

Dell™ Failover Clusters With
Microsoft® Windows Server® 2008
and Windows Server 2008 R2

Software Installation and Troubleshooting Guide

Notes, Cautions, and Warnings



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



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



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

Information in this document is subject to change without notice.

© 2008-2009 Dell Inc. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, *PowerVault*, and *OpenManage* are trademarks of Dell Inc.; *Active Directory*, *Microsoft*, *Windows*, *Windows Server*, and *Windows NT* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *EMC* and *Access Logix* are registered trademarks are trademarks of *EMC* Corporation

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	7
	Features of Failover Clusters Running Windows Server 2008	7
	Supported Cluster Configurations	10
	Cluster Components and Requirements	10
	Operating System.	10
	System Requirements	10
	Cluster Nodes.	10
	Cluster Storage.	12
	Other Documents You May Need	12
2	Preparing Your Systems for Clustering	15
	Cluster Configuration Overview	15
	Installation Overview	17
	Selecting a Domain Model	19
	Configuring Internal Drives in the Cluster Nodes.	19
	Installing and Configuring the Windows Operating System	20

Configuring Windows Networking	21
Assigning Static IP Addresses to Cluster Resources and Components	22
Verifying Communications Between Nodes	26
Installing the Storage Connection Ports and Drivers	27
Installing and Configuring the Shared Storage System	27
Configuring Hard Drive Letters When Using Multiple Shared Storage Systems	28
Formatting and Assigning Drive Letters and Volume Labels to the Disks	28
Configuring Your Failover Cluster	29
Validating Your Failover Cluster Configuration	30
Configuring Windows Server 2008 Failover Clustering	31
Adding a Cluster Node to the Failover Cluster.	31
Configuring the Quorum Disk in Failover Clustering	32
Configuring a Service or Application for High Availability.	33
Verifying Failover of a Clustered Service or Application	34
Modifying Properties of a Clustered Service or Application	35

3	Installing Your Cluster Management Software	37
	Microsoft Failover Cluster Management Console	37
	Running Failover Cluster Management on a Remote Console	37
	Launching Failover Cluster Management Console on a Remote Console	38
4	Understanding Your Failover Cluster	39
	Cluster Objects	39
	Cluster Networks	39
	Preventing Network Failure	39
	Network Interfaces	40
	Cluster Nodes	40
	Forming a New Cluster	41
	Joining an Existing Cluster	41
	Cluster Resources	41
	Setting Resource Properties	41
	Resource Dependencies	42
	Creating a Resource	43
	Resource Failure	44
	Replacing a Failed Disk	45
	Configuring Active and Passive Cluster Nodes	45
	Failover Policies	46
	Windows Server 2008 Cluster Configurations	46
	Failover and Failback Capabilities	52

5	Maintaining Your Cluster	55
	Adding a storage to a Failover Cluster Node	55
	Configuring Network Settings of a Failover Cluster Node	55
	Maintaining a Clustered Service or Application	56
	Starting or Stopping Cluster Service on Cluster Nodes	56
	Running chkdsk on a Clustered Disk in Maintenance Mode	57
	Displaying Event Logs for a Failover Cluster	57
6	Upgrading to a Cluster Configuration	59
	Before You Begin	59
	Supported Cluster Configurations	59
	Completing the Upgrade	60
A	Troubleshooting	61
	Index	69

Introduction

Dell™ Failover Cluster is a group of systems working together to run a common set of applications that presents a single logical system to client applications. The systems (or nodes) in the cluster are physically connected by either local area network (LAN) or wide area network (WAN) and are configured with the cluster software. If a system or the network connections in the cluster fail, the services on the active node failover to the passive node in the cluster.



NOTE: In this document, Microsoft® Windows Server® 2008 refers to either Microsoft Windows Server 2008 or Microsoft Windows Server 2008 R2. For the list of Dell-validated operating systems for a Failover Cluster, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at www.dell.com/ha.

Failover Clusters configured with Microsoft Windows Server 2008 operating systems provide high availability and scalability for mission-critical applications such as databases, messaging systems, file and print services, and virtualized workloads. If a node in a cluster becomes unavailable (as a result of failure or having been taken down for maintenance), another node in the cluster provides the same service. Users accessing the service continue their work and are unaware of any service disruption.

Windows Server 2008 includes functionality to simplify the cluster creation and administration. You can create an entire cluster in one seamless step through a wizard interface.

Features of Failover Clusters Running Windows Server 2008

The Failover Cluster running Windows Server 2008 implements up to 16 nodes in a cluster, depending on the storage array used, and provides the following features:

- A shared storage bus featuring Fibre Channel, Serial Attached SCSI (SAS), or Internet Small Computer System Interface (iSCSI) technology
- High availability of 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 node or storage system without taking the entire cluster offline

The services and capabilities that are included with Failover Clusters running Windows Server 2008 are:

- The Failover Cluster Management Interface — The Failover Cluster Management Interface is a task-oriented tool. To access the management interfaces, **Microsoft Management Console 3.0** and **cluadmin.msc**, go to **Start**→**Programs**→**Administrative Tools**.
- The **Validate a Configuration Wizard** — The cluster tools in Windows Server 2008 include the built-in cluster **Validate a Configuration** wizard to help detect the issue of a cluster failing due to configuration complexity. The **Validate a Configuration** wizard runs a set of tests on the systems in a cluster, and performs the following functions:
 - Checks the software inventory
 - Tests the network and attached storage
 - Validates system configuration
- New method to create clusters — You can install the Failover Clustering feature through the **Initial Configurations Task (ICT)** interface or with the **Server Manager** interface in **Administrative Tools**. You can also uninstall clustering using **Server Manager** interface. For systems running Windows Server 2008, you must use the **Add Feature Wizard** to install the Failover Clustering feature.
- Migrating legacy clusters — You can migrate your cluster that is running the Windows Server 2003 operating system to the Windows Server 2008 operating system. To access the migration functionality in Windows Server 2008, see the **Migrate Services and Applications** wizard. After you run the **Migrate Services and Applications** wizard, a report containing information about the migration tasks is created.



NOTE: You cannot configure nodes running the Windows Server 2003 operating system and nodes running the Windows Server 2008 operating system in the same cluster. In addition, Failover Cluster nodes must be joined to an Microsoft Active Directory® based domain and not a Windows NT 4.0-based domain.

- Improvements in Scoping and Managing Shares — The process of creating a highly-available share with Failover Cluster running Windows Server 2008 is very simple when you use the **Add a Shared Folder** wizard. You can also use the **Browse** button to quickly and reliably identify the folder you want to use for the highly-available share.
- Better Storage and Backup Support — The architecture of Failover Cluster running Windows Server 2008 has undergone storage related changes to improve stability and scalability.
- Enhanced Maintenance Mode — Use the **Maintenance** mode to perform maintenance and administrative tasks; like Volume Snapshots, ChkDsk, and so on; on the cluster disk resources. The **Maintenance** mode turns off cluster health monitoring on the cluster disk for a period of time so that it does not fail while maintenance is in-progress on the cluster disk.
- Superior Scalability — The Failover Cluster running Windows Server 2008 x64 can support 16 nodes. The Failover Cluster running Windows Server 2008 can also supports disks which use GUID Partition Table (GPT) disk partitioning system. GPT disks allow for 128 primary partitions as opposed to 4 in Master Boot Record (MBR) disks. Also, the partition size for GPT disks can be more than 2 TB (the limit for an MBR disk).
- Quorum Model — The Windows Server 2008 Failover Clustering Quorum model is redesigned to eliminate the single point of failure which existed in previous versions. The four ways to establish a quorum are:
 - No Majority - Disk Only (similar to Windows Server 2003 shared disk quorum)
 - Node Majority (similar to Windows Server 2003 Majority Node Set)
 - Node and Disk Majority
 - Node and File Share Majority
- Networking Capabilities — The Failover Cluster running Windows Server 2008 employs a new networking model which includes improved support for:
 - Geographically distributed clusters
 - Ability to have cluster nodes on different subnets
 - DHCP server to assign IP addresses to cluster interfaces
 - Improved cluster heartbeat mechanism and support for IPv6

Supported Cluster Configurations

For the list of Dell-validated hardware, firmware, and software components for a Failover Cluster running Windows Server 2008, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at www.dell.com/ha.

Cluster Components and Requirements

Your cluster requires the following components:

- Operating System
- Cluster nodes (servers)
- Cluster Storage

Operating System

Dell Failover Clusters supports Windows Server 2008 with x64 bit Enterprise Edition only. For a complete list of features, see the documentation for Windows Server 2008, x64 bit, Enterprise Edition.



NOTE: Running different operating systems in a cluster is supported only during a rolling upgrade. You cannot upgrade your Failover Cluster running a different operating system to Windows Server 2008, Enterprise x64 Edition. Only a new cluster installation is permitted for Windows Server 2008, Enterprise x64 Edition.

