

Dell™ PowerEdge™ Cluster
FE550W/FE650W Systems

Installation and Troubleshooting Guide



Notes, Notices, and Cautions



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



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



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

Information in this document is subject to change without notice.

© 2007 Dell Inc. All rights reserved.

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

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge*, and *OpenManage* are trademarks of Dell Inc.; *VERITAS* is a registered trademark of VERITAS Software Corporation; *Microsoft*, *Windows*, *Windows NT*, and *Windows Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries; *EMC*, *CLARiON*, *PowerPath*, and *Navisphere* are registered 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	9
	Overview	9
	Virtual Servers and Resource Groups	9
	Quorum Resource	10
	Shared Storage Systems	10
	Cluster Solution	11
	Operating System	11
	Storage System	12
	Storage Management Software	14
	Fibre Channel Overview	16
	Cluster Configurations	17
	Direct-Attached Cluster	17
	SAN-Attached Cluster	18
	Configuring Active and Passive Cluster Nodes	19
	Failover Policies	21
	Windows 2000 Advanced Server Cluster Configurations	21
	Windows Server 2003 Cluster Configurations	22
	System Requirements	27
	Cluster Nodes	28
	Cluster Storage	29
	Other Documents You May Need	30

2	Cabling Your Cluster Hardware	33
	Cabling the Mouse, Keyboard, and Monitor	33
	Cabling the Power Supplies	33
	Fibre Channel Cable Connectors	35
	Cabling Your Cluster for Public and Private Networks	36
	Cabling to the Public Network	37
	Cabling the Private Network	37
	NIC Teaming	38
3	Preparing Your Systems for Clustering	39
	Before You Begin	39
	Installation Overview	39
	Selecting a Domain Model	41
	Configuring the Nodes as Domain Controllers	41
	Configuring Internal Drives in the Cluster Nodes	42
	Installing and Configuring the Windows Operating System	42
	Configuring Windows Networking	43
	Assigning Static IP Addresses to Cluster Resources and Components	43
	Configuring IP Addresses for the Private Network	44
	Verifying Communications Between Nodes	47

Configuring the Virtual Disks for Cluster Nodes	47
Assigning the Virtual Disks to a Cluster Node	48
Using the Windows Dynamic Disks and Volumes	48
Naming and Formatting Drives on the Shared Storage System	49
Configuring Hard Drive Letters When Using Multiple Shared Storage Systems	52
Installing and Configuring MSCS	54
Verifying Cluster Readiness	54
Installing Applications in the Cluster Group	55
Installing the Quorum Resource	55
Installing and Configuring Windows 2000 Cluster Service (MSCS)	55
Configuring Cluster Networks Running Windows Server 2003	56
Configuring Windows Server 2003 Cluster Service (MSCS)	57
Verifying MSCS Operation	59
Verifying Cluster Functionality	59
Verifying Cluster Resource Availability	60
Troubleshooting Failed Resources	60
4 Upgrading to a Cluster Configuration . . .	61
Before You Begin	61
Supported Cluster Configurations	61
Completing the Upgrade	62
Upgrading Your Operating System	62
Performing a Rolling Upgrade	63

5	Installing Your Cluster in a Direct-Attached Environment	67
	Before You Begin	67
	Connecting the Storage Systems to Your Cluster	67
	Direct-Attached Cluster	67
	Cabling One Cluster to a Dell EMC Storage System	69
	Cabling Two Clusters to a Dell EMC Storage System	71
	Installing the Windows Operating System and Cluster Management Software	71
6	Installing Your Cluster in a SAN-Attached Environment	73
	SAN Overview	73
	Connecting the Storage Systems to Your Cluster	74
	SAN-Attached Cluster	74
	Cabling One SAN-Attached Cluster to an AX100/150 Storage System	76
	Cabling Multiple SAN-Attached Clusters to a Dell EMC Storage System	79
	Connecting a PowerEdge Cluster to Multiple Storage Systems	80
	Implementing Port Zoning on a Fibre Channel Switched Fabric	81
	Using Port Zoning	82
	Using Worldwide Name zoning	82
	Single Initiator Zoning	83

