

Dell™ PowerVault™ MD3000i Storage
Arrays with Microsoft® Windows
Server® Failover Clusters

Hardware 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*, and *PowerVault* are trademarks of Dell Inc.; *Microsoft*, *Active Directory*, *Windows*, *Windows Server*, and *Windows NT* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or 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	Introduction	5
	Overview	5
	Cluster Solution	6
	Cluster Hardware requirements:	6
	Cluster Nodes	7
	Cluster Storage	8
	Cluster Storage Management Software	9
	Supported Cluster Configurations	11
	Other Documents You May Need	14
2	Cabling Your Cluster Hardware	17
	Cabling the Mouse, Keyboard, and Monitor	17
	Cabling the Power Supplies	17
	Cabling Your Public and Private Networks	20
	Cabling Your Public Network	21
	Cabling Your Private Network	21
	Using Dual-Port Network Adapters for Your Private Network	22
	NIC Teaming	22
	Cabling the Storage Systems	23
	Cabling the Cluster in a Non-Redundant Configuration	23

	Cabling the Cluster in Direct-Attached Redundant Configuration	25
	Cabling the Cluster in Network-Attached Redundant Configuration	27
	Connecting a PowerEdge Cluster to Multiple PowerVault MD3000i Storage Systems.	29
3	Preparing Your Systems for Clustering	31
	Cluster Configuration Overview	31
	Installation Overview	33
	Installing the iSCSI NICs	34
	Installing the Microsoft iSCSI Software Initiator.	35
	Installing and Configuring the Storage Management Software	35
	Installing and Configuring the Shared Storage System	37
	Installing and Configuring a Failover Cluster.	63
A	Troubleshooting	65
	Troubleshooting Tools	72
	Known Issues	74
B	Cluster Data Form	75
C	iSCSI Configuration Worksheet.	77
	Index	79

Introduction

This document provides information for installing and managing your Cluster solution using Dell™ PowerVault™ MD3000i storage system. It is intended for experienced IT professionals who need to configure the cluster solution, and for trained service technicians who perform upgrade and maintenance procedures. This document also addresses readers who are new to clustering.

Overview

A Microsoft® Windows Server® failover cluster combines specific hardware and software components to provide enhanced availability for applications and services that are run on the cluster. A failover cluster is designed to reduce the possibility of any single point of failure within the system that can cause the clustered applications or services to become unavailable. It is recommended that you use redundant components like system and storage power supplies, connections between the nodes and the storage array(s), connections to client systems, or other systems in the multi-tier enterprise application architecture in your cluster.

This guide addresses the configuration of your Dell MD3000i iSCSI storage array for use with one or more Windows Server failover clusters. It provides information and specific configuration tasks that enable you to deploy the shared storage for your cluster.

For more information on deploying your cluster with a specific variant of the Windows Server operating system (for example: Windows Server 2003 or Windows Server 2008), see *Dell Failover Clusters with Microsoft Windows Server Installation and Troubleshooting Guide*.

For a list of recommended operating systems, hardware components, and driver or firmware versions for your Dell Windows Server Failover Cluster, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at www.dell.com/ha.

Cluster Solution

Your iSCSI cluster implements a minimum of two node-clustering to a maximum of either eight nodes (for Windows Server 2003) or sixteen nodes (for Windows Server 2008) clustering and provides the following features:

- Internet Small Computer System Interface (iSCSI) technology
- High availability of system services and resources to network clients
- Redundant paths to the shared storage
- Failure recovery for applications and services
- Flexible maintenance capabilities, allowing you to repair, maintain, or upgrade a cluster node without taking the entire cluster offline

Implementing iSCSI technology in a cluster provides the following advantages:

- **Flexibility** - iSCSI as it is based on TCP/IP allows cluster nodes and storage systems to be located at different sites.
- **Availability** - iSCSI components use redundant connections, providing multiple data paths and greater availability for clients.
- **Connectivity** - iSCSI allows more device connections than SCSI. Because iSCSI devices are hot-pluggable, you can add or remove devices from the nodes without bringing down the cluster.

Cluster Hardware requirements:

Your cluster requires the following hardware components:

- Servers (Cluster nodes)
- Storage and storage management software

Cluster Nodes

Table 1-1 lists hardware requirements for the cluster docs.

Table 1-1. Cluster Node Requirements

Component	Minimum Requirement
Processor	At least one processor for each cluster node.
Cluster Nodes	A minimum of two identical Power Edge systems are required. The maximum number of nodes that is supported depends on the variant of the Windows Server operating system used in your cluster, and on the physical topology in which the storage system and nodes are interconnected.
RAM	The variant of the Windows Server operating system that is installed on your cluster nodes determines the minimum required amount of system RAM.
iSCSI Initiator	Complete installation of the iSCSI port driver, Initiator Service, and Software Initiator on each node. NOTE: Microsoft Multipath I/O (MPIO) Multipathing Support for iSCSI is not installed.
Network Interface Cards (NICs) for iSCSI access	Two iSCSI NICs or NIC ports per node. Place the NICs on separate PCI buses to improve availability and performance. TCP/IP Offload Engine (TOE) NICs are also supported for iSCSI traffic. For a list of recommended operating systems, hardware components, and driver or firmware versions for your Dell Windows Server Failover Cluster, see the <i>Dell Cluster Configuration Support Matrices</i> located on the Dell High Availability Clustering website at www.dell.com/ha .

Table 1-1. Cluster Node Requirements (continued)

Component	Minimum Requirement
NICs (public and private)	At least two NICs: one NIC for the public network and another NIC for the private network. NOTE: 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. Use any supported Redundant Array of Independent Disk (RAID) controller or disk controller. Two physical disks are required for mirroring (RAID 1) and at least three are required for disk striping with parity (RAID 5). NOTE: It is strongly recommended that you use hardware-based RAID or software-based disk-fault tolerance for the internal drives.

Cluster Storage

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

Table 1-2. Cluster Storage Requirements

Hardware Components	Minimum Requirement
Supported storage systems	One Dell PowerVault MD3000i RAID enclosure. Up to two Dell PowerVault MD1000 expansion enclosures.
Power and cooling requirements	Two integrated hot-pluggable power supply/cooling fan modules.
Physical disks	At least two physical disks in the PowerVault MD3000i RAID enclosure.
Multiple clusters and stand-alone systems	In a switched-attached configuration, clusters and stand-alone systems can share one or more PowerVault MD3000i systems.



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.

Cluster Storage Management Software

Modular Disk Storage Manager Client

The software runs on the management station to centrally manage the PowerVault MD3000i RAID enclosures. You can use Dell PowerVault Modular Disk Storage Manager to perform tasks such as creating or managing RAID arrays, binding virtual disks, and downloading firmware.

Modular Disk Storage Manager Agent

This software resides on each cluster node to collect system-based topology data that can be managed by the Modular Disk Storage Manager Client.

Multipath Software

Multipath I/O software (also referred to as the failover driver) is a software residing on each cluster node that provides management of the redundant data path between the system and the RAID enclosure. For the multipath software to correctly manage a redundant path, the configuration must provide for redundant NICs and cabling.

The multipath software 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 multipath software 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 by default. 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).

Advanced Features

Advanced features for the PowerVault MD3000i RAID enclosure include:

- **Snapshot Virtual Disk** - Captures point-in-time images of a virtual disk for backup, testing, or data processing without affecting the contents of the source virtual disk.

- **Virtual Disk Copy** - generates a full copy of data from the source virtual disk to the target virtual disk in a storage array. You can use Virtual Disk Copy to back up data, copy data from disk groups that use smaller-capacity physical disks to disk groups using greater capacity physical disks, or restore snapshot virtual disk data to the source virtual disk.



NOTE: For instructions on deploying the correct options in the cluster environment, see "Using Advanced (Premium) PowerVault Modular Disk Storage Manager Features" on page 54.

See *Installing and Configuring the Shared Storage System* and the *Dell PowerVault Modular Disk Storage Manager User's Guide* for more information about Modular Disk Storage Manager, Snapshot Virtual Disk, and Virtual Disk Copy.

Supported Cluster Configurations

Figure 1-1. Non-Redundant Cluster Configuration




Figure 1-2. Redundant Direct-Attached Cluster Configuration





Figure 1-3. Redundant Network-Attached Cluster Configuration



NOTE: n=8 for Windows Server 2003 and n=16 for Windows Server 2008

Other Documents You May Need



CAUTION: The safety information that shipped with your computer provides important safety and regulatory information. Warranty information may be included within this document or as a separate document.



NOTE: To configure Dell blade system modules in a Dell PowerEdge Cluster, see the *Using Dell Blade Servers in a Dell PowerEdge High Availability Cluster* document located on Dell Support website at support.dell.com.

- The *Rack Installation Guide* included with your rack solution describes how to install your system into a rack.
- The *Getting Started Guide* provides an overview to initially set up your system.
- The *Dell Failover Clusters with Microsoft Windows Server Installation and Troubleshooting Guide* provides more information on deploying your cluster with a specific variant of the Windows Server operating system (for example: Windows Server 2003 or Windows Server 2008).
- The *Dell Cluster Configuration Support Matrices* provides a list of recommended operating systems, hardware components, and driver or firmware versions for your Dell Windows Server Failover Cluster.
- The *Setting Up Your System* document provides an overview of initially setting up your system.
- Systems management software documentation describes the features, requirements, installation, and basic operation of the software.
- Operating system documentation describes how to install (if necessary), configure, and use the operating system software.
- The *Dell PowerVault Modular Disk Storage Manager* documentation provides instructions for using the array management software to configure RAID systems.
- Documentation for any components you purchased separately provides information to configure and install those options.
- The Dell PowerVault™ tape library documentation provides information for installing, troubleshooting, and upgrading the tape library.
- Updates are sometimes included with the system to describe changes to the system, software, and/or documentation.

