ways:
  - Insert the media into a cluster node. If X is running, the media mounts automatically under the **/media** mount point and launches the main menu. If the media does not mount automatically, execute the following command:

```
[root]# mount /dev/cdrom {/path/to/mount/point}
```

where `{/path/to/mount/point}` is a directory that is used to access the contents.
  - If you have downloaded the ISO image, it is not necessary to create physical media. Mount the ISO image and access the contents with the following command:

```
[root]# mount -o loop {ISO filename}
{/path/to/mount_point} where
{/path/to/mount_point} is a directory that will be used to
access the contents.
```

- 2** The installer launches automatically. If the installer did not launch automatically, run the *Dell PowerVault MD3000 Resource Media Linux* installation script manually:

```
[root]# ./autorun
```

If the main menu does not appear, change the directory using the following command:

```
[root]# cd {/path/to/mount/point/linux/}
```

Run the installer:

```
[root]# ./install.sh
```

The **Dell PowerVault MD3000 Resource Media** menu is displayed.

```
#####
```

```
Dell PowerVault MD3000 Resource CD
```

```
#####
```

1. View MD3000 Readme
2. Install SAS 5/E Adapter Driver
3. Install MD3000 Storage Manager
4. Install Multi-Pathing driver
5. Install MD3000 Documentation
6. View MD3000 Documentation
7. Dell Support
8. View End User License Agreement

Enter the number to select a component from the above list.

Enter q to quit.

Enter:

- 3** You must complete the steps in the following sections to install the SAS HBA drivers and Multi-Path Proxy driver.

## Installing the SAS 5/E HBA Drivers



**NOTE:** The SAS 5/E driver is provided in the form of a Dynamic Kernel Module Support (DKMS) package. For more information on DKMS, visit the project page at [www.linux.dell.com/dkms](http://www.linux.dell.com/dkms).



**NOTE:** Perform the following steps on your cluster nodes only. Do not install the SAS HBA driver on any dedicated management node.

- 1 Complete the steps in "Launching the Main Menu" on page 43.
- 2 Select **Option 2. Install SAS 5/E Adapter Driver**. Follow the instructions on the screen to complete installation. The following message appears:

```
DKMS: install Completed
```

```
You must restart your computer for the new  
settings to take effect.
```

For more information view documentation on the main menu.

## Installing the PowerVault Modular Disk Storage Manager



**NOTE:** Perform the following steps on any management nodes. If you are using in-band management, then perform the steps in this section on any nodes that also manage the PowerVault MD3000 storage array.

The PowerVault Modular Disk Storage Manager contains the Multi-Path Proxy driver needed to configure your nodes. It also contains the necessary management software to manage and configure the PowerVault MD3000 storage array:

- 1 Complete the steps in "Launching the Main Menu" on page 43.
- 2 Select **Option 3. Install MD3000 Storage Manager**.
- 3 The Java-based installer appears.



**NOTE:** If you encounter any errors or do not see the Java installer, ensure that you have completed the steps in "Pre-requisites" on page 42. For more information, go to the main menu and view documentation.

- 4 Follow the instructions on the screen. To continue with installation, accept the license terms.
- 5 Select an installation location. The default value is `/opt/dell/mdstoragemanager`.

**6** The **Select an Installation Type** window appears and displays the following options:

- \* **Typical (Full)**
- \* **Management Station**
- \* **Host**

If you are using in-band management and do not have a dedicated management node, then select **Typical (Full)** for your cluster nodes. If you are using out-of-band management on a dedicated management node, then select **Management station** for the management node, and **Host** for the cluster nodes.

For more information on out-of-band and in-band management, see "Selecting a Management Infrastructure" on page 29.



**NOTE:** If you choose **Typical (Full)** or **Host**, the Multi-Path Proxy driver warning is displayed. It is safe to ignore this warning. For more information, see "Installing the Redundant Disk Array Controller (RDAC) Multi-Path Proxy Driver" on page 46.

### **Installing the Redundant Disk Array Controller (RDAC) Multi-Path Proxy Driver**

Each cluster node has multiple paths to the same virtual disk through SAS cables. This may cause issues as the cluster nodes are able to access the same data through different paths that could lead to data corruption. Multi-Path Proxy driver consolidates all paths to the virtual disk into a pseudo device. This allows the nodes to access the pseudo device instead of a path directly. If a path failure occurs, the Multi-Path Proxy driver automatically switches paths, and the node continues to access the same data through the same pseudo device.

Red Hat Enterprise Linux 5.1 Advanced Platform includes Multi-Path Proxy driver as part of the base operating system with the package `device-mapper-multipath`. However, Dell provides Multi-Path Proxy driver that is specific to the PowerVault MD3000 storage array. It is recommended that you use the Dell-specific RDAC Multi-Path Proxy (MPP) driver with the PowerEdge Cluster SE600L system.