	SAN-Attached Cluster Configurations	83
	Installing the Windows Operating System and Cluster Management Software	85
7	Maintaining Your Cluster	87
	Adding a Network Adapter to a Cluster Node	87
	Changing the IP Address of a Cluster Node on the Same IP Subnet	88
	Uninstalling MSCS From Clusters Running Windows 2000 Advanced Server	89
	Removing Nodes From Clusters Running Windows Server 2003	89
	Running chkdsk /f on a Quorum Resource	90
	Recovering From a Corrupt Quorum Disk	90
	Changing the MSCS Account Password in Windows Server 2003	91
	Reformatting a Cluster Disk	92
8	Using Microsoft Cluster Service (MSCS)	95
	Cluster Objects	95
	Cluster Networks	95
	Preventing Network Failure	95
	Node-to-Node Communication	96
	Network Interfaces	96

Cluster Nodes	96
Forming a New Cluster	97
Joining an Existing Cluster	97
Cluster Resources	97
Setting Resource Properties	98
Resource Dependencies	98
Setting Advanced Resource Properties	99
Resource Parameters	99
Quorum Resource	100
Resource Failure	101
Resource Dependencies	102
Creating a New Resource	102
Deleting a Resource	103
File Share Resource Type	104
Failover and Failback Capabilities	104
Failover	104
Failback	105
A Troubleshooting	107
B Cluster Data Form	111
Index	113

Introduction

This document provides information for installing and managing your cluster solution. It is intended for experienced IT professionals who need to configure the cluster solution, and for trained service technicians who perform upgrades and maintenance procedures. This document also addresses readers who are new to clustering.

Overview

Clustering uses specific hardware and software to join multiple systems together to function as a single system and provide an automatic failover solution. If one of the clustered systems (also known as a cluster node) fails, resources running on the failed system are moved (or failed over) to one or more systems in the cluster by the Microsoft® Cluster Service (MSCS) software. MSCS is the failover software component in specific versions of the Microsoft Windows® operating system.

When the failed system is repaired and brought back online, resources automatically transfer back (or fail back) to the repaired system or remain on the failover system, depending on how MSCS is configured. For more information, see "Failover and Failback Capabilities" in "Using Microsoft Cluster Service (MSCS)".

Virtual Servers and Resource Groups

In a cluster environment, users do not access a physical server; they access a virtual server, which is managed by MSCS. Each virtual server has its own IP address, name, and hard drive(s) in the shared storage system. MSCS manages the virtual server as a *resource group*, which contains the cluster resources. Ownership of virtual servers and resource groups is transparent to users. For more information about resource groups, see "Cluster Resources" in "Using Microsoft Cluster Service (MSCS)".

When MSCS detects a failed node or failed application, it moves the entire resource group to another node and remaps the virtual server to the new network connection. Users of an application in the virtual server experience only a momentary delay in accessing resources while MSCS re-establishes a network connection to the virtual server and restarts the application.

Quorum Resource

A single shared disk, which is designated as the quorum resource, maintains the configuration data (including all the changes that have been applied to a cluster database) necessary for recovery when a node fails.

The quorum resource can be any resource with the following attributes:

- Enables a single node to gain and defend its physical control of the quorum resource
- Provides physical storage that is accessible by any node in the cluster
- Uses the Windows NT[®] file system (NTFS)

For more information, see "Quorum Resource" in "Using Microsoft Cluster Service (MSCS)" and the MSCS online documentation.



NOTE: Dell™ PowerEdge™ FE550W and FE650W clusters do not support the Majority Node Set Quorum resource type.

Shared Storage Systems


Cluster nodes can share access to external storage systems; however, only one of the nodes can own any virtual disk in the external storage system at any time. MSCS controls which node has access to each virtual disk in the shared storage system.

Each storage system in the cluster is centrally managed by one host system (also called a *management station*) running EMC[®] Navisphere™ Express—a centralized storage management application used to configure Dell | EMC storage systems.