System Requirements

The following sections list the requirements for cluster nodes and storage systems in a Failover Cluster running Windows Server 2008.

Cluster Nodes

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

Table 1-1. Cluster Node Requirements

Component	Minimum Requirement
Cluster nodes	At least two and up to 16 PowerEdge systems running the Windows Server 2008 operating system.
RAM	At least 512 MB of RAM installed on each cluster node.

Table 1-1. Cluster Node Requirements (continued)

Component	Minimum Requirement
NICs	<p>At least two NICs: one NIC for the public network and another NIC for the private network.</p> <p>NOTE: It is recommended that the NICs on each public network are identical, and that the NICs on each private network are identical.</p>
Internal disk controller	<p>One controller connected to at least two internal hard drives for each node. Use any supported RAID controller or disk controller.</p> <p>Two hard drives are required for mirroring (RAID 1) and at least three are required for disk striping with parity (RAID 5).</p> <p>NOTE: It is strongly recommended that you use hardware-based RAID or software-based disk-fault tolerance for the internal drives.</p>
HBA ports	<ul style="list-style-type: none">• For clusters with Fibre Channel storage, two Fibre Channel HBAs per node, unless the server employs an integrated or supported dual-port Fibre Channel HBA.• For clusters with SAS storage, one or two SAS 5/E HBAs per node. <p>NOTE: Where possible, place the HBAs on separate PCI buses to improve availability and performance. For information about supported systems and HBAs, see the <i>Dell Cluster Configuration Support Matrices</i> located on the Dell High Availability Clustering website at www.dell.com/ha.</p>
iSCSI Initiator and NICs for iSCSI Access	<p>For clusters with iSCSI storage, the iSCSI Software Initiator (including iSCSI port driver and Initiator Service) is installed with the operating system.</p> <p>Two iSCSI NICs or Gigabit Ethernet NIC ports per node. NICs with a TCP/IP Off-load Engine (TOE) or iSCSI Off-load capability may also be used for iSCSI traffic.</p> <p>NOTE: Where possible, place the NICs on separate PCI buses to improve availability and performance. For information about supported systems and HBAs, see <i>Dell Cluster Configuration Support Matrices</i> on the Dell High Availability Clustering website at www.dell.com/ha.</p>

Cluster Storage


While configuring your Dell Failover Cluster with Windows Server 2008, attach all cluster nodes to a common shared storage. The type of storage array and topology in which the array is deployed can influence the design of your cluster. For example, a direct-attached SAS storage array may offer support for two cluster nodes whereas a SAN-attached Fibre Channel or iSCSI array has the ability to support sixteen cluster nodes.


A shared storage array enables data for clustered applications and services to be stored in a common location that is accessible by each cluster node. Although only one node can access or control a given disk volume at a point in time, the shared storage array enables other nodes to gain control of these volumes in the event that a node failure occurs. This also helps facilitate the ability of other cluster resources, which may depend upon the disk volume, to failover to the remaining nodes.

Additionally, it is recommended that you attach each node to the shared storage array using redundant paths. Providing multiple connections (or paths) between the node and the storage array reduces the number of single points of failure that could otherwise impact the availability of the clustered applications or services.

For details and recommendations related to deploying a Dell Failover Cluster solution with a storage array, see the "Cabling Your Cluster Hardware" section in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.

Other Documents You May Need

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

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

- The *Dell Windows Server Failover Cluster Hardware Installation and Troubleshooting Guide* provides information on specific configuration tasks that enable you to deploy the shared storage for your cluster.


- The *Dell Cluster Configuration Support Matrices* list the Dell validated hardware, firmware, and software components for a Failover Cluster environment.
- 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 of initially setting up your system.
- The HBA documentation provides installation instructions for the HBAs.
- 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.
- 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.




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 or documentation, or advanced technical reference material intended for experienced users or technicians.


Preparing Your Systems for Clustering

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

Cluster Configuration Overview

 **NOTE:** For more information on step 1, step 2, and step 9, see the "Preparing Your Systems for Clustering" section of the *Dell Failover Hardware Installation and Troubleshooting Guide* for the specific storage array on the Dell Support site at support.dell.com/manuals. For more information on step 3 to step 7 and step 10 to step 14, see this chapter.

- 1 Ensure that your site can handle the cluster's power requirements. For more information.
Contact your sales representative for information about your region's power requirements.
- 2 Install the systems, the shared storage array(s), and the interconnect switches (example: in an equipment rack), and ensure that all these components are powered on.
- 3 Deploy the operating system (including any relevant service pack and hotfixes), network adapter drivers, and storage adapter drivers (including MPIIO drivers) on each of the systems that you want to configure as 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:** To help planning and deployment of your cluster, record the relevant cluster configuration information in the "Cluster Data Form" and "Zoning Configuration Form" of Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.

- 4 Establish the physical network topology and the TCP/IP settings for network adapters on each cluster node to provide access to the cluster public and private networks.
- 5 Configure each cluster node as a member in the same Microsoft® Active Directory® Domain.
- 6 Establish the physical storage topology and any required storage network settings to provide connectivity between the storage array and the systems that are configured as cluster nodes. For more information on configuring the storage system(s), see 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 witness disk for Microsoft Windows Server® 2008 Failover Cluster. Ensure that this LUN is presented to the systems that are configured as cluster nodes.



NOTE: For security reasons, it is recommended that you configure the LUN on a single node, 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.



NOTE: For Failover Clusters configured with Windows Server 2008 operating system, 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.
- 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 Failover Clusters configured with the Windows Server 2008 operating system, 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.
- 13 Test the failover capabilities of the new resources.
- 14 Configure client systems to access the highly-available applications and services that are hosted on your Failover Cluster.

Installation Overview

This section provides installation overview procedures for configuring a cluster running the Windows Server 2008 operating system.



NOTE: The Storage management software may use different terms than those in this guide to refer to similar entities. For example, the terms "LUN" and "Virtual Disk" are often used interchangeably to designate an individual RAID volume that is provided to the cluster nodes by the storage array.

- 1 Ensure that the cluster meets the requirements as described in "Cluster Configuration Overview" on page 15.
- 2 Select a domain model that is appropriate for the corporate network and operating system.
See "Selecting a Domain Model" on page 19.
- 3 Reserve static IP addresses for the cluster resources and components, including:
 - Public network
 - Private network
 - Cluster virtual servers

Use these IP addresses when you install the Windows® operating system and Windows Server 2008 Failover Clustering (WSFC).



NOTE: WSFC supports configuring cluster IP address resources to obtain IP address from a DHCP server in addition to through static entries. It is recommended that you use static IP addresses.

- 4 Configure the internal hard drives.
See "Configuring Internal Drives in the Cluster Nodes" on page 19.
- 5 Install and configure the Windows operating system.
The Windows operating system must be installed on all the cluster nodes. Each node must have a licensed copy of the Windows operating system, and a Certificate of Authenticity.
See "Installing and Configuring the Windows Operating System" on page 20.

6 Install or update the storage connection drivers.

For more information on connecting your cluster nodes to a shared storage array, see "Preparing Your Systems for Clustering" in the Dell *Failover Cluster Hardware Installation and Troubleshooting Guide* that corresponds to your storage array on the Dell Support website at support.dell.com/manuals.

For more information on the corresponding supported adapters and driver versions, see the *Dell Cluster Configuration Support Matrices* located on the Dell High Availability Clustering website at www.dell.com/ha.

7 Install and configure the storage management software.

See the documentation included with your storage system or available at the Dell Support website at support.dell.com/manuals.

8 Configure the hard drives on the shared storage system(s).

See "Configuring and Managing LUNs" in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide corresponding to your storage array on the Dell Support website at support.dell.com/manuals.

9 Install and configure Failover Clustering feature.

See "Configuring Your Failover Cluster" on page 29.

10 Verify cluster functionality. Ensure that:

- The cluster components are communicating properly.
- Cluster Service is started.

11 Verify cluster resource availability.

Use the Failover Cluster MMC to check the running state of each resource group.

The following subsections provide detailed information for the steps in the "Installation Overview" on page 17 that is specific to the Windows Server 2008 operating system.

Selecting a Domain Model

On a cluster running the Microsoft Windows operating system, all nodes must belong to a common domain or directory model. The following configurations are supported:

- It is recommended that all nodes of High Availability applications are member systems in an Microsoft Active Directory® domain.
- All nodes are domain controllers in an Active Directory domain.
- At least one node is a domain controller in an Active Directory and the remaining nodes are member systems.



NOTE: If a node is configured as a domain controller, client system access to its cluster resources can continue even if the node cannot contact other domain controllers. However, domain controller functions can cause additional overhead, such as log on, authentication, and replication traffic. If a node is not configured as a domain controller and the node cannot contact a domain controller, the node cannot authenticate client system requests.