You can now safely connect your SAS cables to your nodes if you disconnected them during the procedure in "Preparing the Operating System" on page 35.



**NOTE:** For more information on cabling your cluster in a redundant setup, see "Cabling Your Cluster" on page 17.



**NOTE:** For more information on device-mapper-multipath, see *Using Device-Mapper Multipath* located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

To install the RDAC MPP driver:

- 1 Complete the steps in "Launching the Main Menu" on page 43.
- 2 Select **4. Install Multi-Pathing driver**
- 3 Follow the steps on the screen to complete installation. After completion, the follow message appears:

```
DKMS: install Completed
```

```
You must restart your computer for the new  
settings to take effect.
```

For additional information about installing RDAC MPP driver, see the `RDACreadme.txt` located on the *Dell PowerVault MD3000 Resource* media.

## Installing and Configuring the Shared Storage System

This section provides information for installing and configuring the shared storage systems.



**NOTE:** For the latest PowerVault MD3000 firmware, see the Dell Support website at [support.dell.com](http://support.dell.com). Use the PowerVault Modular Disk Storage Manager to perform the firmware update. See the section *Installing the PowerVault Modular Disk Storage Manager* and the documentation that came with your storage array for more information.

### Setting Up Your Storage Array

The **Perform Initial Setup Tasks** link located on the **Summary** tab provides links to the basic steps you should follow when initially setting up a storage array in PowerVault Modular Disk Storage Manager.

Initial setup tasks include:

- 1 Blinking the Storage Array** — Find the physical location of the storage array on your network. The storage array can then be identified with a label.
- 2 Renaming the Storage Array** — Provide a unique and memorable name to help you easily identify the storage array.
- 3 Setting a Storage Array Password** — Prevent unauthorized manipulation of the storage array, such as deletion of a virtual disk.
- 4 Setting up Alert Notifications** — Enable e-mail and SNMP alerts to notify administrators about storage array conditions that require attention.
  - a Configure Sender E-mail Settings** — Provide the SMTP, e-mail address, and contact information that PowerVault Modular Disk Storage Manager uses to send e-mail alerts.
  - b Add or Edit E-mail Addresses** — Provide information about accounts that should receive e-mail-based alerts.
  - c Set Up SNMP Alerts** — Provide information about hosts that should receive SNMP-based alerts.
- 5 Configuring Host Access and Create a Host Group** — Set up all the cluster nodes to access the same virtual disk in the storage array. See "Configuring Host Access" on page 48 and "Creating a Host Group" on page 49 for more information.
- 6 Configuring and Manage Virtual Disks** — See "Creating Disk Groups and Virtual Disks" on page 50 for more information.
- 7 Changing Network Configuration (Optional)** — Change your network configuration by changing RAID controller network settings or obtain the network configuration from a DHCP server. Ensure that your storage array is on the same private network as the management node for out-of-band management.

## Configuring Host Access

Ensure that the PowerVault Modular Disk Storage Manager Agent service is running on your cluster nodes.

Configuring host access allows you to permit or deny access to a storage array for specific hosts. On the PowerEdge Cluster SE600L system, configure host groups to allow your nodes to access the same virtual disk(s).



When you permit host access, that host can then be mapped to a virtual disk on the storage array. On the **Summary** tab, the **Hosts & Mappings** area indicates the number of hosts that are configured to access the array. Click the **Configured Hosts** link in this area to see the names of these hosts.

Host access configuration is the first step in setting up your storage array. You must complete this task during initial setup and anytime you connect a new host.

To begin configuring host access, click the **Configure** tab and then click the **Configure Host Access** link. PowerVault Modular Disk Storage Manager scans the array and displays a list of the hosts that have not yet been configured for access to the array. Click the **View Hosts that currently have access to the storage array** link to see hosts that have already been configured.

To automatically configure a host for access to the storage array:

- 1 Click the **Configure** tab and then click the **Configure Host Access** link.
- 2 Select both the cluster nodes individually, or by clicking the **Select All** check box beneath the list.
- 3 Set the Host Type as **Linux** for all the HBA ports on each cluster node. Click the **View Details** button (next to the list) and choose the appropriate host type. Use the SAS Address recorded in the step "Configuring the Cluster Nodes" on page 29 to match SAS HBA to node.
- 4 Click **OK** to configure access to the array for the hosts you selected.

## Creating a Host Group

- 1 Click the **Modify** tab and then click the **Modify Host Topology** link.
- 2 Click the **Create Host Group** link on the **Modify Host Topology** window. The **Create Host Group** window appears.
- 3 Type a name for the new host group in the text box.
- 4 In the **Select Hosts to Add** list, click the names of the first cluster node, then click the **Add** button located to the right of the list. The host moves to the **Hosts in Group List**.