- The User's Guide for your PowerEdge or PowerVault system describes system features and technical specifications, the System Setup program (if applicable), software support, and the system configuration utility.
- The *Dell PowerVault MD3000i Hardware Owner's Manual* provides information about the hardware enclosure.
- The *PowerVault Modular Disk Storage Manager CLI Guide* provides information about using the command line interface (CLI).
- The *Dell PowerVault MD3000i Resource* media provides documentation for configuration and management tools, as well as the full documentation set included here.
- The *Dell PowerVault Modular Disk Storage Manager User's Guide* provides instructions for using the array management software to configure RAID systems.
- 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.



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

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

Cabling Your Cluster Hardware

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

Cabling the Mouse, Keyboard, and Monitor

When installing a cluster configuration in a rack, you must include a switch box to connect the mouse, keyboard, and monitor to the nodes. See the documentation included with your rack for instructions on cabling each node's connections to the switch box.

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:

- For nodes with multiple power supplies, plug each power supply into a separate AC circuit.
- 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 onto one or two circuits and the redundant power supplies are grouped onto a different circuit.

Figure 2-1. Power Cabling Examples 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

Network Connection	Description
Public Network	All connections to the client LAN. At least one public network must be configured for mixed mode (public mode and private mode) for private network failover.
Private Network	A dedicated connection for sharing cluster health and status information between the cluster nodes. Network adapters connected to the LAN can also provide redundancy at the communications level in case the cluster interconnect fails. See your Microsoft® Failover Cluster Services documentation for more information on private network redundancy.

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 public network segments. You can install additional network adapters to support additional public network segments or to provide redundancy in the event of a faulty primary 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.

Table 2-2 lists the required hardware components and connection method for three possible private network configurations.

Table 2-2. Private Network Hardware Components and Connections

Method	Hardware Components	Connection
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 (two node cluster only)	Fast Ethernet network adapters.	Connect a crossover Ethernet cable between the Gigabit Ethernet network adapters in both cluster nodes.
Point-to-Point	Copper Gigabit Ethernet network adapters.	Connect a standard Ethernet cable between the Gigabit Ethernet network adapters in both cluster nodes.



NOTE: Throughout this document, Gigabit Ethernet is used to refer to either Gigabit Ethernet or 10 Gigabit Ethernet.

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 simultaneously to support both the public and private networks.


NIC Teaming


Network Interface Card (NIC) teaming combines two or more NICs to provide load balancing and/or fault tolerance. Your cluster supports NIC teaming, but only in a public network; NIC teaming is not supported in a private network.

You should use the same brand of NICs in a team, and you cannot mix brands of teaming drivers.


Cabling the Storage Systems

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

 **NOTE:** To configure Dell blade system modules in a Dell PowerEdge Cluster, see *Using Dell Blade Servers in a Dell PowerEdge High Availability Cluster* located on Dell Support website at support.dell.com.

 **NOTE:** For more details on storage hardware settings and descriptions, see *Dell PowerVault™ MD3000i RAID Enclosure Hardware Owner's Manual*.

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

 **NOTE:** It is recommended that you configure your Dell PowerVault MD3000i to use out-of-band management.


Cabling the Cluster in a Non-Redundant Configuration




The non-redundant configuration provides only a single data path from a host to the RAID enclosure, and is not recommended for applications that need critical data storage. Path failure from a failed cable or a failed iSCSI NIC may cause the application(s) to failover to the other cluster node. A failed RAID controller module may cause loss of host access to storage and thus may result in the cluster going down.

To cable the cluster:

- 1 Install a network cable from the cluster node 1 iSCSI NIC to the RAID module 0 port In-0.
- 2 Install a network cable from the cluster node 2 iSCSI NIC to the RAID module 0 port In-1.

Figure 2-4. Non-Redundant Cluster Configuration



-  **NOTE:** Multipath is required for this configuration.
-  **NOTE:** Only PowerVault MD3000i with one controller is supported in this configuration.
-  **NOTE:** The SAS out port provides SAS connection for cabling to MD1000 Expansion Enclosure.

Cabling the Cluster in Direct-Attached Redundant Configuration


In the direct-attached redundant configuration, each cluster node is directly attached to the PowerVault MD3000i RAID controller modules using two network cables, and either one dual-port NIC or two single-port NICs.

If a component fails in the storage path such as the port, the cable, or the storage controller, the multipath software automatically re-routes the I/O requests to the alternate path so that the storage array continues to operate without interruption. The configuration with two single-port NICs provides higher availability; a NIC failure does not cause failover cluster to move cluster resources to the other cluster node.

To cable the cluster:

- 1** Connect cluster node 1 to the storage system:
 - a** Install a network cable from the cluster node 1 iSCSI NIC 1 (or NIC port 1) to the RAID controller module 0 port In-0.
 - b** Install a network cable from the cluster node 1 iSCSI NIC 2 (or NIC port 2) to the RAID controller module 1 port In-1.
- 2** Connect cluster node 2 to the storage system:
 - a** Install a network cable from the cluster node 2 iSCSI NIC 1 (or NIC port 1) to the RAID controller module 0 port In-1.
 - b** Install a network cable from the cluster node 2 iSCSI NIC 2 (or NIC port 2) to the RAID controller module 1 port In-0.

Figure 2-5. Direct-Attached Redundant Cluster Configuration



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

Cabling the Cluster in Network-Attached Redundant Configuration

In the network-attached redundant configuration, each cluster node attaches to the storage system using redundant IP storage area network (SAN) industry-standard 1 Gb Ethernet switches, and either with one dual-port iSCSI NIC or two single-port iSCSI NICs. If a component fails in the storage path such as the iSCSI NIC, the cable, the switch, or the storage controller, the multipath software automatically re-routes the I/O requests to the alternate path so that the storage array continues to operate without interruption. The configuration with 2 single-port NICs provides higher availability; a NIC failure does not cause Microsoft Failover Cluster to move cluster resources to the other cluster node.

This configuration can support up to 16 hosts simultaneously and requires dual controller modules. Examples of this configuration are:

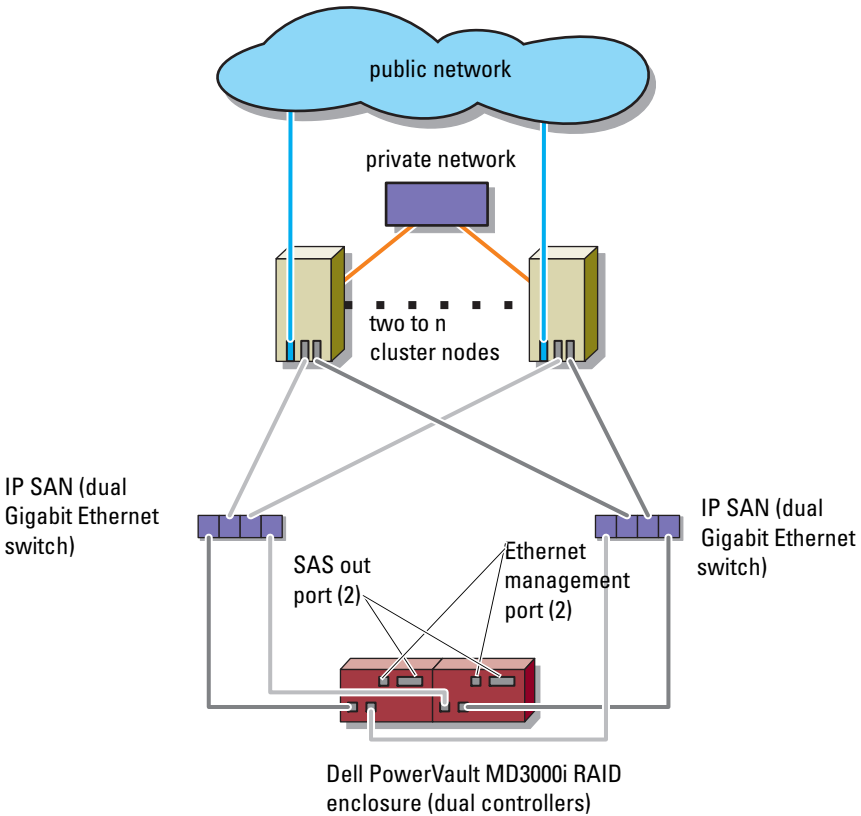
- One cluster up to eight nodes (applicable for Windows Server 2003 & Windows Server 2008).
- One cluster up to sixteen nodes (applicable for Windows Server 2008 x64 only).
- Two eight-node clusters (applicable for Windows Server 2003 & Windows Server 2008).
- One eight-node clusters, one two-node cluster, and one standalone system (applicable for Windows Server 2003 & Windows Server 2008).

To cable the cluster:

- 1** Connect the storage system to the iSCSI network:
 - a** Install a network cable from switch 1 to controller 0 port In-0.
 - b** Install a network cable from switch 1 to controller 1 port In-0.
 - c** Install a network cable from switch 2 to controller 0 port In-1.
 - d** Install a network cable from switch 2 to controller 1 port In-1.
- 2** Connect the cluster to the iSCSI network:
 - a** Install a network cable from the cluster node 1 iSCSI NIC 1 (or NIC port 1) to the network switch 1.
 - b** Install a network cable from the cluster node 1 iSCSI NIC 2 (or NIC port 2) to the network switch 2.
 - c** Repeat step a and b for each additional node.

- 3 Connect each additional cluster or standalone system to the iSCSI network, similar to step 2.