Configuring Internal Drives in the Cluster Nodes

If your system uses a hardware-based RAID solution and you have added new internal hard drives to your system or you are setting up the RAID configuration for the first time, you must configure the RAID array using the RAID controller's BIOS configuration utility before installing the operating system.

For the best balance of fault tolerance and performance, use RAID 1. See the RAID controller documentation for more information on RAID configurations.




NOTE: It is strongly recommended that you use hardware based RAID solution. Alternately you can use the Microsoft Windows Disk Management tool to provide software-based redundancy.

Installing and Configuring the Windows Operating System

 **CAUTION:** Windows standby mode and hibernation mode are not supported in cluster configurations. Do not enable either mode.

- 1 Ensure that the cluster configuration meets the requirements listed in "Cluster Configuration Overview" on page 15.
- 2 Cable the hardware.

 **NOTE:** Do not connect the nodes to the shared storage systems at this time.

For more information on cabling your cluster hardware and the storage array that you are using, see "Cabling Your Cluster Hardware" in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.

- 3 Install and configure the Windows Server 2008 operating system on each node.
- 4 Ensure that the latest supported version of network adapter drivers is installed on each cluster node.
- 5 Configure the public and private network adapter interconnects in each node, and place the interconnects on separate IP subnetworks using static IP addresses. See "Configuring Windows Networking" on page 21.

For information on required drivers, see *Dell Cluster Configuration Support Matrices* located on the Dell High Availability website at www.dell.com/ha.

- 6 Turn off all the cluster nodes and connect each cluster node to the shared storage.

For more information on cabling your cluster hardware and the storage array that you are using, see "Cabling Your Cluster Hardware" in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.

- 7 If required, configure the storage software.
- 8 Reboot node 1.

- 9 From node 1, go to the **Windows Disk Management** application, write the disk signature, partition the disk, format the disk, and assign drive letters and volume labels to the hard drives in the storage system.

For more information, see "Preparing Your Systems for Clustering" in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.

- 10 On node 1, verify disk accessibility and functionality on all shared disks. Verify disk access by performing the following steps on the second node:
 - a Turn on the node.
 - b Modify the drive letters to match the drive letters on node 1.


This procedure allows the Windows operating system to mount the volumes.
 - c Close and reopen **Disk Management**.
 - d Verify that the Windows operating system can access the file systems and the volume labels.
- 11 Install and configure the Failover Clustering feature from the **Server Manager**.
- 12 If required, install and setup the application programs.
- 13 Enter the cluster configuration information on the **Cluster Data Form** in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for your corresponding storage array on the Dell Support website at support.dell.com/manuals (optional).

Configuring Windows Networking

You must configure the public and private networks in each node before you install Failover Clustering on the nodes. The following subsections introduce you to some principles and procedures necessary for the networking prerequisites.

Windows Server 2008 also introduces IPv6 support for clustering. You can have both node-to-node (private) as well as node-to-client (public) communication over IPv6. For more details on using IPv6, see "Configuring IPv6 addresses for Cluster Nodes" on page 24.

Assigning Static IP Addresses to Cluster Resources and Components

 **NOTE:** WSFC supports configuring cluster IP address resources to obtain IP address from a DHCP server in addition to through static entries. It is recommended that you use static IP addresses.

A static IP address is an Internet address that a network administrator assigns exclusively to a system or a resource. The address assignment remains in effect until it is changed by the network administrator.

The IP address assignments for the cluster's public LAN segments depend on the environment's configuration. Configurations running the Windows operating system require static IP addresses assigned to hardware and software applications in the cluster, as listed in Table 2-1.

Table 2-1. Applications and Hardware Requiring IP Address Assignments

Application/Hardware	Description
Cluster IP address	The cluster IP address is used for cluster management and must correspond to the cluster name. Because each server has at least two network adapters, the minimum number of static IP addresses required for a cluster configuration is five (one for each network adapter and one for the cluster). Additional IP addresses are required when WSFC is configured with application programs that require IP addresses, such as file sharing.
Cluster-aware applications running on the cluster	These applications include Microsoft SQL Server Enterprise Edition, Microsoft Exchange Server, and Internet Information Server (IIS). For example, SQL Server Enterprise Edition requires at least one static IP address for the virtual server as SQL Server does not use the cluster's IP address. Also, each IIS Virtual Root or IIS Server instance configured for failover needs a unique static IP address.

Table 2-1. Applications and Hardware Requiring IP Address Assignments (continued)

Application/Hardware	Description
Cluster node network adapters	<p>For cluster operation, two network adapters are required: one for the public network (LAN/WAN) and another for the private network (sharing heartbeat information between the nodes).</p> <p>For more information on cabling your cluster hardware and the storage array that you are using, see "Cabling Your Cluster Hardware" in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.</p> <p>NOTE: To ensure operation during a DHCP server failure, use static IP addresses.</p>

Configuring IPv4 Addresses for Cluster Nodes

Use the static IP address assignments for the network adapters used for the cluster nodes.




NOTE: The IPv4 addresses in Table 2-2 are used as examples only.

Table 2-2. Examples of IP Address Assignments

Usage	Cluster Node 1	Cluster Node 2
Public network static IP address (for client and domain controller communications)	192.168.1.101	192.168.1.102
Public network subnet mask	255.255.255.0	255.255.255.0
Default gateway	192.168.1.1	192.168.1.1
DNS servers	Primary	Primary
	192.168.1.21	192.168.1.21
	Secondary	Secondary
	192.168.1.22	192.168.1.22


Table 2-2. Examples of IP Address Assignments (continued)

Usage	Cluster Node 1	Cluster Node 2
Private network static IP address cluster interconnect (for node-to-node communications)	10.0.0.1	10.0.0.2
Private network subnet mask	255.255.255.0	255.255.255.0

 **NOTE:** Do not configure Default Gateway, NetBIOS, WINS, and DNS on the private network.

If multiple cluster interconnect network adapters are connected to a network switch, ensure that all of the NICs configured for private network have a unique address. You can continue the IP address scheme in Table 2-2 with 10.0.0.3, 10.0.0.4, and so on for the NICs configured for the private network or NIC teams of the other clusters connected to the same switch.

You can improve fault tolerance by using network adapters that support adapter teaming or by having multiple LAN segments. To avoid communication problems, do not use dual-port NICs for the cluster interconnect.

 **NOTE:** NIC teaming is supported only on a public network, not on a private network.

Configuring IPv6 addresses for Cluster Nodes

Windows Server 2008 supports IPv6 and hence Failover Clustering also supports IPv6. Therefore, you can configure all IPv6 address resources, all IPv4 address resources, or a combination of IPv4 and IPv6 address resources.

Failover Clustering supports only IPv6 addresses that allow for dynamic registration in DNS (AAAA host records and the IP6.ARPA reverse look-up zone).

The left most bits of the IPv6 address are called the Format Prefix (FP), which indicates the specific type of IPv6 address. IPv6 accommodates many address type, including:

- Unicast addresses — Unicast addresses are used for one-to-one communication between two hosts.
- Multicast addresses — Multicast addresses are used for one-to-many communication. In Multicast addressing a single IP packet is sent to multiple hosts in a group.

- Anycast addresses — Anycast addresses used for one-to-one-of-many communication. In Anycast addressing an IP packet is sent to the nearest member in a group.

Unicast Addresses have the following types:

- 1** Global unicast addresses — This address can be identified by format prefix (FP) of 001. The global unicast addresses are equivalent to public IPv4 addresses and can be used for Public Interfaces. They are globally routable and reachable on the IPv6 portion of Internet. The 128 bits of the unicast address space can be divided into three sections:
 - Prefix — The network ID or prefix of the address, used for routing. The first 48 bits are used for prefix.
 - Subnet Identifier — A number that identifies a subnet within a site. After the Prefix, 16 bits are used for Subnet ID.
 - Interface ID — Unique identifier for a particular interface (host or other device). This ID is unique within the specific prefix and subnet. After the Subnet ID, the 64 bits are used for Interface ID.
- 2** Link local addresses — Identified by format prefix (FP) of 1111 1110 10. These addresses are used by nodes when they are communicating with neighboring nodes on the same link. Dynamic registrations do not occur for link local addresses and therefore cannot be used in a cluster.
- 3** Site Local addresses — Identified by format prefix (FP) of 1111 1110 11. These addresses are equivalent to IPv4 private address space. Use these addresses between nodes that communicate with other nodes in the same site.

Creating Separate Subnets for the Public and Private Networks

The NICs that are configured for the public and private networks and are installed in the same cluster node must reside on separate IP subnetworks. Therefore, the private network used to exchange heartbeat information between the nodes must have a separate IP subnet or a different network ID than the public network, which is used for client connections.