**NOTE:** If the cluster node name does not appear, ensure that the SAS address is the same value that was configured in "Configuring the Cluster Nodes" on page 29.

- 5 Repeat step 4 to add the additional cluster nodes to the host group.
- 6 To create the host group, click **OK**.

## Creating Disk Groups and Virtual Disks

In some cases, the virtual disks may have been bound when the system was shipped. It is still important, however, to install the management software and to verify that the desired virtual disk configuration exists.

You can manage your virtual disks remotely using PowerVault Modular Disk Storage Manager.

Disk groups are created in the non-configured capacity of a storage array, and virtual disks are created in the free capacity of a disk group. The nodes attached to the storage array read and write data to the virtual disks.



**NOTE:** Before you create virtual disks, you must first organize the physical disks into disk groups and configure host access. You can then create virtual disks within a disk group. Only 30 physical drives are supported in a disk group.

To create a virtual disk, use one of the following methods:

- Automatic Configuration
- Manual Configuration

For more information on how to create Disk Groups and Virtual Disks, see your PowerVault Modular Disk Storage Manager documentation.



**NOTE:** It is recommended that you use a RAID level other than RAID 0 (which is commonly called striping). RAID 0 configurations provide higher performance, but do not provide any fault tolerance. For more information about setting up RAID levels for the system, see the documentation for your storage system.

## Creating Host-to-Virtual Disk Mappings

Create host-to-virtual disk mappings to assign virtual disks to the host groups containing cluster nodes:

- 1 Click the **Configure** tab and select the **Create Host-to-Virtual Disk Mappings** link.

When you click this link, the PowerVault Modular Disk Storage Manager displays a series of pages in which you select the Host Group containing the cluster nodes and virtual disks to be mapped.

- 2 After you complete this configuration, verify the mapping by clicking the **Host-to-Virtual Disk Mappings** link on the **Summary** tab to ensure that the configuration was created correctly.

# Configuring Your Cluster

This section discusses how to install and configure Red Hat® Cluster Suite and Global File System on your Dell™ PowerEdge™ cluster SE600L system using Conga and the CLI Tools.

Conga is a newer configuration and management suite based on a server/agent model. You can access the management server (**luci**) using a standard web browser from anywhere on the network. Conga communicates through the client agent (**ricci**) on the nodes and installs all required packages, synchronizes the configuration file, and manages the storage cluster. Though there are many possible methods, including the command line interface, it is recommended that you use Conga to configure and monitor your cluster.

## Setting Up a High-Availability Cluster

The following sections provide information about installing and configuring your high-availability Red Hat® cluster using Conga.

For more information on using Conga, see the section *Configuring Red Hat Cluster With Conga* in the *Cluster Administration* guide at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

### Preparing Cluster Nodes for Conga

The SE600L script or the `groupinstall` command run earlier would have installed the correct packages. However, you can use the following sections to verify that the correct software is installed.

To prepare the cluster nodes:

- 1** Run the following command on each cluster node to install the Conga client agent **ricci**:
  - a** Install **ricci**.

```
[root]# yum install ricci
```
  - b** Configure **ricci** to start on boot

```
[root]# chkconfig ricci on
```

- c** Start the `ricci` agent

```
[root]# service ricci start
```
- 2** Run the following command on the management node to install the Conga server `luci`. On the management node:
  - a** Install the Conga server agent `luci`:

```
[root]# yum install luci
```



**NOTE:** You can configure the `luci` server on any node, but it is recommended that you install `luci` on a dedicated management node. If you do have a dedicated management node, that node requires to have access to the Cluster group. You can also login to RHN and manually download `luci` for installation on your management node. For any issues, refer to "Troubleshooting Conga" on page 69.

- b** Initialize the `luci` server agent:

```
[root]# luci_admin init
```
- c** Configure the `luci` server agent to start on boot:

```
[root]# chkconfig luci on
```
- d** Start the `luci` server agent:

```
[root]# service luci start
```

For more information on configuring your cluster with Conga, see the section "Configuring Red Hat Cluster With Conga" on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise) or locally in the `Cluster_Administration-en-US` package on each node.


## Creating Your Cluster Using Conga

Conga automatically installs the software required for clustering on all cluster nodes. Ensure that you have completed the steps in "Preparing Cluster Nodes for Conga" on page 51 before proceeding.


- 1 Connect to the **luci** server agent from any browser on the same network as the management node. In your web browser, type:

```
https://{management node hostname or ip  
address}:8084
```

Where *{management node hostname or ip address}* is the hostname or IP address of the optional dedicated management node, or the cluster node running the **luci** server agent.

 **NOTE:** If you encounter any errors, see section "Troubleshooting Conga" on page 69.

- 2 Enter your username and password to securely log in to the **luci** server agent.
- 3 Go to the **cluster** tab.
- 4 Click **Create a New Cluster**.
- 5 Enter a cluster name of 15 characters or less.
- 6 Add the fully qualified private hostname or IP address and root password for each cluster node.