You can use Navisphere Express to perform tasks such as creating disk pools, binding the virtual disks, and downloading the firmware. Additionally, you can use Snapshot Management to capture point-in-time images of a virtual disk for backups or testing without affecting the contents of the source virtual disk.

Cluster Solution


Your cluster solution implements two-node to eight-node clustering and provides the following features.

 **NOTE:** Some SAN-attached cluster configurations can support only six nodes due to limited ports on the Fibre Channel switches. See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for the supported Fibre Channel switches.

- 2 Gb/sec Fibre Channel 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

Operating System

Table 1-1 provides an overview of the supported operating systems. See your operating system documentation for a complete list of features.

 **NOTE:** Some of the core services are common to both operating systems.


 **NOTE:** Throughout this guide, "Windows Server® 2003" refers to either Windows Server 2003, Enterprise Edition or Windows Server 2003 R2, Enterprise Edition. See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for additional information about the supported operating system versions and service packs.


Table 1-1. Windows Operating System Features


Windows 2000 Advanced Server	Windows Server 2003
Supports two-node clusters	Supports up to two nodes with: <ul style="list-style-type: none">• Single storage processor (SP)• Dual SP in direct-attached configuration only or Supports up to eight SAN-attached nodes per cluster using the Dell EMC AX100/150 (dual SP) storage system
Supports up to 8 GB of RAM per node	Supports up to 32 GB of RAM per node

Table 1-1. Windows Operating System Features (continued)

Windows 2000 Advanced Server	Windows Server 2003
Cluster configuration and management using control panel utilities	Cluster configuration and management using Configure Your Server and Manage Your Server wizards
	Metadirectory Services

 **NOTE:** Throughout this document, Dell | EMC AX100/150 storage system refers to both Dell | EMC AX100 and AX150-series storage systems, unless explicitly mentioned.

 **NOTE:** Running multiple operating systems in a cluster is not supported.

 **NOTE:** MSCS and Network Load Balancing (NLB) features cannot coexist on the same node, but can be used together in a multitiered cluster. For more information, see the Dell PowerEdge Clusters website at www.dell.com/clusters or the Microsoft Support website at support.microsoft.com.

Storage System

Your cluster supports AX100/150-series storage systems in a direct-attached or SAN-attached environment. These storage systems work together with the following hardware components:

- Storage processor (SP)
- Uninterrupted power supply (UPS)
- Host bus adapter (HBA)
- Fibre Channel switches (for SAN-attached configurations)

Processor Enclosure

The processor enclosure is configured with SP that control the disk pools in the storage system and provides various storage functionalities, including snapshots. Supported processor enclosures vary in performance and configuration.

The AX100/150-series processor enclosures vary in configuration and application. Table 1-2 provides a list of features for each enclosure.

Table 1-2. Processor Enclosure Features

AX100SC/150SC Processor Enclosure	AX100/150 Processor Enclosure
One SP	Two SPs
Supports direct-attached clusters only	Supports direct-attached and SAN-attached clusters
One power supply (AX100SC) One or two power supplies (AX150SC)	Two power supplies
One data path	Redundant data paths
Write cache protection using an internal cache card	Cache Vaulting using an external UPS
12 hot-swappable SATA drives	12 hot-swappable SATA drives
Supports up to six disk pools	Supports up to six disk pools
Supports RAID 1, RAID 5, RAID 1/0	Supports RAID 1, RAID 5, RAID 1/0

UPS

UPS is required for the dual storage systems (AX100/150 storage systems). It supplies battery power to SP-A running long enough to write the cached data to the disk during a complete power outage.



NOTE: UPS does not support the AX100SC/150SC storage systems.

HBA

The HBA is an I/O adapter that connects the server's expansion bus to the storage components. Cluster solutions that are configured with the AX100SC/150SC (single SP) storage system require one HBA in each server to transfer data from the server to the storage system. Cluster solutions that are configured with the AX100/150 (dual SP) storage system require two HBAs in each server to provide redundant paths and load balance the I/O data transfer from the server to the storage system.