Configuring the Network Interface Binding Order for Clusters Running Windows Server 2008

After configuring the IP addresses for networks on your Failover Cluster, configure the Network Interface Binding Order:

- 1 Click **Start**→**Control Panel**→double-click **Network And Sharing Center**.
- 2 In the **Tasks** pane, click the **Manage Network Connections**.
The **Network Connections** window appears.
- 3 Click the **Advanced** menu and then click **Advanced Settings**.
The **Advanced Settings** window appears.
- 4 In the **Adapters and Bindings** tab, ensure that the **Public** connection is at the top of the list followed by the **Private** connection.

To change the connection order:

- a Click **Public** or **Private**.
- b Click the up-arrow or down-arrow to move the connection to the top or bottom of the **Connections** box.
- c Click **OK**.
- d Close the **Network Connections** window.



NOTE: If your cluster node has more network adapters other than the public and private network adapters, they can be listed in any order, in the **Adapters and Bindings** tab, under the **Private** connection.

Dual-Port NICs and NIC Teaming in the Private Network

Dual-port NICs and NIC teaming are not supported in the private network. They are supported only in the public network.

Verifying Communications Between Nodes

- 1 Open a command prompt on each cluster node.
- 2 At the prompt, type:
`ipconfig /all`
- 3 Press <Enter>.
All known IP addresses for each local server appear on the screen.
- 4 Issue the **ping** command from each remote system.

Ensure that each local server responds to the **ping** command. If the IP assignments are not set up correctly, the nodes may not be able to communicate with the domain. For more information on this issue, see "Troubleshooting" on page 5.

Installing the Storage Connection Ports and Drivers

Before you connect each cluster node to the shared storage:

- Ensure that an appropriate storage connection exists on the nodes.
- Ensure that the cluster nodes have a complementary technology that enables proper interaction between the nodes and shared Fibre Channel, SAS, or iSCSI storage array.
- You may also require operating system drivers and Multipath Input/Output (MPIO) drivers to ensure proper interaction between the cluster nodes and the shared storage array.

For more information, see "Preparing Your Systems for Clustering" in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.

Installing and Configuring the Shared Storage System

The shared storage array consists of disk volumes that are used in your cluster. The storage management software for each supported shared storage array provides a way to create disk volumes and assigns these volumes to all the nodes in your cluster.

For more information, see the "Preparing Your Systems for Clustering" section in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.

Configuring Hard Drive Letters When Using Multiple Shared Storage Systems

Before creating the cluster, ensure that both nodes have the same view of the shared storage systems. Because each node has access to hard drives that are in a common storage array, each node must have identical drive letters assigned to each hard drive. Using volume mount points in Windows Server 2008, your cluster can access more than 22 volumes.



NOTE: Drive letters A through D are reserved for the local system.

To ensure that hard drive letter assignments are identical:

- 1 Ensure that your cables are attached to the shared storage devices in the proper sequence.

You can view all storage devices using the Windows Server 2008 Disk Management console.

- 2 To maintain proper drive letter assignments, ensure that each storage connection port on the storage device is enumerated by each node. Also, ensure that each storage connection port is connected to the same RAID controller, storage processor, or SAN switch.

For more information on the location of the RAID controllers or storage processors on your shared storage array, see "Cabling Your Cluster Hardware" in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.

- 3 Go to "Formatting and Assigning Drive Letters and Volume Labels to the Disks" on page 28.

Formatting and Assigning Drive Letters and Volume Labels to the Disks

- 1 Open Disk Management on the Server Manager.
- 2 Ensure that the shared disks are **Online**. Format the disks, assign the drive letters and volume labels on node 1 by using the **Windows Disk Management** utility.

For example, create volumes labeled "Volume Y" for disk Y and "Volume Z" for disk Z.

- 3 Perform the following steps on all the node(s):
 - a Open **Disk Management** on the **Server Manager**.
 - b Assign the drive letters for the drives.
 - c Reassign the drive letter, if necessary.

To reassign the drive letter:

- With the mouse pointer on the same icon, right-click and select **Change Drive Letter and Path** from the submenu.
- Click **Change**, select the letter you want to assign the drive (for example, Z), and then click **OK**.
- Click **Yes** to confirm the changes.

If the cables are connected properly, the drive order is the same as is on each node, and the drive letter assignments of all the cluster nodes follow the same order as is on node 1. The volume labels can also be used to double-check the drive order by ensuring that the disk with volume label "Volume Z" is assigned to drive letter Z and so on for each disk on each node. Assign drive letters on each of the shared disks, even if the disk displays the drive letter correctly.

For more information about the storage array management software, see your storage array documentation located on the Dell Support website at support.dell.com/manuals.

Configuring Your Failover Cluster

The **Failover Clustering** feature included in Windows Server 2008 installs the Cluster Service. The Cluster Service performs the basic cluster functionality including membership, communication, and failover management.



NOTE: For installing and performing operations on Failover Clusters, you must be logged on to the domain with an account that has local Administrator rights on all the systems that you want to configure as cluster nodes.



NOTE: Dell supports running Windows Server 2008 Failover Clustering in a non-redundant configuration on Dell PowerVault MD3000 storage system using single port SAS controller only if the **WS08clusprepcfg.xml** file is installed. Ensure that this file is installed before you run the **Validate a Configuration** wizard. For more information on downloading the file, see the Dell Support site at support.dell.com/manuals.

To install Failover Clustering feature:

- 1 Click **Start**→**Administrative Tools**→**Server Manager**, if **Server Manager** is not running,
- 2 If prompted for permission to continue, click **Continue**.
- 3 Under **Features Summary**, click **Add Features**.
- 4 In the **Add Features Wizard**, click **Failover Clustering and Multipath I/O** and then click **Install**.
- 5 Click **Close** to close the wizard.
- 6 Repeat the step 1 to step 5 for each system that you want to configure as cluster node.

When you install **Failover Clustering**, the Cluster Service starts on each node and responds automatically if one of the nodes fails or goes offline. To provide application failover for the cluster, the Failover Clustering feature must be installed on each cluster node.

For more information, see "Understanding Your Failover Cluster" on page 5.

Validating Your Failover Cluster Configuration



NOTE: It is strongly recommended that you validate your configuration by running all tests in the **Validate a Configuration Wizard** before you create a Failover Cluster. By running the tests, you can confirm that your hardware and settings are compatible with Failover Clustering. You can run the tests on a set of systems and storage devices either before or after you have configured them in a Failover Cluster. Dell and Microsoft support a cluster configuration as long as it has passed all the tests included in **Validate a Configuration Wizard**.

To run **Validate a Configuration** wizard:

- 1 After you have installed Failover Clustering Feature on all your nodes, connected your storage and configured drive letters and volumes, open the Failover Cluster MMC from **Start**→**Administrative Tools**→**Failover Cluster Management**.
- 2 Go to the **Action** tab and select **Validate a Configuration Wizard**.
- 3 Click **Next** and on the **Select Servers or a Cluster** window. Enter the names of the systems that you want to be a part of your cluster and click **Add**. Click **Next**.

- 4 On the **Testing Options** window, select the specific tests you want to run or select **Run all tests** (recommended).
- 5 On the last screen of the **Validate a Configuration** wizard, click **Next** to confirm.

This runs a list of validation tests and prompts the errors or warnings that are present in the configuration in the form of a **Summary** window.

Configuring Windows Server 2008 Failover Clustering

The cluster setup files are installed on the system disk when you add **Failover Clustering** feature from the **Add Feature** wizard. To configure the Failover Cluster feature in Windows Server 2008:

- 1 In the Failover Cluster Management console, select **Failover Cluster Management**, under **Management** option, click **Create a Cluster**.
- 2 Follow the instructions in the wizard, and click **Finish**.
- 3 To view a report of the tasks that the wizard performed, click **View Report**.

Adding a Cluster Node to the Failover Cluster

To add a cluster node:

- 1 Turn on the node(s).
- 2 Run the **Validate a Configuration** on the new node along with the existing node(s).
- 3 If the cluster you want to configure is not displayed in the console tree of the **Failover Cluster Management** screen, right-click **Failover Cluster Management**, click **Manage a Cluster**, and select or specify the cluster that you want to configure.
- 4 Click **Add Node** in the **Actions** pane. Follow the instructions on the screen to complete the process.

Configuring the Quorum Disk in Failover Clustering

The Quorum configuration determines the maximum number of failures a Failover Cluster can sustain without stopping the Cluster Service. In your cluster configured with the Windows Server 2008 operating system, you do not have to configure a shared storage resource for the Quorum disk. The following terms are commonly used in the concept of Quorum resource in the Windows Server 2008 operating system:

- The Quorum disk in the Failover Cluster must have the *votes* from a majority of nodes in the cluster
- Depending on the quorum configuration you select, the term *failure* may refer to the failure of a node, failure of a disk designated as the *witness*, or the failure of a file share designated as the *witness*.