 **NOTE:** The password is sent in an encrypted SSL session, and not saved.
- 7 Ensure that the option for **Enable Shared Storage Support** is selected and click **Submit**.

Conga downloads and installs all the required cluster software, creates a configuration file, and reboots each cluster node. Watch the Conga status window for details on the progress of each cluster node

## Configuring Fencing Using Conga

In your PowerEdge Cluster SE600L system, the network power switches provide the most reliable fencing method. For more information on configuring your network power switches for remote access, see the documentation that came with your network power switches.

Configure your network power switches and remote access controllers such as Dell Remote Access Controller (DRAC) or Intelligent Platform Management Interface (IPMI) on the same private network as the cluster nodes. Refer to the section "Configuring Remote Access" on page 32 for more information.

Remote access controllers such as DRAC or IPMI should be used as secondary methods, unless no network power switches are available. In this case configure a secondary method as **manual**. Use a log watching utility to notify you if your primary fencing method is failing. For more information, see "Testing Fencing Mechanisms" on page 72 and see the section "Fencing" in the "Cluster Suite Overview" located on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

To configure fencing:

- 1 Go to **Cluster** → **Nodes**.
- 2 Select one of the cluster nodes.
- 3 Under the section for **Main Fencing Method** section, click **Add**.
- 4 Configure primary and secondary fencing.



**NOTE:** If you have optional network power switches, they will be configured as shared devices.

## Setting Up a Storage Cluster

This section describes the procedure to set up a Global File System (GFS) that is shared between the cluster nodes. Verify that the high-availability Red Hat cluster is running before setting up the storage cluster. The PowerEdge Cluster SE600L system comprises of a Red Hat Cluster Suite high-availability cluster and GFS storage cluster.



**NOTE:** Before setting up the storage, complete the steps in "Installing the Redundant Disk Array Controller (RDAC) Multi-Path Proxy Driver" on page 46.

Configuring shared storage consists of the following steps:

- 1 Configuring a Clustered Logical Volume (CLVM).
- 2 Configuring Global File System (GFS).

For more information, see the *LVM Administrator's Guide* and *Global File System* on [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

The procedure for configuring a Storage Cluster is documented using both Conga and CLI tools. You can use any one of the methods.

## Configuring a Storage Cluster With Conga

To configure a storage cluster with Conga:

- 1 Log in to Conga.
- 2 Click the **Storage** tab.
- 3 In the **System list** section, select a node. The hard drives that are visible to the node are displayed. Repeat this step for the second node. The same hard-drive list must be displayed. If both nodes cannot access the same hard drives, see the PowerVault MD3000 storage documentation to ensure both nodes are viewing the same virtual disk.

## Partitioning Shared Storage With Conga



**NOTICE:** Running the commands in the following steps erases all partitions and data from your hard drive



**NOTE:** Run all commands from one cluster node only. All nodes have access to the same virtual disk and the following steps must be performed only one time.

- 1 In the **system list** section, select a node.
- 2 Select **partition tables**.
- 3 Select **New Partition Table**. For **label** ensure the option **GPT** is selected. On the right pane, check the box for your virtual disk and click **create**.
- 4 If your virtual disk is not listed, follow the steps in "Shared Storage Issues" on page 71
- 5 Select **Unused Space**.
- 6 On the left page, select **System List** then the other node and view the partitions it has access to. If the data displayed is inconsistent, click **reprobe storage**. It may be necessary to run the command **partprobe** on both nodes to ensure consistent view. See "Shared Storage Issues" on page 71 for more information.

To create a data partition:

- 7 In the **System List** section, select a node.
- 8 Select **Partition Tables**.

- 9 Select **New Partition Table**. For **Label** ensure the option **GPT** is selected. On the right pane, check the box for your virtual disk and click **Create**. If your virtual disk is not listed, follow the steps in "Shared Storage Issues" on page 71.
- 10 Select **Unused Space**.
- 11 Do not change the size that is displayed. Select the following values:
  - Partition Type—Primary
  - Content—Do not enter any value
- 12 On the left page, select **System List** then the other node and view the partitions it has access to. If the data displayed is inconsistent, click **Reprobe Storage**. It may be necessary to run the command **Partprobe** on both nodes to ensure consistent view. See "Shared Storage Issues" on page 71 for more information.

### Creating Clustered Logical Volume With Conga


- 1 On the left-pane, click **Volume Groups** and click **New Volume Group**.
- 2 In Volume Group Name enter a name that identifies your shared storage volume group. For example, `clustered_vg`.
- 3 Ensure that the option **Clustered** is set to **True**.
- 4 In the right-pane, select the newly-created data partition. In this example, the partition is `/dev/sdb1`.
- 5 Click **Create**.