Figure 2-6. Network-Attached Redundant Cluster Configuration



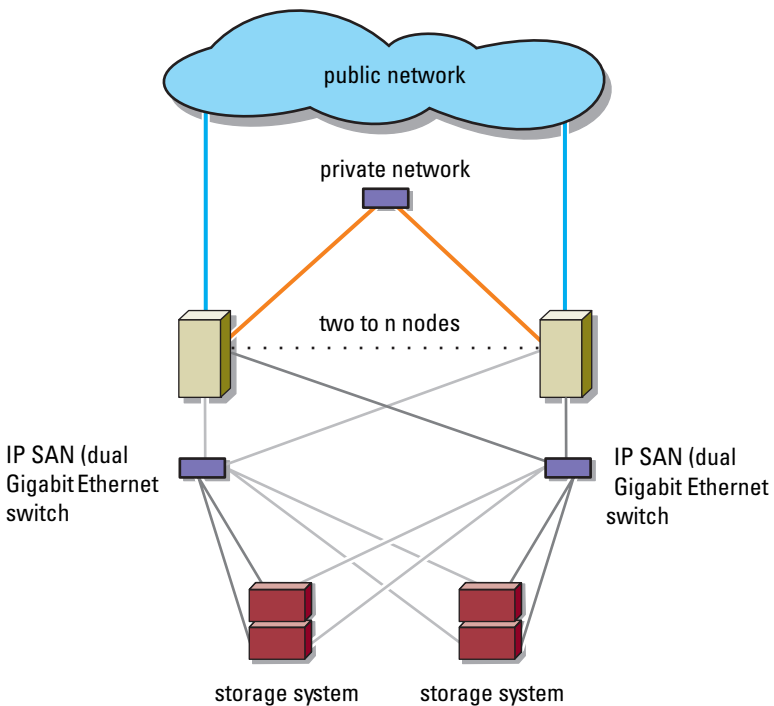
NOTE: n=8 for Windows Server 2003 and n=16 for Windows Server 2008

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

Connecting a PowerEdge Cluster to Multiple PowerVault MD3000i Storage Systems

You can increase your cluster storage capacity by attaching multiple storage systems to your cluster using redundant network switches. The PowerEdge cluster systems support configurations with multiple PowerVault MD3000i storage systems attached to clustered systems. In this scenario, the Failover Cluster software can fail over disk drives in any cluster-attached shared storage system between the cluster nodes.

Figure 2-7. Network-Attached Redundant Cluster Configuration With Multiple Storage Arrays




NOTE: n=8 for Windows Server 2003 and n=16 for Windows Server 2008

When attaching multiple PowerVault MD3000i storage systems with your cluster, the following rules apply:

- A maximum of four Power Vault MD3000i storage systems per cluster.
- The shared storage systems and firmware must be identical. Using dissimilar storage systems and firmware for your shared storage is not supported.
- Windows limits access to drives using limited drive letters which is 22. Because drive letters A through D are reserved for local disks, a maximum of 22 drive letters (E to Z) can be used for your storage system disks.
- Windows Server 2003 R2 & Enterprise & also Windows Server 2008, Enterprise Edition supports mount points, allowing greater than 22 drives per cluster.

Preparing Your Systems for Clustering

 **CAUTION:** Only trained service technicians are authorized to remove and access any of the components inside the system. See the safety information that shipped with your computer for complete information about safety precautions, working inside the computer, and protecting against electrostatic discharge.

Cluster Configuration Overview

- 1 Ensure that your site can handle the cluster's power requirements.
Contact your sales representative for information about your region's power requirements.
- 2 Install the servers, the shared storage array(s), and the interconnect switches (example: in an equipment rack), and ensure that all these components are powered on.



NOTE: For more information on step 3 to step 7 and step 10 to step 12, see the "Preparing your systems for clustering" section of the *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* or the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

- 3 Deploy the operating system (including any relevant service pack and hotfixes), network adapter drivers, and storage adapter drivers (including Multipath I/O drivers (MPIO)) on each of the servers that will become cluster nodes. Depending on the deployment method that is used, it may be necessary to provide a network connection to successfully complete this step.



NOTE: You can record the Cluster configuration and Zoning configuration (if relevant) to the Cluster Data Form and Zoning Configuration Form, respectively to help in planning and deployment of your cluster. For more information, see the "Cluster Data Form" on page 75 and the iSCSI configuration information in the worksheet located at "iSCSI Configuration Worksheet" on page 77.

- 4 Establish the physical network topology and the TCP/IP settings for network adapters on each server node to provide access to the cluster public and private networks.
- 5 Configure each server node as a member server in the same Windows Active Directory® Domain.



NOTE: You can configure the cluster nodes as Domain Controllers. For more information, see the "Selecting a Domain Model" section of the *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* or the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* located on the Dell Support website at support.dell.com.

- 6 Establish the physical storage topology and any required storage network settings to provide connectivity between the storage array and the servers that will be configured as cluster nodes. Configure the storage system(s) as described in your storage system documentation.
- 7 Use storage array management tools to create at least one logical unit number (LUN). The LUN is used as a cluster quorum disk for Windows Server 2003 Failover cluster and as a witness disk for Windows Server 2008 Failover cluster. Ensure that this LUN is presented to the servers that will be configured as cluster nodes.



NOTE: It is recommended that you configure the LUN on a single node, for security reasons, as mentioned in step 8 when you are setting up the cluster. Later, you can configure the LUN as mentioned in step 9 so that other cluster nodes can access it.

- 8 Select one of the systems and form a new failover cluster by configuring the cluster name, cluster management IP, and quorum resource. For more information, see "Preparing Your Systems for Clustering" on page 31.



NOTE: For Windows Server 2008 Failover Clusters, run the **Cluster Validation Wizard** to ensure that your system is ready to form the cluster.

- 9 Join the remaining node(s) to the failover cluster. For more information, see "Preparing Your Systems for Clustering" on page 31.
- 10 Configure roles for cluster networks. Take any network interfaces that are used for iSCSI storage (or for other purposes outside of the cluster) out of the control of the cluster.

- 11 Test the failover capabilities of your new cluster.



NOTE: For Windows Server 2008 Failover Clusters, you can also use the **Cluster Validation Wizard**.

- 12 Configure highly-available applications and services on your failover cluster. Depending on your configuration, this may also require providing additional LUNs to the cluster or creating new cluster resource groups. Test the failover capabilities of the new resources.
- 13 Configure client systems to access the highly available applications and services that are hosted on your failover cluster.

Installation Overview

Each node in your Dell Windows Server failover cluster should have the same release, edition, service pack, and processor architecture of the Windows Server operating system installed. For example, all nodes in your cluster may be configured with Windows Server 2003 R2, Enterprise x64 Edition. If the operating system varies among nodes, it is not possible to configure a failover cluster successfully. It is recommended to establish system roles prior to configuring a failover cluster, depending on the operating system configured on your cluster.

For a list of recommended operating systems, hardware components, and driver or firmware versions for your Dell Windows Server Failover Cluster, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at www.dell.com/ha.

For more information on deploying your cluster with the Windows Server 2003 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide* on the Dell Support website at support.dell.com.

For more information on deploying your cluster with the Windows Server 2008 operating systems, see the *Dell Failover Clusters with Microsoft Windows Server 2008 Installation and Troubleshooting Guide* on the Dell Support website at support.dell.com.

The following sub-sections describe steps that enable you to establish communication between the cluster nodes and your shared MD3000i storage array, and to present disks from the storage array to the cluster.

- 1 Installing the iSCSI NICs
- 2 Installing the Microsoft iSCSI Software Initiator
- 3 Installing and Configuring the Storage Management Software
- 4 Installing and Configuring the Shared Storage System
- 5 Installing and Configuring a Failover Cluster

Installing the iSCSI NICs

It is recommended that you install the latest supported version of the driver. If the NIC driver requires any service packs or hotfixes to be installed along with the operating system, install them at this time.

For a list of recommended operating systems, hardware components, and driver or firmware versions for your Dell Windows Server Failover Cluster, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at www.dell.com/ha.

Enabling TOE NIC

The purpose of TOE is to take the TCP/IP packets to be processed by the system microprocessor(s) and offload them on the NIC. The TOE eliminates the bottlenecks with applications that generate significant network traffic, freeing up CPU cycles, and the amount of available main memory bandwidth. TOE NICs provide increased performance for iSCSI traffic.



NOTE: All the nodes in a cluster solution must use similar NICs (TOE NICs or regular NICs) for iSCSI traffic. Combining TOE NICs and regular NICs is not supported in a cluster solution.

Enabling TOE NIC requires the installation of the Microsoft Scalable Networking Pack (SNP) hotfix and the TOE NIC driver package on the node. For more information, see the TOE NIC setup in your system's Installation Guide, located on the Dell Support website at support.dell.com.

You must configure the public, private, and iSCSI networks in each node before you install MSCS. The following sections explain the principles and procedures related to the networking prerequisites.

Installing the Microsoft iSCSI Software Initiator

- 1 Use a web browser and go to the Microsoft Download Center website at www.microsoft.com/downloads.
- 2 Search for *iscsi initiator*.
- 3 Select and download the latest supported initiator software and related documentation for your operating system.



NOTE: For the latest supported Software Initiator version, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at www.dell.com/ha.


- 4 Double-click the executable file. The installation wizard launches. In the **Welcome** screen, click **Next**.
- 5 In the following screens select the **Initiator Service**, **Software Initiator**, and **Microsoft MPIO Multipathing Support for iSCSI** options. Click **Next** to continue with the installation.
- 6 Read and accept the license agreement and click **Next** to install the software.
- 7 At the completion screen, click **Finish** to complete the installation.
- 8 Select the **restart now** option to reboot the system.


Installing and Configuring the Storage Management Software

To install and configure the PowerVault MD3000i RAID enclosure in your cluster:


- 1 Ensure that the PowerVault MD3000i has the latest firmware. See your PowerVault MD3000i document for more information.
- 2 Install the Host Software (multipath software and Modular Disk Storage Manager Agent) on each cluster node, and Modular Disk Storage Manager Client software on the management station. For more information, see the *Dell PowerVault Modular Disk 3000i Systems Installation Guide*.
- 3 Set the correct failback mode on each cluster node—you must merge the **PowerVault MD3000i Stand Alone to Cluster.reg** file located in the `windows\utility` directory of the *Dell PowerVault MD3000i Resource CD* into the registry of each node.