Failover clustering provides the following four options for quorum configuration:

- 1 Node and Disk Majority** - The cluster nodes, along with a disk resource on shared storage that is designated as a *witness*, are given one vote each. The cluster is online even if half of the total number of nodes rounded up are alive and the witness disk is online. If the witness disk fails, the cluster can sustain the failure of one node less than half of the total number of nodes rounded up.

If N denotes the number of nodes in your cluster which have one vote each and one vote is allotted to the witness disk, then, the total number of votes for the cluster to be up:


$$\{N+1 \text{ (witness disk) } + 1\}/2$$




NOTE: If your cluster has odd number of Cluster Nodes, truncate the decimal output to the next whole number. For example, if your cluster has 5 nodes, then the total number of votes for the cluster is $5+1+1= 7/2 = 3.5$. Round off 3.5 to next whole number which is 4. Therefore, the total number of votes required for the cluster to be online is 4.


- 2 Node and File Share Majority** - This model is similar to Node and Disk Majority in behavior but the witness disk is replaced by a File share to act as the witness. You can configure the File Share Witness as any network share that is accessible by all the cluster nodes.

- 3 **Node Majority:** The cluster nodes determine the maximum number of failures that the cluster can sustain (similar to the Majority Node Set feature in the Windows Server 2003 operating system). A Node Majority Cluster can sustain the failure of one node less than half the total number of nodes rounded up in the cluster.
- 4 **No Majority - Disk only:** This model cannot sustain the failure of the quorum disk (similar to the Shared Disk Quorum feature in the Windows Server 2003 operating system), which becomes a single point of failure, and is therefore not recommended.

 **NOTE:** It is recommended that you use **Node and Disk Majority** option for quorum configurations.


 **NOTE:** If your cluster setup has odd number of nodes while creating the Failover Cluster then WSFC uses **Node Majority** quorum model. You can manually change the quorum model to **Node and Disk Majority**.

Creating a LUN for the Witness Disk for Node and Disk Majority and No Majority Models

 **NOTE:** It is recommended that you create a separate LUN of 512 MB (approximately) for the witness disk for the **Node and Disk Majority** and **No Majority** quorum models. The witness disk in these models has a copy of the cluster configuration.

When you create the LUN for the witness disk on the shared storage ensure that you have:

- Formatted the LUN with NTFS.
- Used the LUN exclusively for the Cluster logs.
- Not stored any application data or user data on the witness disk.

 **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 they do not provide the level of availability required for the witness disk.

Configuring a Service or Application for High Availability

Failover Clustering in Windows Server 2008 allows you to configure a service or an application for high availability by running the **Configure a Service or Application for High Availability** wizard. To configure a highly-available service or application:

- 1 In the **Failover Cluster Management** console, right-click on **Failover Cluster Management**, click **Manage a Cluster**, and select or specify the cluster you want to configure.
- 2 Click **Services and Applications** and click **Configure a Service or Application** under **Actions**.
- 3 Follow the instructions in the wizard to specify the service or application that you want to configure for high availability. When prompted, enter the following information:
 - A name for the clustered service or application. This name is registered in DNS and associated with the IP address for this clustered service or application.
 - Any IP address information that is not automatically supplied by your DHCP settings.
 - The storage volume or volumes that the clustered service or application should use.
 - Any specific information for the service or application that you are configuring.
- 4 Click **View Report** if you want to see the report for the tasks performed.

Verifying Failover of a Clustered Service or Application

After you have configured a service or application for high availability using the instructions above, you can verify the failover of the clustered service or application:

- 1 In the **Failover Cluster Management** console, right-click on **Failover Cluster Management**, click **Manage a Cluster**, and select or specify the cluster you want to configure.
- 2 Under **Services and Applications**, click on the service or application for which you want to test failover.
- 3 Under **Actions**, click **Move** this service or application to another node.
- 4 Check if the service or application comes online on the other node.

Modifying Properties of a Clustered Service or Application

Failover Clustering allows you to modify the failover behavior of a clustered service or application. To modify the clustered service properties:

- 1 Right-click on the clustered service or application and click on **Properties**.
- 2 Select from the two tabs **General** and **Failover**. The following options available under these tabs:
 - **Preferred owners** - This is an option under **General** tab. This option allows you to designate one or more nodes in your cluster as the Preferred Owner for the clustered service or application. It also allows you to set the order of nodes.
 - **Prevent Failback/Allow Failback** - This is an option under **General** tab. This option enables you to specify if the service or application automatically fails back to the most preferred owner.
 - **Maximum failures in the specified period and Period (hours)** - This is an option under **Failover** tab. This option enables you to specify the number of times the cluster service must attempt to restart or failover the service or application in a specified time period. If the service or application fails more than the maximum limit specified, the service or application is left in the failed state.

Installing Your Cluster Management Software

This section provides information on configuring and administering your cluster using Microsoft® Failover Cluster Management console.

Microsoft Failover Cluster Management Console

Failover Cluster Management console is Microsoft's tool for configuring and administering a cluster. The following sections describe the procedures to run Failover Cluster Management console locally on a cluster node and to install the tool on a remote console. To launch the Failover Cluster Management console, click **Start**→**Programs**→**Administrative Tools**→**Failover Cluster Management**

Running Failover Cluster Management on a Remote Console

You can administer and monitor the Cluster Service remotely by installing the Remote Server Administration Tools (RSAT) and Failover Clustering feature on a remote console (or management station) running the Microsoft Windows® operating system.

The RSAT tool for Failover Clustering includes the **Failover Cluster Manager** console and the **cluster.exe** command line tool.

To install **Remote Server Administration Tools** package on a remote console:

- 1 On a system running any Windows operating system that you want to configure as the remote console:
Click **Start**→**Server Manager**→**Features**→**Add Features**.
- 2 Expand the **Remote Server Administration Tools** tab and then expand the **Feature Administration Tools**.
- 3 Select **Failover Clustering Tools** option and click **Next**.
- 4 Click **Install**.

After a few minutes, the **Remote Server Administration Tools** package is successfully installed.

Launching Failover Cluster Management Console on a Remote Console

Perform the following steps on the remote console:

- 1** Ensure that the **Failover Clustering Tools** is installed from RSAT on the system.
- 2** Click **Start** and select **Administrative Tools**.
- 3** Select **Failover Cluster Management**.
- 4** Click **Action** tab in the console and select **Manage a Cluster** option.
- 5** Provide the name of the cluster you want to manage and click **OK**.

Understanding Your Failover Cluster

Cluster Objects

Cluster objects are the physical and logical units managed by a cluster. Each object is associated with the following:

- Properties that define the object and its behavior within the cluster
- A set of cluster control codes used to manipulate the object's properties
- A set of object management functions to manage the object through Microsoft® Windows Server® 2008 Failover Cluster (WSFC).

Cluster Networks

A cluster network provides a communications link between the cluster nodes (private network), the client systems in a local area network (public network), or a combination of the above (public-and-private network).

Preventing Network Failure

When you install the **Failover Clustering** feature provided by the Microsoft Windows Server 2008 operating system, identify the public and private network segments connected to your cluster nodes. To ensure cluster failover and non-interrupted communications, perform the following procedures:

- Configure the private network for internal communications.
- Configure the public network for all communications to provide a redundant path if all of the private networks fail.
- Configure subsequent network adapters for client system use only or for all communications.

You can set priorities and roles of the networks when you install the Failover Clustering feature.

Network Interfaces

You can use the **Failover Cluster Management** console to view the state of all cluster network interfaces.

Cluster Nodes

A cluster node is a system in a cluster running the Microsoft Windows® operating system and WSFC.

Each node in a cluster:

- Attaches to one or more cluster storage devices
- Communicates with the other nodes through network adapters
- Is aware of systems that join or leave the cluster
- Is aware of the resources that are running on each node
- Is grouped with the remaining nodes under a common cluster name, which is used to access and manage the cluster

Table 4-1 defines states of a node during cluster operation.

Table 4-1. Node States and Definitions

State	Definition
Down	The node is not actively participating in cluster operations.
Joining	The node is becoming an active participant in the cluster operations.
Paused	The node is actively participating in cluster operations but cannot take ownership of resource groups or bring resources online.
Up	The node is actively participating in all cluster operations, including hosting cluster groups.
Unknown	The node state cannot be determined.

When Failover Clustering is configured on a node, the administrator chooses whether that node forms its own cluster or joins an existing cluster. When the cluster service is started, the node searches for other active nodes on networks that are enabled for internal cluster communications.

Forming a New Cluster

Failover Clustering maintains a current copy of the cluster database on all active nodes. If a node cannot join a cluster, the node attempts to gain control of the witness disk resource in Node and Disk Majority model and forms a cluster. The node uses the recovery logs in the quorum resource to update its cluster database.