The new volume group is displayed. Verify that the volume group size is correct.
- 6 Click **New Logical Volume**. Enter a value in the **Logical Volume Name** field to identify your shared storage logical volume. For example, `clustered_lv`.
- 7 Select a size. All the available space is used by default, however you can create several logical volumes to meet your specific needs.

### Creating Global File System With Conga

- 1 In the **Content** field, select **GFS1 - Global FS v.1**.
- 2 The GFS entry dialogue appears.



- 3 Verify that the value in the **Cluster Name** field is the same as the value listed in the **Cluster** tab.
- 4 Enter a unique **GFS Name**. You do not have to specify a mount point or list the mount point in `/etc/fstab`.
- 5 In **Number of Journals**, enter the number of cluster nodes plus one. In this example, the value is **3**.
- 6 Verify that the clustered value is set to **True**.  
Do not change any other values unless you need to customize the cluster as per your requirements.  
 **NOTE:** If any errors are displayed, see "Shared Storage Issues" on page 71.
- 7 Skip to "Managing Services" on page 64.

## Configuring a Storage Cluster With CLI Tools

The following is an example of creating a clustered logical volume for use with GFS and a PowerVault MD3000 storage array using Multi-Path Proxy driver using CLI commands. For instructions on using Conga, see "Configuring a Storage Cluster With Conga" on page 55.

### Partitioning Shared Storage With CLI Tools

To partition shared storage:

- 1 Verify that the devices on each node are consistent by executing the following command:

```
[root]# cat /proc/partitions
```

All devices must be accessible from both nodes.



**NOTICE:** Executing the command in step 2 erases all partitions and data from your hard drive.



**NOTE:** Run all commands from one cluster node only. All nodes have access to the same virtual disk and the following steps must be performed one time only.

- 2 From one node only, set up shared storage partitioning. Create a partition by executing the following command. For Red Hat® Cluster Suite and Global File System, it is typical for one partition to consume the entire disk.

```
[root@node1 ~]# parted {your mpp device}
```

Where *{your mpp device}* is the mpp pseudo device that points to the shared virtual disk. If you are unsure which device is your virtual disk, follow the steps in "Determining Virtual Disk" on page 70. For example using the command:

```
[root@node1 ~]# parted /dev/sdab
```

The following screen is displayed:

```
GNU Parted 1.8.1
```

```
Using /dev/sdb
```

```
Welcome to GNU Parted! Type 'help' to view a  
list of commands.
```

- 3 Change the partition type to GPT. The following screen is displayed.

```
(parted) mklabel gpt
```

View the disk information with the print command:

```
(parted) print
```

```
Model: DELL MD Virtual Disk (scsi)
```


```
Disk /dev/sdb: 1000GB
```

```
Disk geometry for /dev/sdb 0kB - 1000GB
```

```
Disk label type: gpt
```

- 4 Create a single partition that uses the entire disk:

```
(parted) mkpart primary 0 -1
```

 **NOTE:** The value "-1" uses the entire disk.

- 5 Set the `lvm` flag for the data partition. Use the `set` command:

```
(parted) set 1 lvm on
```

- 6 Verify that the partitions are set up properly:

```
(parted) print
```

```
Disk geometry for /dev/sdb: 0kB - 1000GB
```

```
For Model: DELL MD Virtual Disk (scsi)
```

```
Disk /dev/sdb: 1000GB
```

```
Disk label type: gpt
```

```
Number Start End Size FileSystemName Flags
```

```
1          17.4kB 1000MB 1000MB primary lvm
```

- 7 Quit `parted`. The changes are saved immediately.

```
(parted) quit
```

- 8 Perform the following step on both nodes. Display devices and verify the partitions are visible to the kernel.

```
cat /proc/partitions
```

- 9 If the partitions are not visible, run the following command to update the kernel:


```
[root]# partprobe {your mpp device}
```

Where *{your mpp device}* is the newly-created partition. For example:


```
[root]# partprobe /dev/sdb
```

To view partitions again:

```
[root]# cat /proc/partitions
```

 **NOTE:** If the newly-created partition is not displayed after executing the `partprobe` command, reboot the system.

## Creating Cluster Logical Volumes With CLI Tools

 **NOTE:** Ensure that Red Hat Cluster Suite is functioning before you create a cluster logical volume.

To create a logical volume, run the following commands in series:

- 1 Create a physical volume (pv):

```
[root@node1 ~]# pvcreate {your mpp device partition}
```

Where *{your mpp device partition}* is the newly-created partition. For example:

```
[root@node1 ~]# pvcreate /dev/sdb1
```

The message below is displayed:

```
Physical volume "/dev/sdb1" successfully created.
```

- 2 Create a volume group (vg):

```
[root@node1 ~]# vgcreate {volume group name} {your mpp device partition}
```

Where *{volume group name}* is a name of your choice to identify this volume group, and *{your mpp device partition}* is the data partition you created. For example:

```
[root@node1 ~]# vgcreate clustered_vg /dev/sdb1
```

The following message is displayed:

```
Volume group clustered_vg successfully created
```

### 3 Create a logical volume (lv):

```
[root@node1 ~]# lvcreate {size of logical volume}
-n {logical volume name} {volume group name}
```

Where *{size of logical volume}* is one of the following options:

- l 100%FREE (consumes all available space)
- l {physical extents} (see man `lvcreate` for more info)
- L {size} (see man `lvcreate` for more info)

You can create several logical volumes based on your specific applications needs, or a single one that consumes all available space.