- 4 If you have third-party applications that use the Microsoft Volume Shadow-copy Service (VSS) or Virtual Disk Service (VDS) Application Programming Interface (API), install the VDS_VSS package located in the windows\VDS_VSS directory on the *PowerVault MD3000i Resource CD*. Separate versions for 32-bit and 64-bit operating systems are provided. The VSS and VDS provider will engage only if it is needed.

 **NOTE:** If you uninstall and reinstall the Multipath I/O (MPIO) software or PowerVault Modular Disk Manager, you must merge the **PowerVault MD3000i Stand Alone** to the **Cluster.reg** file into the registry again.

 **NOTE:** If you are reconfiguring a cluster node into a standalone host, you must merge the **PowerVault MD3000i Cluster to Stand Alone.reg** file located in the `windows\utility` directory of the *Dell PowerVault MD3000i Resource CD* into the host registry.


These registry files enable correct failback operation on the host.

 **NOTE:** The cluster node can be used as a management station.

You can manage a storage array in two ways:

- Out-of-band management
- In-band management

When you use out-of-band management, you must set the network configuration for each RAID controller module including its IP address, subnet mask, and gateway. If you are using a DHCP system, you can enable automatic network configuration, but if you are not using a DHCP system, you must enter the network configuration manually.

 **NOTE:** It is recommended that you use out-of-band management.

Adding Storage Arrays

To add a storage array to the Modular Disk Storage Manager, click the **New** link in the **Array Selector** area. A window is displayed that allows you to choose the automatic or manual process to add a new storage array.

You can add the Storage Arrays using either **Automatic Discovery** or **Manual Discovery**

Installing and Configuring the Shared Storage System

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

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. If you follow these steps, you can be certain that you have completed all the basic steps to configure your storage array.

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 unapproved 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 the 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 iSCSI connections** — Set up one or more hosts to access the storage array. See *Configuring iSCSI Connections* for more information.
- 6 Configuring and Manage Virtual Disks** — See "Creating Disk Groups and Virtual Disks" on page 49 for more information.

- 7 View and Enable Premium Features (Optional)** — If you have purchased premium features, including Snapshot Virtual Disks and virtual disk copies, check the premium features that are currently available and enable them if they are turned off. See *Using Advanced (Premium) PowerVault Modular Disk Storage Manager Features* for more information.
- 8 Changing Network Configuration (Optional)** — Change your network configuration by changing RAID controller network settings or obtain the network configuration from a DHCP system.

Configuring iSCSI Connections

To use the storage array, you must configure iSCSI on both the host server(s) and the storage array. Step-by-step instructions for configuring iSCSI are described in this section. However, before configuring iSCSI, you must install the Microsoft iSCSI initiator and the MD Storage Manager software. For more information on installing, see "Installing the Microsoft iSCSI Software Initiator" on page 35 and "Installing and Configuring the Storage Management Software" on page 35.



NOTE: Although some of these steps shown in this section can be performed in MD Storage Manager from a management station, the iSCSI initiator must be installed and configured on each host server.

Before You Start

Before you begin configuring iSCSI, you should fill out the "iSCSI Configuration Worksheet" on page 77. Gathering this type of information about your network prior to starting the configuration steps helps you complete the process faster.

Terminology

The table below outlines the terminology used in the iSCSI configuration steps later in this section.

Table 3-1. Standard Terminology Used in iSCSI Configuration

Term	Definition
CHAP (Challenge Handshake Authentication Protocol)	An optional security protocol used to control access to an iSCSI storage system by restricting use of the iSCSI data ports on both the host server and storage array. For more information on the types of CHAP authentication supported, see "Creating a Host Group" on page 48.
host or host server	A server connected to the storage array through iSCSI ports.
host server port	iSCSI port on the host server used to connect it to the storage array.
iSCSI initiator	The iSCSI-specific software installed on the host server that controls communications between the host server and the storage array.
iSCSI host port	The iSCSI port (two per controller) on the storage array.
iSNS (Microsoft Internet Storage Naming Service)	An automated discovery, management, and configuration tool used by some iSCSI devices.
management station	The system from which you manage your host server/storage array configuration.
storage array	The enclosure containing the storage data accessed by the host server.
target	An iSCSI port on the storage array that accepts and responds to requests from the iSCSI initiator installed on the host server.

The "iSCSI Configuration Worksheet" on page 77 helps you plan your configuration. Recording host server and storage array IP addresses at a single location will help you configure your setup faster and more efficiently.

Configuring iSCSI on Your Storage Array

The following sections contains step-by-step instructions for configuring iSCSI on your storage array. However, before beginning, it is important to understand where each of these steps occur in relation to your host server/storage array environment.

Table 3-2 contains the sequence of steps for configuring each specific iSCSI connections and where it occurs. The following subsections describe each of the steps in more detail.

Table 3-2. Host Server vs. Storage Array

This step is performed on the HOST SERVER using the Microsoft iSCSI Initiator:	This step is performed on the STORAGE ARRAY using MD Storage Manager:
Perform Target Discovery From the iSCSI Initiator	Discover the Storage Array (Out-of-Band Management Only)
Configuring CHAP Authentication on the Host Server (Optional)	Configure the iSCSI Ports on the Storage Array
Connect to the Target Storage Array From the Host Server	Configure Host Access
	Configuring CHAP Authentication on the Storage Array (Optional)
	Set Up In-Band Management (Optional)

Using iSNS

Internet Storage Naming Service (iSNS) Server, supported only on Windows iSCSI environments, eliminates the need to manually configure each individual storage array with a specific list of initiators and target IP addresses. Instead, iSNS automatically discovers, manages, and configures all iSCSI devices in your environment.

For more information on iSNS, including installation and configuration, see www.microsoft.com.

Discover the Storage Array (Out-of-Band Management Only)

The following subsections describe the methods of discovering the storage array.

Default Management Port Settings

By default, the storage array management ports will be set to DHCP configuration. If the controller(s) on your storage array is unable to get IP configuration from a DHCP system, it will timeout and revert to a default static IP address. The default IP configuration is:

Controller 0: IP: 192.168.128.101 Subnet Mask: 255.255.255.0

Controller 1: IP: 192.168.128.102 Subnet Mask: 255.255.255.0



NOTE: The default gateway is not set.



NOTE: If DHCP is not used, you must perform initial configuration of the management station on the same physical subnet as the storage array. Also, during initial configuration, you must configure at least one network adapter on the same IP subnet as the storage array's default management port (192.168.128.101 or 192.168.128.102). After initial configuration (use MD Storage Manager to manage the ports), you can change the management station's IP address back to its previous settings.



NOTE: This procedure applies to out-of-band management only. If you choose to set up in-band management, you must complete this step and then refer to "Set Up In-Band Management (Optional)" on page 62.

You can discover the storage array automatically or manually. Choose one and complete the steps below.

Automatic Storage Array Discovery

- 1** Launch MD Storage Manager.

If this is the first storage array to be set up, the **Add New Storage Array** window appears.

- 2** Choose **Automatic** and click **OK**.

It may take several minutes for the discovery process to complete. If you close the Discovery Status window before the discovery is complete, the discovery process is cancelled.

- 3** After discovery is complete, a confirmation screen appears. Click **Close** to close the screen.

Manual Storage Array Discovery

- 1** Launch MD Storage Manager.
If this is the first storage array to be set up, the **Add New Storage Array** window appears.
- 2** Select **Manual** and click **OK**.
- 3** Select **Out-of-band management** and enter the host server name(s) or IP address(es) of the iSCSI storage array controller.
- 4** Click **Add**.

Out-of-band management should now be successfully configured.

After discovery is complete, a confirmation screen appears. Click **Close** to close the screen.

Setting Up the Array

- 1** When discovery is complete, the name of the first storage array found appears under the **Summary** tab in MD Storage Manager.
- 2** The default name for the newly discovered storage array is *Unnamed*. If another name appears, click the down arrow next to that name and choose **Unnamed** in the drop-down list.
- 3** Click the **Initial Setup Tasks** option to see links to the remaining post-installation tasks. For more information about each task, see the *User's Guide*. Perform these tasks in the order shown in Table 3-3.



NOTE: Before configuring the storage array, check the status icons on the **Summary** tab to ensure that the enclosures in the storage array are in an optimal status. For more information on the status icons, see "Troubleshooting Tools" on page 72.

Table 3-3. Initial Storage Array Setup Tasks

Task	Purpose	Information Needed
<p>Rename the storage array.</p> <p>If you need to physically find the device, click Blink the storage array on the Initial Setup Tasks dialog box or click the Tools tab and choose Blink. Lights on the front of the storage array blink intermittently to identify the array.</p>	<p>To provide more a meaningful name than the software-assigned label of Unnamed.</p>	<p>A unique, clear name with no more than 30 characters that may include letters, numbers, and no special characters other than underscore (<u>_</u>), minus (-), or pound sign (#).</p> <p>NOTE: MD Storage Manager does not check for duplicate names. Names are not case sensitive.</p>
<p>Set a storage array password.</p>	<p>To restrict unauthorized access, MD Storage Manager asks for a password before changing the configuration or performing a destructive operation.</p>	<p>A case-sensitive password that meets the security requirements of your enterprise.</p>
<p>Set the management port IP addresses on each controller.</p>	<p>To set the management port IP addresses to match your public network configuration. Although DHCP is supported, static IP addressing is recommended.</p>	<p>In MD Storage Manager, select Initial Setup Tasks → Configure Ethernet Management Ports, then specify the IP configuration for each management port on the storage array controllers.</p> <p>NOTE: If you change a management port IP address, you may need to update your management station configuration and/or repeat storage array discovery.</p>

Table 3-3. Initial Storage Array Setup Tasks (continued)

Task	Purpose	Information Needed
Set up alert notifications. <ul style="list-style-type: none">• Set up e-mail alerts.• Set up SNMP alerts. The Status area in the Summary tab shows if alerts have been set for the selected array.	To arrange to notify individuals (by e-mail) and/or storage management stations (by SNMP) when a storage array component degrades or fails, or an adverse environmental condition occurs.	E-mail — Sender (sender's SMTP gateway and e-mail address) and recipients (fully qualified e-mail addresses) SNMP — <ol style="list-style-type: none">1 A community name2 A known set of storage management stations set by administrator as an ASCII string in the management console (default: "public")3 A trap destination4 IP address or host name of a management console running an SNMP service

Configure the iSCSI Ports on the Storage Array

By default, the iSCSI ports on the storage array are set as given below.