Joining an Existing Cluster

A node can join a cluster if it can communicate with another active node in the cluster. When a node joins a cluster, the node is updated with the latest copy of the cluster database. Failover Clustering validates the node's name, verifies version compatibility, and the node joins the cluster.

Cluster Resources

A cluster resource is any physical or logical component that can be:

- Brought online and taken offline
- Managed in a cluster
- Hosted by one managed system at a time

Setting Resource Properties

Failover Clustering allows you to set the properties for any resource that is a part of the cluster.

To change or modify resource properties:

- 1 Right-click the resource you want to modify and click on **Properties**. The resource properties are listed under four tabs:
 - **General**- Allows you to rename the resource or use the **Repair** option to stop using a failed disk and assign a different disk.
 - **Dependencies**- Failover Clustering uses the resource dependencies list when bringing resources online and offline.

For example, if a group with a physical disk and a file share is brought online together, the physical disk containing the file share must be brought online before the file share. Under the **Dependencies** tab, you can specify the resources that must be brought online before your resource can come online. You can add multiple dependencies using

AND or OR. If you use AND, all the dependent resources must come online before your resource can come online. If you use OR, any one of the dependent resources must be online before your resource can come online.

- **Policies-** Allows you to define your desired response to a failure of your resource. You can also specify the Pending time-out value here which is the length of time your resource can take to change states between online and offline before the Cluster Service puts it in a **Failed** state.
- **Advanced Policies-** Allows you to select the possible nodes that can own the resource. The cluster also checks the health of the resource using either basic health check or thorough health check. You can define the **Basic resource health check interval** and **Thorough resource health check interval** in **Advanced Policies** window.

Resource Dependencies

Failover Clustering uses the **Resource Dependencies** list when bringing resources online and offline. For example, if a group with a physical disk and a file share is brought online together, the physical disk containing the file share must be brought online before the file share. Table 4-2 shows resources and their dependencies.

Table 4-2. Cluster Resources and Required Dependencies

Resource	Required Dependencies
File share	Network name (only if configured as a distributed file system [DFS] root)
IP address	None
Network name	IP address that corresponds to the network name
Physical disk	None

A dependent resource requires another resource to operate. Table 4-3 describes resource dependencies.

Table 4-3. Resource Dependencies

Term	Definition
Dependent resource	A resource that depends on other resources.
Dependency	A resource on which another resource depends.
Dependency tree	A series of dependency relationships or hierarchy. The following rules apply to a dependency tree: <ul style="list-style-type: none">• A dependent resource and its dependencies must be in the same group.• A dependent resource is taken offline before its dependencies and brought online after its dependencies, as determined by the dependency hierarchy.

Creating a Resource

- 1 Open the Failover Clustering console, right-click on the resource that you want to modify, and click **Properties**.
- 2 Under the **Dependencies** tab, you can specify the resources that must be brought online before your resource can come online. You can add multiple dependencies using **AND** or **OR**.

If you use **AND**, all the dependent resources must come online before your resource can come online. If you use **OR**, any one of the dependent resources must come online before your resource can come online.

- 3 To view the **Resource Dependency** diagram, right-click on the resource and select the **Show Dependency Report** option.



NOTE: You must configure the required dependencies before you create the resource.

Resource Failure

Failover Clustering periodically checks if a resource is functioning properly using either basic health check or thorough health check.

- 1 In the Failover Clustering console, right-click on the resource you want to modify, and click **Properties**.
- 2 Under the **Advanced Policies** tab, you can define the **Basic resource health check interval** and **Thorough resource health check interval**.

The **Thorough resource health check interval** requests a more thorough check of the resource's state and therefore is typically longer than the **Basic resource health check interval**.



NOTE: Do not adjust the **Basic resource health check interval** and **Through resource health check interval settings** unless instructed to do so by Dell Technical Support.

Adjusting the Resource Failure Policies

To define your desired response to a failure of your resource:

- 1 Right-click on the resource and click **Properties** and click the **Policies** tab.
- 2 In the Policies tab you can specify the **Pending time-out** value.

Pending time-out value is the length of time your resource can take to change states between online and offline before the Cluster Service puts it in a **Failed** state.

Additionally you can set the following failure policies:

- If the resource fails, you can choose not to restart the resource on the current node or attempt restart on the current node. You can specify the **period for restarts** and the **maximum restarts** in the specified time period.
- If restart is unsuccessful, you can choose to failover all resources along with their dependent resources.
- If the cluster service exceeds the maximum number of restart attempts within the specified time period and the failed resource has not been restarted, Failover Clustering considers the resource to be failed.



NOTE: To configure the **Looks Alive**, **Is Alive**, **Threshold**, and **Period** values for a particular resource, see "Setting Resource Properties" on page 7.



NOTE: Do not adjust the **Basic resource health check interval** and **Through resource health check interval settings** unless instructed to do so by Dell technical Support.

Replacing a Failed Disk

If a disk in a Failover Cluster has failed, you can assign a different disk. To replace the failed disk:

- 1 Right-click on the resource and click **Properties**.
- 2 In the **General** tab, click **Repair**, and select a new disk that you want to use. The new disk that you assign must be one that can be clustered.



NOTE: The **Repair** option does not recover data. You can restore data to the disk before using the **Repair** option.

File Share Resource Type

If you want to use your cluster solution as a high-availability file server, select one of the following types of file share for your resource:

- Basic file share — Publishes a file folder to the network under a single name.
- Share subdirectories — Publishes several network names—one for each file folder and all of its immediate subfolders. This method is an efficient way to create large numbers of related file shares on a file server.
- DFS root — Creates a resource that manages a stand-alone DFS root. Fault-tolerant DFS roots cannot be managed by this resource. A DFS root file share resource has required dependencies on a network name and an IP address. The network name can be either the cluster name or any other network name for a virtual server.

Configuring Active and Passive Cluster Nodes

Active nodes process application requests and provide client services. Passive nodes are backup nodes that ensure that client applications and services are available if a hardware or software failure occurs. Cluster configurations may include both active and passive nodes.



NOTE: Passive nodes must be configured with appropriate processing power and storage capacity to support the resources that are running on the active nodes.

Your cluster solution supports variations of active/active and active/passive configurations.

Cluster solutions running the Windows Server 2008 operating system can support up to sixteen nodes in multiple configurations.

An active/active configuration contains virtual servers running separate applications or services on each node. When an application is running on node 1, the remaining node(s) do not have to wait for node 1 to fail. Those node(s) can run their own cluster-aware applications (or another instance of the same application) while providing failover for the resources on node 1. For example, multiway failover is an active/active failover solution because running applications from a failed node can migrate to multiple active nodes in the cluster. However, you must ensure that adequate resources are available on each node to handle the increased load if one node fails.

In an active/passive configuration, one or more active cluster nodes are processing requests for a clustered application while the passive cluster nodes only wait for the active node(s) to fail.

Failover Policies

When implementing a failover policy, configure failback if the cluster node lacks the resources (such as memory or processing power) to support cluster node failures.

Windows Server 2008 Cluster Configurations

Cluster configurations running Windows Server 2008 provide the following failover policies:

- N (number of active nodes) + I (number of inactive nodes) failover
- Failover pair
- Multiway failover
- Failover ring

Table 4-4 provides an overview of the failover policies implemented with Windows Server 2008. For more information, see the sections that follow this table.

Table 4-4. Windows Server 2008 Failover Policies

Failover Policy	Description	Advantage	Disadvantage(s)
N + I	One or more nodes provides backup for multiple systems.	Highest resource availability.	<ul style="list-style-type: none">• May not handle more than one backup node failure.• May not fully utilize all of the nodes.
Failover pair	Applications can failover between the two nodes.	Easy to plan the capacity of each node.	Applications on the pair cannot tolerate two node failures.
Multiway	Running applications migrate to multiple nodes in the cluster.	Application load-balancing.	Must ensure that the failover nodes have ample resources available to handle the additional workload.
Failover ring	Running applications migrate to the next preassigned node.	Easy to scope node capacity for one server failure.	The next node for failover may not have ample resources to handle the workload.

N + I Failover

N + I failover is an active/passive policy where dedicated passive cluster node(s) provide backup for the active cluster node(s). This solution is best for critical applications that require dedicated resources. However, backup nodes add a higher cost of ownership because they remain idle and do not provide the cluster with additional network resources.

Figure 4-1 shows an example of a 6 + 2 (N + I) failover configuration with six active nodes and two passive nodes. Table 4-5 provides an N + I failover matrix for Figure 4-1.