Lastly where *{logical volume name}* is a name of your choice. For example:

```
[root@node1 ~]# lvcreate -l 100%FREE -n
clustered_lv clustered_vg
```

The message below is displayed:

```
Logical volume "clustered_lv" created
```

### 4 Verify the lv was created successfully:

```
[root@node1 ~]# lvdisplay
/dev/clustered_vg/clustered_lv
```

The following screen is displayed:

```
--- Logical volume ---
LV Name  /dev/clustered_vg/clustered_lv
VG Name  clustered_vg
LV UUID  zcdt4M-VI2U-19bA-h2mJ-dmfo-
fFBE-ub3RYk
LV Write Access  read/write
LV Status        available
# open           0
LV Size          1020.00 MB
Current LE       255
```

```
Segments                1
Allocation               inherit
Read ahead sectors      0
Block device             253:2
```

- 5 On both nodes, edit the file `/etc/lvm/lvm.conf`.

Search for **cluster**. The following entry is displayed:

```
# Type of locking to use. Defaults to local file-
based locking (1).
```

```
# Turn locking off by setting to 0 (dangerous:
risks metadata corruption
```

```
# if LVM2 commands get run concurrently).
```

```
# Type 2 uses the external shared library
locking_library.
```

```
# Type 3 uses built-in clustered locking.
```

```
locking_type = 1
```

Change the `locking_type` to `3` to enable clustered lvm support, as indicated in the comments:

```
locking_type = 3
```

- 6 Start the clustered lvm daemon (`clvmd`) by typing:

```
service clvmd start
```

Ensure that `clvmd` starts on boot:

```
chkconfig clvmd on
```

## Creating Global File System Using CLI Tools



**NOTE:** For more information on configuring GFS, see the *Global File System Guide* on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise) or the package `Global_File_System-en-US` locally on each node. For instructions on using Conga to configure GFS, see the section "Configuring a Storage Cluster With Conga" on page 55.

Before you create a GFS:

- Verify that Red Hat Cluster Suite is functioning properly and ensure that a clustered logical volume is created.
- Gather the following information:
  - {cluster name}—is listed in Conga in the **Cluster** tab, and is also the `/etc/cluster/cluster.conf` file in each node.
  - {gfs filesystem name}—A unique identifier among GFS file systems.
  - {number of journals}—One journal per node is required. However, at least one additional journal is recommended.
  - {name of clustered logical volume}—The cluster logical volume that you created. For example  
`/dev/clustered_vg/clustered_lv`.

To create a file system:

- 1 From one system only run the following command:

```
[root@node1 ~]# gfs_mkfs -t {cluster name}:{gfs
filesystem name} -p lock_dlm -j {number of
journals} {name of clustered logical volume}
```

For example:

```
[root@node1 ~]# gfs_mkfs -t SE600L:MD3000_gfs -p
lock_dlm -j 3 /dev/clustered_vg/clustered_lv
```

The following warning message is displayed:

```
This will destroy any data on
/dev/clustered_vg/clustered_lv.
```

```
Are you sure you want to proceed? [y/n]
```

- 2 Type `<y>`.
- 3 A confirmation message is displayed indicating the creation of a file system.

```
Device:
/dev/clustered_vg/clustered_lv
```

```
Blocksize: 4096
```

```
Filesystem Size: 162772
```

Journals: 3  
Resource Groups: 8  
Locking Protocol: lock\_dlm  
Lock Table: SE600L:MD3000\_gfs  
Syncing...  
All Done

## Managing Services

This section describes the procedure to create and test services on your PowerEdge Cluster SE600L system.

### Creating Services

The following sets provide an overview of procedure to create services:

- 1 Click **Cluster List**.
- 2 Select the cluster name and click **Services** on the left-pane.
- 3 Select **Add a service**.
- 4 Choose a service name that describes the function that you are creating the service for.
- 5 Click **Add a resource to this service**.
- 6 Select either **New Local** or **Existing Global Resource**.
- 7 Select the type of resource you need to create and enter in the required data in the appropriate fields.



**NOTE:** For more information, see *Adding a Cluster Service to the Cluster* located on the Red Hat website at [www.redhat.com](http://www.redhat.com).