IPv4 Addresses (Default IP addresses):

Controller 0, Port 0: IP: 192.168.130.101 Subnet Mask: 255.255.255.0
Port: 3260

Controller 0, Port 1: IP: 192.168.131.101 Subnet Mask: 255.255.255.0
Port: 3260

Controller 1, Port 0: IP: 192.168.130.102 Subnet Mask: 255.255.255.0
Port: 3260

Controller 1, Port 1: IP: 192.168.131.102 Subnet Mask: 255.255.255.0
Port: 3260

IPv6 addresses (Default IP addresses):

Controller 0, Port 0: IP Configuration: Obtain configuration automatically
Port: 3260

Controller 0, Port 1: IP Configuration: Obtain configuration automatically
Port: 3260

Controller 1, Port 0: IP Configuration: Obtain configuration automatically
Port: 3260

Controller 1, Port 1: IP Configuration: Obtain configuration automatically
Port: 3260



NOTE: By default, the gateway value is not set.

To configure the iSCSI ports on the storage array, complete the following steps:

- 1 From MD Storage Manager, click iSCSI and select **Configure iSCSI Host Ports**.
- 2 Configure the iSCSI ports on the storage array.



NOTE: Using static IP addressing is recommended, although DHCP is supported.

Setting Static IP addresses:

IPv4 Configurations – Only one IP address can be configured per controller per port.

IPv6 Configurations – Two routable IP addresses can also be configured in addition to one Link-local IP address assigned to each port of both the controllers.

The typical format used for the IP addresses is as below:

Link Local IP address:

FE80:0000:0000:0000:XXXX:XXXX:XXXX:XXXX

Routable IP address 1:

XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX

Routable IP address 2:

XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX

Router IP address:

XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX

Where 'X' is a hexadecimal number between 0 and F.

The following settings are available (depending on your specific configuration) by clicking the **Advanced** button:

- **Virtual LAN (VLAN) support**

A VLAN is a network of different systems that behave as if they are connected to the same segments of a local area network (LAN) and are supported by the same switches and routers. When configured as a VLAN, a device can be moved to another location without being reconfigured. To use VLAN on your storage array, obtain the VLAN ID from your network administrator and enter it here.

- **Ethernet priority**

This parameter is set to determine a network access priority.

- **TCP listening port**

The port number used by the controller on the storage array to listen for iSCSI logins from host server iSCSI initiators.



NOTE: The TCP listening port for the iSNS system is the port number the storage array controller uses to connect to an iSNS system. This allows the iSNS system to register the iSCSI target and portals of the storage array so that the host server initiators can identify them.

- **Jumbo frames**

Jumbo Ethernet frames are created when the maximum transmission units (MTUs) are larger than 1500 bytes per frame. This setting is adjustable port-by-port.

- 3 To enable ICMP PING responses for all ports, select **Enable ICMP PING responses**.
- 4 Click **OK** when all iSCSI storage array port configurations are complete.
- 5 Test the connection by performing a ping command on each iSCSI storage array port.

Perform Target Discovery From the iSCSI Initiator

This step identifies the iSCSI ports on the storage array to the host server. To perform target discovery from the iSCSI initiator:

- 1 Click **Start**→ **Programs**→ **Microsoft iSCSI Initiator**.
- 2 Click the **Discovery** tab.
- 3 Under **Target Portals**, click **Add** and enter the **IP address or DNS name** of the iSCSI port on the storage array.
- 4 If the iSCSI storage array uses a custom TCP/IP port, change the **Port** number. The default port number is 3260.
- 5 Click **Advanced** and set the following values on the **General** tab:
 - **Local Adapter:** Must be set to **Microsoft iSCSI Initiator**.
 - **Source IP:** The source IP address of the host you want to connect with.
 - **Data Digest and Header Digest:** Optionally, you can specify that a digest of data or header information be compiled during transmission to assist in troubleshooting.
 - **CHAP logon information:** Leave this option unselected and do not enter CHAP information at this point, unless you are adding the storage array to a SAN that has target CHAP already configured.



NOTE: IPSec is not supported.

Click **OK** to exit the **Advanced** menu. Click **OK** to close the **Add Target Portals** screen.

- 6 To exit the **Discovery** tab, click **OK**.

If you plan to configure **CHAP authentication**, do not perform discovery on more than one iSCSI port at this point. Go to "Configure Host Access" on page 48.

If you do not plan to configure **CHAP authentication**, repeat step 1 through step 6 (above) for all iSCSI ports on the storage array.

Configure Host Access

This section helps you configure the connection between the host servers and the virtual disks on the storage array. You must perform this step before mapping virtual disks to host servers, or whenever you connect new host servers to the storage array.

- 1 Launch MD Storage Manager.
- 2 Click **Configure** and select **Configure Host Access (Manual)**.
- 3 At **Enter host name**, enter the host server to be available to the storage array for virtual disk mapping.

This can be an informal name, not necessarily a name used to identify the host server to the network.

- 4 In the **Select host type** drop-down menu, select the host type.

For a Non-Redundant Configuration, select **Windows MSCS Cluster - Single Path**.

For a Redundant Configuration with Dual SAS 5/E HBAs, select **Windows 2000/Server 2003 Clustered**.

- 5 Click **OK** to configure access to the array for the hosts you selected.
- 6 If your iSCSI initiator shows up in the list of **Known iSCSI Initiators**, ensure it is highlighted. Click **Add**, and then click **Next**. Otherwise, click **New** and enter the **iSCSI Initiator Name**. The iSCSI initiator name can be found on the **General** tab of the **iSCSI Initiator Properties** window. Click **Next** to proceed.
- 7 Choose whether the host server will be part of a host server group that will share access to the same virtual disks as other host servers. Select **Yes** only if the host is part of a Microsoft cluster. Click **Next**.
- 8 Click **Finish**.

Creating a Host Group

- 1 Click **Modify** 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.
- 5 Repeat step 4 to add the second cluster node 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. A minimum of one virtual disk is required for an active/passive cluster configuration; at least two virtual disks are required for an active/active cluster configuration.

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 hosts attached to the storage array read and write data to the virtual disks.



NOTE: Before you can 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.

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

- Automatic Configuration
- Manual Configuration

See the *Dell PowerVault Modular Disk Storage Manager User's Guide* for more information on how to create Disk Groups and Virtual Disks.

It is recommended that you create at least one virtual disk for each application. If multiple NTFS volumes are created on a single virtual disk using Windows Disk Management, the volumes failover together, rather than individually from node-to-node.



NOTE: It is recommended that you use a RAID level other than RAID 0 (which is commonly called striping). RAID 0 configurations provide very high performance, but do not provide the level of availability required for the quorum resource. See the documentation for your storage system for more information about setting up RAID levels for the system.

Creating Host-to-Virtual Disk Mappings

Create host-to-virtual disk mappings to assign virtual disks to the host groups containing cluster nodes by clicking the **Configure** tab, then clicking 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. 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.

Loading RAID Controller Module NVSRAM for Non-Redundant Configuration

For proper operation, the non-redundant configuration requires a special NVSRAM file be loaded onto the PowerVault MD3000i RAID enclosure. The NVSRAM file is located at the \utility\NVSRAM\ directory on the PowerVaultMD3000i Resource Media, with a prefix of Non-redundant-MSCS. To load the NVSRAM file to the PowerVault MD3000i RAID enclosure, from the storage management station, open the PowerVault Modular Disk Storage Manager Client.

- 1** Click **Support** and then, **Download Firmware**.
- 2** In the Download Firmware window, click **Download RAID Controller Module NVSRAM**. The current controller firmware and NVSRAM versions in use are displayed.
- 3** Click **Select File** to browse to the file you want to download. By default, only firmware images compatible with the current storage array configuration are listed.
- 4** Select the appropriate file in the **File Selection** window and click **OK**. If the file you selected is not valid or is incompatible with the current storage array configuration, an error message is displayed. Click **OK** to close the message and select another file.
- 5** Click **Transfer...** A **Confirm Download** dialog box is displayed showing the RAID controller and NVSRAM firmware you selected.
- 6** To complete the download, click **Yes**.

Configuring the RAID Level for the Shared Storage Subsystem

The virtual disks in your shared storage subsystem must be configured into disk groups or virtual disks using the Dell PowerVault Modular Storage Manager software. All virtual disks, especially if they are used for the quorum resource, should be bound and should incorporate the appropriate RAID level to ensure high availability.



NOTE: It is recommended that you use a RAID level other than RAID 0 (which is commonly called striping). RAID 0 configurations provide very high performance, but do not provide the level of availability required for the quorum resource. See the documentation for your storage system for more information about setting up RAID levels for your system.

Windows Operating System and Dynamic Volumes

The Windows operating system does not support dynamic disks (upgraded disks) or volumes as shared cluster storage. If the shared cluster storage is configured as a dynamic disk, the Cluster Configuration wizard is not able to discover the disks, preventing the cluster and network clients from accessing the disks.

Assigning Drive Letters and Mount Points

A mount point is a drive attached to an empty folder on an NTFS volume. A mount point functions the same as a normal drive but is assigned a label or name instead of a drive letter. Using mount points, a cluster can support more shared disks than the number of available drive letters.

The cluster installation procedure does not automatically add the mount point into the disks managed by the cluster. To add the mount point to the cluster, create a physical disk resource in the cluster resource group for each mount point. Ensure that the new physical disk resource is in the same cluster resource group and is dependent on the root disk (i.e., the disk from which the mount point is attached).