Figure 4-1. Example of an N+1 Failover Configuration for an Eight-Node Cluster




Table 4-5. Example of an N+1 Failover Configuration for an Eight-Node Cluster

Cluster Resource Group	Primary Node	AntiAffinityClassNames Value
A	Node 1	AString
B	Node 2	AString
C	Node 3	AString
D	Node 4	AString
E	Node 5	AString
F	Node 6	AString

Configuring Group Affinity

On $N + I$ (active/passive) Failover Clusters running Windows Server 2008, some resource groups may conflict with other groups if they are running on the same node. For example, running more than one Microsoft Exchange virtual server on the same node may generate application conflicts. Use Windows Server 2008 to assign a public property (or attribute) to a dependency between groups to ensure that they failover to similar or separate nodes. This property is called *group affinity*.

Group affinity uses the `AntiAffinityClassNames` public property, which ensures that designated resources are running on *separate nodes*, if possible.

For example, in Table 4-5, the `AntiAffinityClassNames` string for cluster resource group A and group B are identical (`AString`), which indicates that these groups are assigned to run on separate nodes, if possible. If node 1 fails, resource group A fails over to the next backup node (node 7). If node 2 then fails, because their `AntiAffinityClassNames` string value (`AString`) identifies group A and group B as conflicting groups, group B skips node 7 and instead fails over to node 8.

To set the public property for the cluster groups shown in Table 4-5:

- 1 Open a command prompt.
- 2 Type the following:

```
cluster group "A" /prop AntiAffinityClassNames=
"AString"
```

- 3 Repeat step 2 for the remaining cluster groups.

To specify group affinity in your $N + I$ cluster configuration, use the **Cluster Data Form** in the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com.

Failover Pair

Failover pair is a policy in which each application can failover between two specific nodes in a multinode cluster. The **Possible Owners** list in Cluster Administrator determines which nodes run the failed over applications.

If you have applications that run well on two-node cluster, and you want to migrate these applications to Windows Server 2008, failover pair is a good policy. This solution is easy to plan and administer, and applications that do not run well on the same server can easily be moved into separate failover pairs. However, in a failover pair, applications on the pair cannot tolerate two node failures.

Figure 4-2 shows an example of a failover pair configuration. Table 4-6 provides a failover configuration for the cluster shown in Figure 4-2.

Figure 4-2. Example of a Failover Pair Configuration




Table 4-6. Example of a Failover Pair Configuration for a Four-Node Cluster

Cluster Resource Group	Possible Owners List
App1	1, 2
App2	3, 4

Multiway Failover

Multiway failover is an active/active policy where running applications from a failed node migrate to multiple nodes in the cluster. This solution provides automatic failover and load-balancing. Ensure that the failover nodes have sufficient resources to handle the workload. Figure 4-3 shows an example of four-node multiway failover configuration.

Table 4-7 shows a four-node multiway failover configuration for the cluster shown in Figure 4-3. For each resource group, the failover order in the **Preferred Owners** list in Failover Cluster Management console outlines the order that you want that resource group to failover. In this example, node 1 owns applications A, B, and C. If node 1 fails, applications A, B, and C failover to cluster nodes 2, 3, and 4. Configure the applications similarly on nodes 2, 3, and 4.

When implementing multiway failover, configure failback to avoid performance degradation. See "Understanding Your Failover Cluster" on page 5 for more information.

Figure 4-3. Example of a Four-Node Multiway Failover Configuration




Table 4-7. Example of a Four-Node Multiway Failover Configuration

Application	Failover Order in the Preferred Owners List
A	Node 2
B	Node 3
C	Node 4


Failover Ring

Failover ring is an active/active policy where all running applications migrate from the failed node to the next preassigned node in the Preferred Owners List. If the failing node is the last node in the list, the failed node's applications failover to the first node.

While this type of failover provides high availability, ensure that the next node for failover has sufficient resources to handle the additional workload.

Figure 4-4 shows an example of a failover ring configuration.

Figure 4-4. Example of a Four-Node Failover Ring Configuration



Failover and Failback Capabilities

Failover

When an application or cluster resource fails, WSFC detects the failure and attempts to restart the resource. If the restart fails, WSFC takes the application offline, moves the application and its resources to another node, and restarts the application on the other node.

See "Setting Resource Properties" on page 7 for more information.

Cluster resources are placed in a group so that WSFC can move the resources as a combined unit, ensuring that the failover and/or failback procedures transfer all resources.

After failover, Cluster Administrator resets the following recovery policies:

- Application dependencies
- Application restart on the same node
- Workload rebalancing (or failback) when a failed node is repaired and brought back online

Failback

Failback returns the resources back to their original node. When the system administrator repairs and restarts the failed node, WSFC takes the running application and its resources offline, moves them from the Failover Cluster node to the original node, and then restarts the application.

You can configure failback to occur immediately, at any given time, or not at all. To minimize the delay until the resources come back online, configure the failback time during off-peak hours.

Modifying Your Failover Policy

Use the following guidelines when you modify your failover policy:

- Define how Failover Clustering detects and responds to group resource failures.
- Establish dependency relationships between the resources to control the order in which the resources are taken offline.
- Specify time-out, failover threshold, and failover period for your cluster resources.

See "Resource Failure" on page 10 for more information.

- Specify a **Possible Owner List** in Failover Clustering MMC for cluster resources. The Possible Owner List for a resource controls which nodes are allowed to host the resource.

Maintaining Your Cluster

This section provides instructions to perform multiple maintenance tasks like adding, configuring, and removing cluster components in your Dell™ Failover Cluster.

Adding a storage to a Failover Cluster Node

Failover Clustering groups all available disks on the shared storage into a group named **Available Storage** group. You can also add storage to an existing Failover Cluster. To add storage to an existing Failover Cluster:

- 1 Open the Failover Cluster Management console and connect to your cluster.
- 2 Right-click on **Storage** and click **Add a Disk**. Any disks that are visible from all the cluster nodes and can be clustered is listed.
- 3 Select the disks you want to add and then click **OK**.

This adds the disk to the Available Storage group of your Failover Cluster.

Configuring Network Settings of a Failover Cluster Node

The cluster nodes should have at least two networks, one for private network traffic and one for public network traffic. Failover Clustering allows you to specify if a network will be used by the cluster, either by the nodes only or by nodes and clients, or not used by the cluster at all. It is a generally accepted practice to enable your private network for use by the cluster nodes only and your public network to be used by cluster nodes and clients connected to the cluster.

To configure the network settings of a Failover Cluster:

- 1 Open the **Failover Cluster Management** console and connect to your cluster.
- 2 Expand **Networks**, right-click the network that you want to configure, and click **Modify**.

- 3 Configure the networks:
 - For your Private network, select **Allow the cluster to use this network only**.
 - For your Public network, select both **Allow the cluster to use this network** and **Allow clients to connect through this network**.
 - For any other network such as iSCSI network that might be configured, select **Do not allow the cluster to use this network**.

Maintaining a Clustered Service or Application

You can take any of your clustered service or applications offline to perform maintenance or diagnostics. The cluster service ensures that all the dependencies are met before making a clustered service or application **Online** or **Offline**.

To bring a Clustered service or application Online or Offline:

- 1 To open the Failover Cluster Management console, click **Start**→**Administrative Tools**.
- 2 Under **Services and Applications**, navigate to the service or application you want to take online or offline.
- 3 Right-click on the service or application and select from the following option:
 - **Bring this service or application online**
 - **Take this service or application offline**



NOTE: To monitor the events for the service or application, right-click on a service or application, select the **Show the critical events for this application**. A list of critical errors that have occurred on this service or application is displayed.

Starting or Stopping Cluster Service on Cluster Nodes

Failover Clustering allows you to stop and restart the Cluster Service on a node for troubleshooting or maintenance operations on the node. All applications or services hosted on that node fails to the other node when you stop the cluster service.

To stop or restart the cluster service on a node:

- 1 Right-click the node that you want to stop or restart in the **Failover Cluster Management** console.
- 2 Click **More Actions** and select from either of the following options that are displayed:
 - Stop Cluster Service
 - Start Cluster Service

Running chkdsk on a Clustered Disk in Maintenance Mode

Failover Clustering allows you to put a disk in **Maintenance** mode without taking the disk offline. The health monitoring on the disk turns off for a period of time while maintenance is carried out on the disk. You can then run **chkdsk** utility on the disk.

To put a disk in Maintenance mode:

- 1 To open the **Failover Cluster Management** console, click **Start**→**Administrative Tools**.
- 2 Right-click the disk that you want to put on maintenance mode.
- 3 Click **More Actions** and select **Turn On Maintenance Mode for this disk** option.
- 4 Ensure that the status of the disk is listed as **Online (Maintenance)** in the MMC.
- 5 Run **chkdsk** utility on the disk.

Displaying Event Logs for a Failover Cluster

All cluster events can be viewed using the Failover Cluster Management console for any node in your Failover Cluster. To display the event logs:

- 1 To open the **Failover Cluster Management** console, click **Start**→**Administrative Tools**.
- 2 In the console tree, right-click **Cluster Events**, and then click **Query**.