See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for a list of supported HBAs in your cluster configuration.

Fibre Channel Switch

The Fibre Channel switch functions as a director, mapping requests and responses between the interconnected devices, and provides a dedicated interconnection between the server and storage system. See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for a list of supported switches in your cluster configuration.

Storage Management Software

This section provides information about the software your cluster uses to manage the communications between the nodes and the storage systems in a SAN environment.

Navisphere Storage System Initialization Utility

The Navisphere Storage System Initialization Utility provides a user interface to initialize your AX100/150-series storage systems. Using the utility, you can configure the IP address, subnet mask, and default gateway address for the storage system's SPs, and assign user names and password for storage system access.



NOTE: For best practice, install the utility on at least one node in the cluster.

Navisphere Server Utility

The Navisphere Server Utility registers the cluster node HBAs with the AX100/150-series storage systems, allowing the nodes to access the cluster storage data. The tool is also used for cluster node maintenance procedures, such as:

- Updating the cluster node host name and/or IP address on the storage array
- Updating file system information
- Adding, removing, or replacing an HBA
- Starting and stopping a snapshot



NOTE: The Navisphere Server Utility must be installed on all cluster nodes.

PowerPath

EMC PowerPath™ detects and re-establishes a failed connection to a processor enclosure by automatically rerouting I/O through an alternate path. PowerPath also provides data load balancing from the server to the storage system.



NOTE: If the AX100SC/150SC storage system is configured with one SP, PowerPath can only detect and display the failed connection.

Navisphere Express

Navisphere Express provides centralized storage management and configuration, allowing you to configure and manage the disks and components in one shared storage system. Navisphere Express is included with the core software (FLARE) on the AX100/150-series storage systems.

Snapshot Management

Snapshot Management captures images of a virtual disk and retains the image independently of subsequent changes to the files. The images can be used to share virtual disks with another system without affecting the contents of the source virtual disk.

Snapshot Management creates copies of either virtual disks or snapshots. Snapshots are virtual copies that create an image of the source virtual disk at the time the snapshot was created. This snapshot is retained independently of subsequent changes to the source virtual disk. You can use snapshots to facilitate backups or to allow multiple hosts to access data without affecting the contents of the source virtual disk.



NOTICE: To avoid data corruption, do not access a snapshot from the same node as its source.

EZ Switch Setup Utility

The Brocade EZ Switch Setup Utility is a wizard-driven utility that configures your Brocade Fibre Channel switches to their original settings and confirms that all cluster nodes and storage systems are attached and logged into each switch. Additionally, the utility restores your switch's zoning configuration to its original settings.

Fibre Channel Overview

This section provides a brief overview of Fibre Channel and the elements of a Fibre Channel network.

Fibre Channel Protocol

Fibre Channel provides high-speed data transfer between the nodes and storage systems, allows multiple server systems to share one or more storage systems, and provides long-distance connectivity and high bandwidth.

Implementing Fibre Channel technology in a cluster provides the following advantages:

- **Availability** — Fibre Channel components use redundant connections, providing multiple data paths and greater availability for clients.
- **Connectivity** — Fibre Channel allows more device connections than SCSI. Because Fibre Channel devices are hot-pluggable, you can add or remove devices from the nodes without bringing down the cluster.

Fibre Channel Switch Fabric

A Fibre Channel switch fabric consists of one or more switches that provide a connection from one device (sender) to another device or switch (receiver) on the network. If the data is sent to another switch, the process repeats itself until a connection is established between the sender and the receiver.

SAN

In addition to direct-attached connections, nodes can also connect to storage systems through a SAN—a high-performance network storage solution. A SAN bypasses traditional network bottlenecks, providing a high-speed, highly available data consolidation solution.



NOTE: Your SAN may require additional hardware and software components that are not listed in this document or the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide*. For information about SAN-compliant hardware and software components, see the EMC Support Matrix located in the EMC Technical Library at www.emc.com/dell-ax100support for the AX100 storage system and www.emc.com/dell-ax150support for the AX150 storage system.