NOTE: Mount points are supported in Microsoft Cluster on the Windows Server 2003 operating systems only. When mounting a drive to an NTFS volume, do not create mount points from the quorum resource or between the clustered disks and the local disks. Mount points must be in the same cluster resource group and must be dependent on the root disk.

Naming and Formatting Drives on the Shared Storage System

Each virtual disk being created in the PowerVault Modular Disk Storage Manager becomes a physical disk in Windows Disk Management. For each physical disk, perform the following:

- Write the disk signature
- Create the partition
- Assign the drive letter
- Format the partition with NTFS



NOTICE: The drive letters are manually assigned from the second node, the shared disks are simultaneously accessible from both nodes. To ensure file system integrity and prevent possible data loss before you install the MSCS/Failover Cluster software, prevent any I/O activity to the shared drives by performing the following procedure on one node at a time and ensuring that the other node is shut down.

The number of drive letters required by individual servers in a cluster may vary. It is recommended that the shared drives be named in reverse alphabetical order beginning with the letter *z*. To assign drive letters and format drives on the shared storage system, perform the following steps:

- 1 Turn off node 2 and open **Disk Management** on node 1.
- 2 Allow Windows to enter a signature on all new physical or logical drives.



NOTE: Do not upgrade or convert your disks to dynamic disks.

- 3 Locate the icon for the first unnamed, unformatted drive on the shared storage system.
- 4 Right-click the icon and select **Create** from the submenu. If the unformatted drives are not visible, verify the following:
 - The iSCSI Initiator target connections are active.
 - The LUNs have been assigned to the hosts.
 - The storage system is properly cabled to the servers.

- 5 In the dialog box, create a partition the size of the entire drive (the default) and then click **OK**.



NOTE: A virtual disk that is mapped or assigned from the storage system to a cluster node(s) is represented as a physical disk within the Windows operating system on each node. Microsoft Cluster allows only one node to access a given physical disk resource at a time. Therefore, if a disk is partitioned and contains multiple NTFS volumes, concurrent access to different volumes is only possible from the cluster node controlling the physical disk resource. If two NTFS volumes need to be controlled by different nodes, these volumes must reside on separate disks.

- 6 Click **Yes** to confirm the partition.
- 7 With the mouse pointer on the same icon, right-click and select **Change Drive Letter and Path** from the submenu.
- 8 Assign a drive letter to an NTFS volume or create a mount point.

To assign a drive letter to an NTFS volume:

- a Click **Edit** and select the letter you want to assign to the drive (for example, z).
- b Click **OK**.
- c Go to step 9.

To create a mount point:

- a Click **Add**.
- b Click **Mount** in the following empty NTFS folder.
- c Type the path to an empty folder on an NTFS volume, or click **Browse** to locate it.
- d Click **OK**.
- e Go to step 9.

- 9 Click **Yes** to confirm the changes.
- 10 Right-click the drive icon again and select **Format** from the submenu.
- 11 Under **Volume Label**, enter a descriptive name for the new volume; for example, `Disk_Z` or `Email_Data`.

12 In the dialog box, change the file system to **NTFS**, select **Quick Format**, and click the **Start** button.



NOTE: The NTFS file system format is required for shared-disk resources under Microsoft Cluster.

13 Click **OK** at the warning.

14 Click **OK** to acknowledge that the format is complete.

15 Click **Close** to close the dialog box.

16 Repeat step 3 through step 15 for each remaining drive.

17 Close **Disk Management**.

18 Turn off node 1.

19 Turn on node 2.

20 On node 2, open **Disk Management**.

21 Ensure that the drive letters for node 2 are correct and re-assign the drive letters, if necessary. To re-assign the drive letters, repeat step 7 through step 9.

Using Advanced (Premium) PowerVault Modular Disk Storage Manager Features

PowerVault Modular Disk Storage Manager includes the following advanced features:

- Snapshot Virtual Disk
- Virtual Disk Copy


To install and enable these premium features, you must purchase a feature key file for each feature and then specify the storage array that will host them. For instructions about this process, see the *Premium Feature Activation* card that shipped along with your Dell PowerVault MD3000i storage system.

These premium features increase the high availability for your cluster solution. It is essential that you follow the instructions below to ensure proper cluster operations.

Snapshot Virtual Disk


Snapshot Virtual Disk captures point-in-time images of a virtual disk for backup, testing, or data processing without affecting the contents of the source virtual disk. You can use either Simple Path or Advanced Path to

create a snapshot for your cluster disk. The Snapshot Virtual Disk can be mapped to the primary node (the node owning the source disk) or the secondary node (the node not owning the source disk) for backup, testing, or data processing.


 **NOTICE:** Avoid mapping the Snapshot Virtual Disk to more than one node in the cluster at any point of time. The Snapshot Virtual Disk is not managed by Cluster Administrator/Failover Cluster Manager, so mapping the Snapshot Virtual Disk to the host group or both nodes in the cluster may allow both nodes to access data concurrently and thus cause data corruption.


To map the Snapshot Virtual Disk to the primary node:

- 1 Use Host-to-Virtual Disk Mapping in the Modular Disk Storage Manager. This ensures that a different disk signature is assigned properly to the Snapshot Virtual Disk.
- 2 Use Windows Disk Management to re-scan for the Snapshot Virtual Disk, assign the drive letter, and start accessing the drive.

 **NOTE:** The disks may be re-scanned several times for the Snapshot Virtual Disk to be detected by Windows Disk Management. If the Snapshot Virtual Disk is not detected, wait for a few minutes and re-scan the disks. Repeat the process until the Snapshot Virtual Disk is detected; do not reboot the server.

If you need to map the Snapshot Virtual Disk to the secondary node (the node not owning the source disk), you must map the Snapshot Virtual Disk to the primary node first, to ensure that the snapshot is assigned a new disk signature. Then, use Modular Disk Storage Manager to unmap the Snapshot Virtual Disk from the primary node, map it to the secondary node, and start accessing it.

 **NOTICE:** Attempts to map the Snapshot Virtual Disk to the secondary node, prior to obtaining the signature from the primary node, may cause the operating system to misidentify the Snapshot Virtual Disk as an existing system volume and that may result in data loss or an inaccessible Snapshot Virtual Disk.

 **NOTE:** For a cluster configuration with multiple Snapshot Virtual Disks, each virtual disk must be mapped to the node owning the associated source disk first. The primary node for a Snapshot Virtual Disk may not be the primary node for another Snapshot Virtual Disk.

Virtual Disk Copy

Virtual Disk Copy generates a full copy of data from the source virtual disk to the target virtual disk in a storage array. You can use Virtual Disk Copy to back up data, copy data from disk groups that use smaller-capacity physical disks to disk groups using greater-capacity physical disks, or restore Snapshot Virtual Disk data to the source virtual disk.

To create a Virtual Disk Copy of a Microsoft Cluster shared disk:

- 1 Create a Snapshot Virtual Disk using the cluster shared disk as a source disk.
- 2 Do not map that Snapshot Virtual Disk to any cluster node. Then, use the newly created Snapshot Virtual Disk as the source disk for the Virtual Disk Copy.



NOTE: When you attempt to create a Virtual Disk Copy of an Microsoft Cluster shared disk directly, the operation fails and displays the following error: The operation cannot complete because the selected virtual disk is not a source virtual disk candidate.

If the cluster shared disk fails and you need to restore it from the target virtual disk, use Cluster Administrator to change the status of the cluster group containing the failed disk to offline, and then use one of the following methods:

- 1 Use Virtual Disk Copy to transfer the data from the target virtual disk to the cluster shared disk.
- 2 Unassign the cluster shared disk from the host group and then map the target virtual disk to the host group.

Understanding CHAP Authentication

Before proceeding to either "Configuring CHAP Authentication on the Storage Array (Optional)" on page 58 or "Configuring CHAP Authentication on the Host Server (Optional)" on page 59, it would be useful to gain an overview of how CHAP authentication works.

What is CHAP?

Challenge Handshake Authentication Protocol (CHAP) is an optional iSCSI authentication method where the storage array (target) authenticates iSCSI initiators on the host server. Two types of CHAP are supported: *target* CHAP and *mutual* CHAP.

Target CHAP

In target CHAP, the storage array authenticates all requests for access issued by the iSCSI initiator(s) on the host server through a CHAP secret. To set up target CHAP authentication, you enter a CHAP secret on the storage array, then configure each iSCSI initiator on the host server to send that secret each time it attempts to access the storage array.

Mutual CHAP

In addition to setting up target CHAP, you can set up mutual CHAP in which both the storage array *and* the iSCSI initiator authenticate each other. To set up mutual CHAP, configure the iSCSI initiator with a CHAP secret that the storage array must send to the host sever in order to establish a connection. In this two-way authentication process, both the host server and the storage array send information that the other must validate before a connection is allowed.

CHAP is an optional feature and is not required to use iSCSI. However, if you do not configure CHAP authentication, any host server connected to the same IP network as the storage array can read from and write to the storage array.



NOTE: If you elect to use CHAP authentication, you must configure it on both the storage array (using MD Storage Manager) and the host server (using the iSCSI initiator) before preparing virtual disks to receive data. If you prepare disks to receive data before you configure CHAP authentication, you will lose visibility to the disks after CHAP is configured.

CHAP Definitions

To summarize the differences between target CHAP and mutual CHAP authentication, see Table 3-4.

Table 3-4. CHAP Types Defined

CHAP Type	Description
Target CHAP	Sets up accounts that iSCSI initiators use to connect to the target storage array. The target storage array then authenticates the iSCSI initiator.
Mutual CHAP	Applied <i>in addition</i> to target CHAP. Mutual CHAP sets up an account that a target storage array uses to connect to an iSCSI initiator. The iSCSI initiator then authenticates the target.

Setting Up CHAP

The next two steps in your iSCSI configuration, "Configuring CHAP Authentication on the Storage Array (Optional)" on page 58 and "Configuring CHAP Authentication on the Host Server (Optional)" on page 59, offer step-by-step procedures for setting up CHAP on your storage array and host server.