## Example of Configuring NFS

Use Conga to configure an NFS service:

- 1 Click **Resources** and add the following resources:
  - IP Address—IP address to be used as a virtual IP address by the NFS service
  - GFS File System—the clustered logical volume
  - NFS Export—name for this export
  - NFS Client—the target network client options
- 2 Click **Services** and enter a service name.
- 3 Click **Add a resource to this service**.
- 4 From the drop-down menu, choose **Use an existing global resource**.
- 5 Select the resource **IP Address** that you created in step 1.
- 6 Add a child resource to the newly-created **IP Address**. Choose **GFS File System** created in step 1.
- 7 Add a child resource to the newly-created **GFS File System**. Choose **NFS Export** created in step 1.
- 8 Add a child resource to the newly-created **NFS Export**. Choose **NFS Client** created in step 1.

This process creates a dependency check among resources. Before attempting to start, all child resources wait for the parent resource they are associated with. For example, this ensures the NFS service does not try to start if there is no mounted file system, and so on.

You can manage services from Conga by clicking on **Cluster** then **Services** and choosing an action.

You can also use the CLI to manage services. Use the following command:

```
[root]# clusvcadm
```

For example, to relocate a service from node 1 to node 2, enter the following command:

```
[root]# clusvcadm -r service_name node2
```

Use the command **clustat** to view cluster status:

```
[root]# clustat
```



# Troubleshooting

This section describes various methods to verify and identify possible issues, and test the high-availability features of your Dell™ PowerEdge™ Cluster SE600L system.

## Physical Connectivity

Check all your physical connections. Disconnect each cable and inspect it for damage, then connect it firmly. If problems persist, try to swap the cable with another one, and see if the issue is with the cable, or is a problem with the device.

## Verifying Node Connectivity

Perform the following to verify connectivity:

On node1:

```
[root@node1]# ping {node2 fully qualified hostname}
```

For example:

```
[root@node1]# ping node2.local
```

On node2:

```
[root@node2]# ping {node1 fully qualified hostname}
```

On the management node:

```
[root@management]# ping {node1 fully qualified  
hostname}
```

```
[root@management]# ping {node2 fully qualified  
hostname}
```

## Networking Issues

Red Hat® Clustering nodes use Multicast to communicate. Your switches must be configured to enable multicast addresses and support the Internet Group Management Protocol (IGMP). See the Cluster Administration guide in section "2.6. Multicast Addresses" on [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise) for more information, and the documentation that came with your switches.

## Cluster Status

Conga will allow you to monitor your cluster. Alternatively, you may run the command `clustat` from any node. For example:

```
[root]# clustat
```

### Logging

Any important messages are logged to `/var/log/messages`.

The following is an example of loss of network connectivity on node1 which causes node2 to fence it.

```
Nov 28 15:37:56 node2 openais[3450]: [TOTEM] previous ring seq 24 rep 172.16.0.1
```

```
Nov 28 15:37:56 node2 openais[3450]: [TOTEM] aru 2d high delivered 2d received flag 1
```

```
Nov 28 15:37:56 node2 openais[3450]: [TOTEM] Did not need to originate any messages in recovery
```

```
.
```

```
Nov 28 15:37:56 node2 openais[3450]: [TOTEM] Sending initial ORF token
```

```
Nov 28 15:37:56 node2 openais[3450]: [CLM ] CLM CONFIGURATION CHANGE
```

```
Nov 28 15:37:56 node2 openais[3450]: [CLM ] New Configuration:
```

```
Nov 28 15:37:56 node2 kernel: dlm: closing connection to node 2
```

```
Nov 28 15:37:56 node2 fenced[3466]: node1.ha.lab not a cluster member after 0 sec post_fail_delay
```

```
Nov 28 15:37:56 node2 openais[3450]: [CLM ] r(0) ip(172.16.0.2)
```

```
Nov 28 15:37:56 node2 fenced[3466]: fencing node "node1.example.com"
```

```
Nov 28 15:37:56 node2 openais[3450]: [CLM ] Members Left:
```

```
Nov 28 15:37:56 node2 openais[3450]: [CLM ] r(0) ip(172.16.0.1)
```

## Troubleshooting Conga

The following sections describe issues you may encounter while creating the cluster initially and the possible work around.