- 3** In the **Cluster Events Filter** dialog window, select the criteria for the events that you want to display and click **OK**.
- 4** To view an event, click on the event, and see the details in the **Event Details** screen.

If you want the cluster logs to be displayed in textual format, then run the following command in the command prompt of each node:

```
cluster log /g.
```

You must be logged in as an Administrator to run the command.

Upgrading to a Cluster Configuration

This section provides instructions to perform upgrading a cluster configuration in your Dell™ Failover Cluster.

Before You Begin

Before you upgrade your non-clustered system to a cluster solution:

- Back up your data.
- Verify that your hardware and storage systems meet the minimum system requirements for a cluster as described in "System Requirements" on page 10.
- Verify that your hardware and storage systems are installed and configured as explained in the following sections:
 - **Cabling Your Cluster Hardware** section of the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array
 - "Preparing Your Systems for Clustering" on page 15
 - "Installing Your Cluster Management Software" on page 5

Supported Cluster Configurations

Dell certifies and supports only solutions that are configured with the Dell products described in this guide. For more information on the corresponding supported adapters and driver versions, see *Dell Cluster Configuration Support Matrices* located on the Dell High Availability website at www.dell.com/ha.

Completing the Upgrade

After installing the required hardware and network adapter upgrades, set up and cable the system hardware.



NOTE: You may need to reconfigure your switch or storage groups so that both nodes in the cluster can access their logical unit numbers (LUNs).

The final phase for upgrading to a cluster solution is to install and configure Microsoft® Windows Server® 2008 with WSFC.

Troubleshooting

This appendix provides troubleshooting information for your cluster configuration.

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. For more information, see the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals .
	One of the cables is faulty.	Replace the faulty cable.
	You are using iSCSI storage array, the challenge handshake authentication protocol (CHAP) password entered is wrong.	Enter correct user-name and password for CHAP, if used.
You are using a Dell™ PowerVault™ MD3000 or PowerVault MD3000i storage array and the Host Group or Host-to Virtual Disk Mappings is not correctly created.	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. 	

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
The nodes cannot access the storage system, or the cluster software is not functioning with the storage system.	You are using a Dell/EMC storage array and Access control is not enabled correctly.	Verify the following: <ul style="list-style-type: none">• EMC® Access Logix™ software is enabled on the storage system.• All logical unit numbers (LUNs) and hosts are assigned to the proper storage groups.
	You are using a Fibre Channel storage array in a Storage Area Network (SAN), and one or more zones are not configured correctly.	Verify the following: <ul style="list-style-type: none">• Each zone contains only one initiator (Fibre Channel daughter card).• Each zone contains the correct initiator and the correct storage port(s).
	You are using a Fibre Channel storage array and the length of the interface cables exceeds the maximum allowable length.	Ensure that the fibre optic cables do not exceed 300 m (for multimode) or 10 km (for single mode switch-to-switch connections only).

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® Server 2008® Failover Cluster (WSFC) and the clustered applications or services.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Attempts to connect to a cluster using Cluster Administrator fail.	<ul style="list-style-type: none">• 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. <p>The cluster network name is not responding on the network because the Internet Connection Firewall is enabled on one or more nodes.</p>	<p>Verify that the 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.</p> <p>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.</p> <p>Configure the Internet Connection Firewall to allow communications that are required by WSFC and the clustered applications or services.</p>
You are prompted to configure one network instead of two during WSFC installation.	The TCP/IP configuration is incorrect.	The node-to-node network and public network must be assigned static IP addresses on different subnets. See "Assigning Static IP Addresses to Cluster Resources and Components" on page 22 for information about assigning the network IPs.
	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.

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• For Fibre Channel storage arrays, check all zone configurations• Check the Access Control settings on the attached storage systems. Verify that the node in question is a member of the correct Storage Group or Host Group.• 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 WSFC and the clustered applications or services.
The disks on the shared cluster storage appear offline in Windows Disk Administration	This situation is normal. By default, the shared disks are offline when the node detects them the first time.	Right-click on the disk and click Online .
The Validate Network Communication test fails on the iSCSI networks	This is a known issue which occurs when you run Cluster Validation on a cluster configuration that is using iSCSI array as Direct attached storage in Windows Server 2008.	See Microsoft Knowledge Base article KB951434 at the Microsoft Support website at support.microsoft.com for more information.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
Cluster Services may not operate correctly on a cluster running Windows Server 2008 when the Internet Firewall enabled.	The Windows Internet Connection Firewall is enabled, which may conflict with Cluster Services.	Perform the following steps: 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 .
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 WSFC and the clustered applications or services.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
You are using a PowerVault MD3000 or PowerVault MD3000i storage array and Virtual Disks failover 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: <ul style="list-style-type: none">• For PowerVault MD3000 storage, merge the Cluster.reg file located in the \utility directory of the Dell PowerVault MD3000 Resource Media into the registry of each node. This file changes the PowerVault MD3000 from <i>Stand Alone</i> to <i>Cluster</i> mode.• For PowerVault MD3000i, merge the Cluster.reg file located in the windows\utility directory of the Dell PowerVault MD3000i resource media into the registry of each node. This file changes the PowerVault MD3000i from <i>Stand Alone</i> to <i>Cluster</i> mode.
You are using a PowerVault MD3000 or PowerVault MD3000i storage array and 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
<p>You are using a PowerVault MD3000 or PowerVault MD3000i storage array and one of the following occurs:</p> <ul style="list-style-type: none">• Unable to assign the drive letter to the snapshot virtual disk.• Unable to access the snapshot virtual disk.• 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.	<p>The snapshot virtual disk has been erroneously mapped to the node that does not own the source disk.</p>	<p>Unmap the snapshot virtual disk from the node not owning the source disk, then assign it to the node that owns the source disk. For more information, see "Using Advanced (Premium) PowerVault Modular Disk Storage Manager Features" section of the Dell Failover Cluster Hardware Installation and Troubleshooting Guide for the specific storage array on the Dell Support website at support.dell.com/manuals.</p>
<p>You are using a PowerVault MD3000 or PowerVault MD3000i storage array in a non-redundant configuration. The Recovery Guru in the Modular Disk Storage Manager Client reports that the virtual disks are not on the preferred controller, and the enclosure status LED is blinking amber.</p>	<p>The NVSRAM for the non-redundant configuration has not been loaded.</p>	<p>For PowerVault MD3000 storage array, load the correct NVSRAM for the non-redundant configuration.</p>

Index

A

active/active
about, 45

C

chkdsk/f
running, 57

cluster
cluster objects, 39
forming a new cluster, 41
joining an existing cluster, 41

Cluster Administrator
about, 37

cluster configurations
active/active, 45
active/passive, 45
supported configurations, 59

cluster nodes
about, 40
states and definitions, 40

cluster objects
about, 39

cluster resources
resource failure, 44
setting resource properties, 41

cluster storage
requirements, 12

D

domain model
selecting, 19

drivers
installing and configuring
Emulex, 27

E

Emulex HBAs
installing and configuring, 27
installing and configuring
drivers, 27

F

failback
about, 53

failover
modifying failover policy, 53
policies, 46

failover configurations
for Windows Server 2003,
Enterprise Edition, 46

- failover policies, 46
 - failover pair, 49
 - failover ring, 52
 - for Windows Server 2003,
Enterprise Edition, 46
 - multiway failover, 50
 - N+I failover, 47

G

- group affinity
 - about, 49
 - configuring, 49

H

- HBA drivers
 - installing and configuring, 27
- host bus adapter
 - configuring the Fibre Channel
HBA, 27

I

- IP address
 - assigning to cluster resources and
components, 22
 - example configuration, 23

M

- Microsoft Cluster Administrator
 - about, 37
- MSCS
 - installing and configuring, 29
- multiway failover, 50

N

- N+I failover
 - configuring group affinity, 47
- network adapters
 - using dual-port for the private
network, 26
- network failure
 - preventing, 39
- network interfaces, 40
- networking
 - configuring Windows, 21

O

- operating system
 - installing, 20
 - upgrading, 60
 - Windows Server 2003, Enterprise
Edition
 - installing, 17

P

- period values
 - adjusting, 44
- private network
 - configuring IP addresses, 23
 - creating separate subnets, 25
 - using dual-port network adapters, 26
- public network
 - creating separate subnets, 25

Q

- quorum resource
 - running chkdisk, 57

R

- resource
 - creating, 43
- resource dependencies, 42

S

- subnets
 - creating, 25

T

- threshold
 - adjusting, 44
- troubleshooting
 - connecting to a cluster, 64
 - shared storage subsystem, 61-62

U

- upgrading
 - operating system, 60
- upgrading to a cluster solution
 - before you begin, 59
 - completing the upgrade, 60

W

- warranty, 12
- Windows Server 2003,
Enterprise Edition
 - cluster configurations, 47, 49-50, 52