Configuring CHAP Authentication on the Storage Array (Optional)

If you are configuring target-only CHAP authentication, complete "Configuring Target CHAP Authentication on the Storage Array" on page 58 and "Configuring CHAP Authentication on the Host Server (Optional)" on page 59.

If you are configuring mutual CHAP authentication, complete "Configuring Mutual CHAP Authentication on the Storage Array" on page 59 and "Configuring CHAP Authentication on the Host Server (Optional)" on page 59.

If you are **not** configuring any type of CHAP, skip to "Connect to the Target Storage Array From the Host Server" on page 61.



NOTE: If you choose to configure mutual CHAP authentication, you must first configure target CHAP.

Remember, in terms of iSCSI configuration, the term *target* always refers to the storage array.

Configuring Target CHAP Authentication on the Storage Array

- 1 From MD Storage Manager, click the **iSCSI** tab and then **Change Target Authentication**.

Make a selection based on the following:

Table 3-5. CHAP Settings

Selection	Description
None	This is the default selection. If None is the only selection, the storage array will allow an iSCSI initiator to log on without supplying any type of CHAP authentication.
None and CHAP	The storage array will allow an iSCSI initiator to log on with or without CHAP authentication.
CHAP	If CHAP is selected and None is not selected, the storage array will require CHAP authentication before allowing access.

- 2 To configure a CHAP secret, select **CHAP** and select **CHAP Secret**.
- 3 Enter the **Target CHAP secret** (or **Generate Random Secret**), confirm it in **Confirm Target CHAP Secret**, and click **OK**.

Although the storage array allows sizes from 12 to 57 characters, many initiators only support CHAP secret sizes up to 16 characters (128-bit).



NOTE: Once entered, a CHAP secret is not retrievable. Ensure that you record the secret in an accessible place. If **Generate Random Secret** is used, copy and paste the secret into a text file for future reference since the same CHAP secret is used to authenticate any new host servers you may add to the storage array. If you forget this CHAP secret, you must disconnect all existing hosts attached to the storage array and repeat the steps in this chapter to add them.

- 4 Click **OK**.

Configuring Mutual CHAP Authentication on the Storage Array

The initiator secret must be unique for each host server that connects to the storage array and must not be the same as the target CHAP secret.

- 1 From MD Storage Manager, click on the **iSCSI** tab, then select **Enter Mutual Authentication Permissions**.
- 2 Select an initiator on the host server and click the **CHAP Secret**.
- 3 Enter the **Initiator CHAP secret**, confirm it in **Confirm initiator CHAP secret**, and click **OK**.



NOTE: In some cases, an initiator CHAP secret may already be defined in your configuration. If so, use it here.

- 4 Click **Close**.



NOTE: To remove a CHAP secret, you must delete the host initiator and add it.

Configuring CHAP Authentication on the Host Server (Optional)

If you configured CHAP authentication in "Configuring Target CHAP Authentication on the Storage Array" on page 58, complete the following steps. If not, skip to "Connect to the Target Storage Array From the Host Server" on page 61.

To optionally configure CHAP authentication on the host server:

- 1 Click **Start**→**Programs**→**Microsoft iSCSI Initiator**.
- 2 If you are NOT using mutual CHAP authentication, skip to step 4.

- 3 If you are using mutual CHAP authentication:
 - a Click the **General** tab.
 - b Select **Secret**.
 - c At the **Enter a secure secret** window, enter the mutual CHAP secret you entered for the storage array.
- 4 Click the **Discovery** tab.
- 5 Under **Target Portals**, select the IP address of the iSCSI port on the storage array and click **Remove**.

The iSCSI port you configured on the storage array during target discovery should disappear. You will reset this IP address under CHAP authentication in the steps that immediately follow.

- 6 Under **Target Portals**, click **Add** and re-enter the **IP address or DNS name** of the iSCSI port on the storage array (removed above).
- 7 Click **Advanced** and set the following values on the **General** tab:
 - **Local Adapter:** Must always be set to **Microsoft iSCSI Initiator**.
 - **Source IP:** The source IP address of the host you want to connect with.
 - **Data Digest and Header Digest:** Optionally, you can specify that a digest of data or header information be compiled during transmission to assist in troubleshooting.
 - **CHAP logon information:** Enter the target CHAP authentication username and secret you entered (for the host server) on the storage array.
 - **Perform mutual authentication:** If mutual CHAP authentication is configured, select this option.



NOTE: IPSec is not supported.

- 8 Click **OK**.

If a discovery session failover is desired, repeat step 5 and step 6 for all iSCSI ports on the storage array. Otherwise, single-host port configuration is sufficient.



NOTE: If the connection fails, ensure that all IP addresses are entered correctly. Incorrectly typed IP addresses are a common cause of connection problems.

Connect to the Target Storage Array From the Host Server

- 1 Click **Start**→ **Programs**→ **Microsoft iSCSI Initiator**.
- 2 Click the **Targets** tab.

If previous target discovery was successful, the *iqn* of the storage array should be displayed under **Targets**.
- 3 Click **Log On**.
- 4 Select **Automatically restore this connection when the system boots**.
- 5 Select **Enable multipath**.
- 6 Click **Advanced** and configure the following settings under the **General** tab:
 - **Local Adapter**: Must be set to **Microsoft iSCSI Initiator**.
 - **Source IP**: The source IP address of the host server you want to connect from.
 - **Target Portal**: Select the iSCSI port on the storage array controller that you want to connect to.
 - **Data Digest and Header Digest**: Optionally, you can specify that a digest of data or header information be compiled during transmission to assist in troubleshooting.
 - **CHAP logon information**: If CHAP authentication is required, select this option and enter the **Target secret**.
 - **Perform mutual authentication**: If mutual CHAP authentication is configured, select this option.



NOTE: IPsec is not supported.

- 7 Click **OK**.

To support storage array controller failover, the host server must be connected to at least one iSCSI port on each controller. Repeat step 3 through step 8 for each iSCSI port on the storage array that you want to establish as failover targets (the **Target Portal** address will be different for each port you connect to).

The **Status** field on the **Targets** tab should now display as **Connected**.

- 8 Click **OK** to close the Microsoft iSCSI initiator.

Viewing the Status of Your iSCSI Connections

In MD Storage Manager, click the **iSCSI** tab and then **Configure iSCSI Host Ports** to view the status of each iSCSI port you attempted to connect to and the configuration state of all IP addresses. If either **Disconnected** or **Unconfigured** is displayed, check the following and repeat the iSCSI configuration steps:

- Are all cables securely attached to each port on the host server and storage array?
- Is TCP/IP correctly configured on all target host ports?
- Is CHAP set up correctly on both the host server and the storage array?

Set Up In-Band Management (Optional)

Out-of-band management (see *Discover the Storage Array (Out-of-Band Management Only)*) is the recommended method for managing the storage array. However, to optionally set up in-band management:

Controller 0: IP: 192.168.128.101 Subnet Mask: 255.255.255.0

Controller 1: IP: 192.168.128.102 Subnet Mask: 255.255.255.0



NOTE: The management station you are using must be configured for network communication to the same IP subnet as the PowerVault MD3000i iSCSI host ports.

- 1 Establish an iSCSI session to the MD3000i RAID storage array.
- 2 Restart the SMagent service.
- 3 Launch MD Storage Manager, and then click **New**.



NOTE: When you set up the first storage array management, the **Add New Storage Array** window appears.

- 4 Select **Manual** and click **OK**.
- 5 Select **In-band management** and enter the host server name(s) or IP address(es) of the attached host that is running the MD Storage Manager software.
- 6 Click **Add**.

In-band management should now be successfully configured.

Installing and Configuring a Failover Cluster

You can configure the operating system services on your Windows Server failover cluster, after you have established the private and public networks and have assigned the shared disks from the storage array to the cluster nodes. The procedures for configuring the failover cluster are different depending on the Windows Server operating system you use.

For more information on deploying your cluster with a specific variant of the Windows Server operating system (for example: Windows Server 2003 or Windows Server 2008), see the *Dell Failover Clusters with Microsoft Windows Server Installation and Troubleshooting Guide* on the Dell Support website at support.dell.com.

Troubleshooting

This appendix provides troubleshooting information for your cluster configurations.

Table A-1 describes general cluster problems you may encounter and the probable causes and solutions for each problem.