### Running luci on a Cluster Node

If you are using a cluster node also as a management node and running **luci**, you have to restart luci manually after the initial configuration. For example:

```
[root]# service luci restart
```

### Issues While Creating a Cluster Initially

If the following error appears when initially installing the cluster:

The following errors occurred:

```
Unable to add the key for node node1.ha.lab to the trusted keys list.
```

```
Unable to add the key for node node2.ha.lab to the trusted keys list.
```

```
Unable to connect to node2.ha.lab: Unable to establish an SSL connection to node2.ha.lab:11111:
```

```
ClientSocket(hostname, port, timeout): connect() failed
```

```
Unable to connect to node1.ha.lab: Unable to establish
an SSL connection to node1.ha.lab:11111:
ClientSocket(hostname, port, timeout): connect()
failed
```

This error occurs when the **luci** server cannot communicate with the **ricci** agent. Verify that **ricci** is installed and started on each node. Ensure that the firewall has been configured correctly, and that Security-Enhanced Linux (SELinux) is not the issue. Check `/var/log/audit/audit.log` for details on SELinux issues. Make sure your nodes have the latest SELinux policy with the following command:

```
[root]# yum update selinux-policy
```

If you continue to encounter errors, it may be necessary to disable SELinux. This is not recommended, and should only be used as a last resort. Disable SELinux with the command:

```
[root]# setenforce 0
```

See "Security and SELinux" in the *Deployment Guide* on [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

## Configuration File Issues

Configuration errors manifest themselves as the following error in `/var/log/messages`:

```
"AIS Executive exiting (-9)"
```

Check for syntax errors in your `/etc/cluster/cluster.conf` file. This is unlikely to happen if you are using Conga to manage your cluster configuration file.

## Determining Virtual Disk

If you are unable to determine the correct corresponding SCSI device (sd) that maps to your virtual disk use this command:

```
[root]# grep -A 5 "MD Virtual" /var/log/dmmsg
```

The last few lines displayed will be the virtual disk. For example:

```
Vendor: DELL
```

```
Model: MD Virtual Disk Rev: 0670
```

Type: Direct-Access

ANSI SCSI revision: 05

```
scsi(5:0:0:0): Enabled tagged queuing, queue depth 62.  
SCSI device sdb: 2032423936 512-byte hdwr  
sector(1040601 MB)  
sdb: Write Protect is off
```

It is also possible to match the World Wide Name from various files located in `/var/mpp`, `/proc/mpp`, and `/proc/scsi/mpp`. For more information, see the documentation that came with your storage array.

## Logical Volume Issues

It may be necessary to restart the clustered logical volume manager with the command:

```
[root]# service clvmd restart
```

Ensure that all nodes have a consistent view of the shared storage with the command `partprobe` or clicking `reprobe storage` in Conga. As a last resort, reboot all nodes, or select `restart cluster` in Conga.

## Shared Storage Issues

If you are experiencing errors when creating the clustered logical volume, you may need to wipe any previous labels from the virtual disk.



**NOTICE:** This will destroy all data on the virtual disk.

Run the following command from one node:

```
[root@node1 ~]# pvremove -ff {/dev/sdXY}
```

Where `{/dev/sdXY}` is the partition intended for data. See the output of `/proc/mpp` to verify. For example:

```
[root@node1 ~]# pvremove -ff /dev/sdb1
```

If you are using Conga, click `reprobe storage`, otherwise type:

```
[root@node1 ~]# partprobe /dev/sdb
```

If you have imaged the nodes with a cloning method, then the unique identifier (uuid) for the system logical volumes may be the same. It may be necessary to change the uuid with the commands `pvchange --uuid` or `vgchange --uuid`. For more information, see *LVM Administrator's Guide* on the Red Hat website at [www.redhat.com/docs/manuals/enterprise](http://www.redhat.com/docs/manuals/enterprise).

## Testing Fencing Mechanisms

Fence each node to ensure that fencing is working properly.

- 1 Watch the logs from node 1 with the following command:  

```
[root@node1]# tail -f /var/log/messages
```
- 2 Fence node 2 by executing the following command:  

```
[root@node1]# fence_node {fully qualified hostname  
or ip address of node2}
```
- 3 View the logs on node 1 and the console node 2. Node 1 should successfully fence node 2.
- 4 Continue to watch the messages file for status changes. You can also use the Cluster Status tool to see the cluster view of a node. The parameter `-i 2` refreshes the tool every two seconds. For more information on clusters see:  

```
[root]# clustat -i 2
```
- 5 After you successfully fence one node, repeat this process for the second node.



# A

## Checklist to Set Up Your Cluster

**Table A-1. Cluster Creation Checklist**

Item	Verified
Physical setup	
Physical cabling	
Remote Access Controllers configured in the BIOS	
DRAC telnet enabled (if applicable)	
Network Power Switches configured (if applicable)	
Operating Systems installed and updated	
Cluster and Cluster Storage	
SAS HBA driver (mptlinux package)	
Multi-Path Proxy driver (linuxrdac package)	
Dell™ PowerVault™ MD3000 storage array configured	
PowerVault MD3000 Storage Manager installed on a node	
Virtual Disk created	
Host Groups defined and mapped to Virtual Disks	
Red Hat® Cluster Suite installed and configured	
Nodes participating in cluster	
Fencing configured on nodes	
Clustered logical volume	
Global File System	
Services created	