Table A-1. General Cluster Troubleshooting

Problem	Probable Cause	Corrective Action
The nodes cannot access the storage system, or the cluster software is not functioning with the storage system.	The storage system is not cabled properly to the nodes or the cabling between the storage components is incorrect.	Ensure that the cables are connected properly from the node to the storage system. See "Cabling Your Cluster Hardware" on page 17 for more information.
	One of the cables is faulty.	Replace the faulty cable.
	Host Group or Host-to-Virtual Disk Mappings is not created correctly.	Verify the following: <ul style="list-style-type: none"> • Host Group is created and the cluster nodes are added to the Host Group. • Host-to-Virtual Disk Mapping is created and the virtual disks are assigned to the Host Group containing the cluster nodes.
	The CHAP password entered is wrong.	If CHAP is used, enter correct user name and password.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
One of the nodes takes a long time to join the cluster. OR One of the nodes fail to join the cluster.	The node-to-node network has failed due to a cabling or hardware failure. Long delays in node-to-node communications may be normal. One or more nodes may have the Internet Connection Firewall enabled, blocking Remote Procedure Call (RPC) communications between the nodes.	Check the network cabling. Ensure that the node-to-node interconnection and the public network are connected to the correct NICs. Verify that the nodes can communicate with each other by running the ping command from each node to the other node. Try both the host name and IP address when using the ping command. Configure the Internet Connection Firewall to allow communications that are required by the Microsoft® Cluster Service (MSCS) and the clustered applications or services. See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at support.microsoft.com for more information.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Attempts to connect to a cluster using Cluster Administrator fail.	The Cluster Service has not been started. A cluster has not been formed on the system. The system has just been booted and services are still starting.	Verify that Cluster Service is running and that a cluster has been formed. Use the Event Viewer and look for the following events logged by the Cluster Service: Microsoft Cluster Service successfully formed a cluster on this node. OR Microsoft Cluster Service successfully joined the cluster. If these events do not appear in Event Viewer, see the Microsoft Cluster Service Administrator's Guide for instructions on setting up the cluster on your system and starting the Cluster Service.
	The cluster network name is not responding on the network because the Internet Connection Firewall is enabled on one or more nodes	Configure the Internet Connection Firewall to allow communications that are required by Microsoft Cluster and the clustered applications or services. See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at support.microsoft.com for more information.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
You are prompted to configure one network instead of two during Microsoft Failover Cluster installation.	The TCP/IP configuration is incorrect.	The node-to-node network and public network must be assigned static IP addresses on different subnets. For more information about assigning the network IPs, see "Assigning Static IP Addresses to Your Cluster Resources and Components" in the <i>Dell Failover Clusters with Microsoft Windows Server 2003 Installation and Troubleshooting Guide</i> .
	The private (point-to-point) network is disconnected.	Ensure that all systems are powered on so that the NICs in the private network are available.
Using Microsoft Windows NT® 4.0 to remotely administer a Windows Server® 2003 cluster generates error messages.	Normal. Some resources in Windows Server 2003 are not supported in Windows NT 4.0.	It is strongly recommended that you use Microsoft Windows® XP Professional or Windows Server 2003 for remote administration of a cluster running Windows Server 2003.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Unable to add a node to the cluster.	The new node cannot access the shared disks. The shared disks are enumerated by the operating system differently on the cluster nodes.	Ensure that the new cluster node can enumerate the cluster disks using Windows Disk Administration. If the disks do not appear in Disk Administration, check the following: <ul style="list-style-type: none">• Check all cable connections• Check all zone configurations• Check the Access Control settings on the attached storage systems• Use the Advanced with Minimum option.
	One or more nodes may have the Internet Connection Firewall enabled, blocking RPC communications between the nodes.	Configure the Internet Connection Firewall to allow communications that are required by the Microsoft Cluster and the clustered applications or services. See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at support.microsoft.com for more information.
The disks from one of the cluster nodes on the shared cluster storage appear unreadable or uninitialized in Windows Disk Administration.	This situation is normal if you stopped the Cluster Service. If you are running Windows Server 2003, this situation is normal if the cluster node does not own the cluster disk.	No action required.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Microsoft Cluster does not operate correctly on a cluster running Windows Server 2003 and the Internet Firewall enabled.	The Windows Internet Connection Firewall is enabled, which may conflict with Microsoft Cluster.	<p>Perform the following steps:</p> <ol style="list-style-type: none">1 On the Windows desktop, right-click My Computer and click Manage.2 In the Computer Management window, double-click Services.3 In the Services window, double-click Cluster Services.4 In the Cluster Services window, click the Recovery tab.5 Click the First Failure drop-down arrow and select Restart the Service.6 Click the Second Failure drop-down arrow and select Restart the Service.7 Click OK. <p>For information on how to configure your cluster with the Windows Internet Connection Firewall enabled, see Microsoft Base (KB) articles 258469 and 883398 at the Microsoft Support website at support.microsoft.com and the Microsoft Windows Server 2003 Technet website at www.microsoft.com/technet.</p>

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Public network clients cannot access the applications or services that are provided by the cluster.	One or more nodes may have the Internet Connection Firewall enabled, blocking RPC communications between the nodes.	Configure the Internet Connection Firewall to allow communications that are required by the Microsoft Cluster and the clustered applications or services. See Microsoft Knowledge Base article KB883398 at the Microsoft Support website at support.microsoft.com for more information.
Virtual Disks fail over continuously between the two storage controllers when a storage path fails.	The failback mode for the cluster node(s) is not set properly.	Set the correct failback mode on each cluster node; you must merge the PowerVault MD3000i Stand Alone to Cluster.reg file located in the <code>windows\utility</code> directory of the <i>Dell PowerVault MD3000i Resource</i> CD into the registry of each node.
Virtual Disk Copy operation fails.	The Virtual Disk Copy operation uses the cluster disk as the source disk.	To perform a Virtual Disk Copy operation on the cluster share disk, create a snapshot of the disk, and then perform a Virtual Disk Copy of the snapshot virtual disk.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Unable to assign the drive letter to the snapshot virtual disk.	The snapshot virtual disk has been erroneously mapped to the node that does not own the source disk.	Unmap the snapshot virtual disk from the node not owning the source disk, then assign it to the node that owns the source disk.
Unable to access the snapshot virtual disk.		See Using Advanced (Premium) PowerVault Modular Disk Storage Manager Features for more information.
System Error Log displays a warning with event 59 from partmgr stating that the snapshot virtual disk is a redundant path of a cluster disk.		

Troubleshooting Tools

The Modular Disk Storage Manager establishes communication with each managed array and determines the current array status. When a problem occurs on a storage array, MD Storage Manager provides several ways to troubleshoot the problem:

- **Recovery Guru** - The Recovery Guru diagnoses critical events on the storage array and recommends step-by-step recovery procedures for problem resolution. To access the Recovery Guru using MD Storage Manager, click **Support? Recover from Failure**. The Recovery Guru can also be accessed from the Status area of the Summary page.



NOTE: SAS Device Miswire Recovery Guru condition can be generated by connecting the host port of one controller to the unused expansion port on the other controller in a MD3000i enclosure.

- **Storage Array Profile** - The Storage Array Profile provides an overview of your storage array configuration, including firmware versions and the current status of all devices on the storage array. To access the Storage Array Profile, click **Support? View storage array profile**. The profile can also be viewed by clicking the Storage array profile link in the Hardware Components area of the Summary tab.

- **Status Icons** - Status icons identify the six possible health status conditions of the storage array. For every non-Optimal status icon, use the Recovery Guru to detect and troubleshoot the problem.
 - **Optimal** - Every component in the managed array is in the desired working condition.
 - **Needs Attention** - A problem exists with the managed array that requires intervention to correct it.
 - **Fixing** - A **Needs Attention** condition has been corrected and the managed array is currently changing to an Optimal status.
 - **Unresponsive** - The storage management station cannot communicate with the array, one controller, or both controllers in the storage array. Wait at least five minutes for the storage array to return to an Optimal status following a recovery procedure.
 - **Contacting Device** - MD Storage Manager is establishing contact with the array.
 - **Needs Upgrade** - The storage array is running a level of firmware that is no longer supported by MD Storage Manager.
- **Support Information Bundle** - The Gather Support Information link on the Support tab saves all storage array data, such as profile and event log information, to a file that you can send if you seek technical assistance for problem resolution.

Known Issues



NOTE: You must schedule downtime for the clusters and applications during an upgrade because an online Dell PowerVault MD3000i storage array controller firmware upgrade is not supported in a Dell PowerEdge Cluster configuration.

Issues Faced When Upgrading the Dell™ PowerVault™ MD3000i Storage Array Controller Firmware in a Dell PowerEdge™ Cluster Configuration

- Cluster(s) and application(s) using the virtual disks on an PowerVault MD3000i storage array may fail while performing an online upgrade of the storage array. Access to the virtual disks on the PowerVault MD3000i may be temporarily lost.

To resolve this issue you must take the cluster nodes offline before upgrading the PowerVault MD3000i storage array controller firmware.

- You must reboot the cluster node(s) to be able to access the virtual disks. Dell is working with its partners to resolve this issue.

Cluster Data Form

You can attach the following form in a convenient location near each cluster node or rack to record information about the cluster. Use the form when you call for technical support.

Table B-1. Cluster Configuration Information

Cluster Information	Cluster Solution
Cluster name and IP address	
Server type	
Installer	
Date installed	
Applications	
Location	
Notes	

Table B-2. Cluster Node Configuration Information

Node Name	Service Tag Number	Public IP Address	Private IP Address

Table B-3. Additional Network Information

Additional Networks

Table B-4. Storage Array Configuration Information

Array	Array Service Tag	IP Address	Number of Attached DAEs	Virtual Disks
1				
2				
3				
4				

iSCSI Configuration Worksheet

A

B

192.168.130.101 (In 0 default)	192.168.128.102 (Mgmt port)
192.168.131.101 (In 1 default)	192.168.131.102 (In 1 default)
192.168.128.101 (Mgmt port)	192.168.130.102 (In 0 default)

Mutual CHAP Secret

Target CHAP Secret

If you need additional space for more than one host server, use an additional sheet.

A	Static IP address (host server)	Subnet <small>(should be different for each NIC)</small>	Default gateway
	iSCSI port 1	_____	_____
	iSCSI port 2	_____	_____
	iSCSI port 3	_____	_____
	iSCSI port 4	_____	_____
	Management port	_____	_____
	Management port	_____	_____

B	Static IP address (storage array)	Subnet	Default gateway
	iSCSI controller 0, In 0	_____	_____
	iSCSI controller 0, In 1	_____	_____
	Management port, cntnl. 0	_____	_____
	iSCSI controller 1, In 0	_____	_____
	iSCSI controller 1, In 1	_____	_____
	Management port, cntnl. 1	_____	_____

Index

A

alerts, 44

C

CHAP, 56
 mutual, 56
 target, 56

I

initial storage array setup, 39
 alerts, 44
installing iSCSI
 Windows host and server, 40
iSCSI, 38
 terminology, 38
iSCSI configuration
 configuring ports, 44
 connect from host server, 61
 discovery, 40
 host access, 48
 in-band management, 62
 set CHAP on host server, 59
 target discovery, 47
iSCSI configuration
 worksheet, 77
iSCSI management, 41
iSNS, 40

M

MSCS
 installing and configuring, 62

O

operating system
 installing, 33

P

password, 43
post-installation
 configuration, 43
PowerVault 22xS storage system
 clustering, 52

S

status, 42
status icons, 42

W

Windows Server 2003,
 Enterprise Edition
 installing, 33