Zones

Fibre Channel fabrics enable you to set up barriers between different devices and operating environments. These barriers create logical fabric subsets (also known as *zones*) that divide a fabric into smaller groups or components, regardless of their proximity to one another.

Cluster Configurations

Direct-Attached Cluster

In a two-node direct-attached cluster, both cluster nodes are directly attached to a single storage system. In this configuration, the RAID controllers (or SPs) on the storage systems are connected by cables directly to one Fibre Channel HBA in each node.

Figure 1-1 shows a basic direct-attached, single-cluster configuration that is configured with the AX100SC/150SC (one SP) storage system.

Figure 1-1. Direct-Attached, Single-Cluster Configuration Using the AX100SC/150SC Storage System

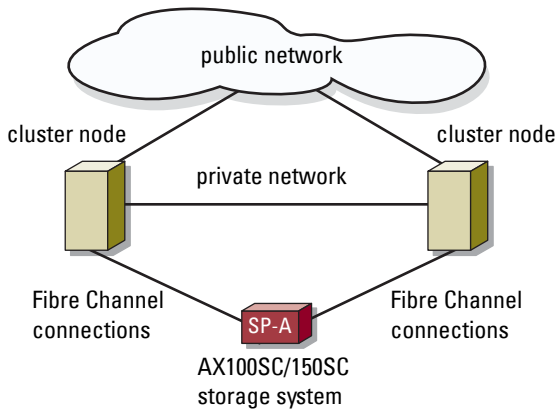



Figure 1-2 shows a basic direct-attached, single-cluster configuration that is configured with the AX100/150 (two SPs) storage system. As AX100/150 storage system includes two SPs, the cluster nodes are configured with redundant paths to the storage system, thereby eliminating a single point of failure.

Figure 1-2. Direct-Attached, Single-Cluster Configuration Using the AX100/150 Storage System



PowerPath Limitations in a Direct-Attached Cluster

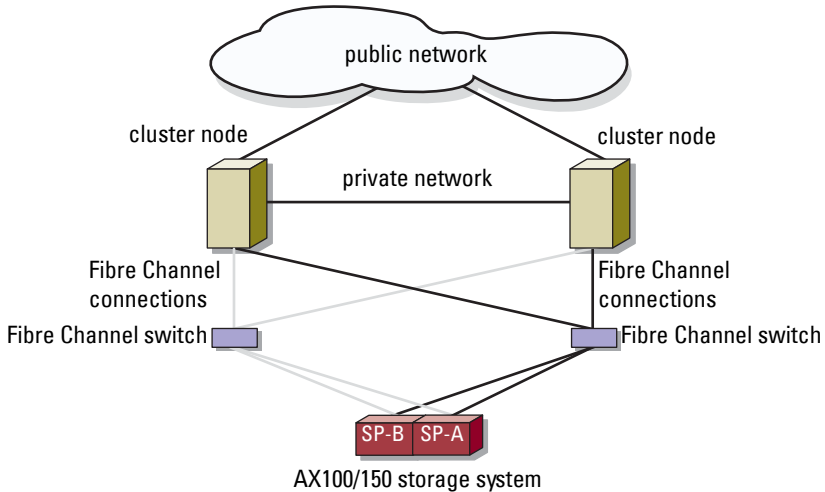
PowerPath provides failover capabilities and multiple path detection, as well as dynamic load balancing between multiple ports on the same SP. If your cluster solution is configured with the AX100SC/150SC storage system, PowerPath only detects and displays a path failure. Because the storage system is configured with one SP and no redundant paths, PowerPath cannot reroute the I/O traffic to another path.

SAN-Attached Cluster

In a SAN-attached cluster, all of the nodes are attached to a single storage system or to multiple storage systems through a SAN using redundant switch fabrics. SAN-attached clusters are superior to direct-attached clusters in configuration flexibility, expandability, and performance.


Figure 1-3 shows a SAN-attached cluster.

Figure 1-3. SAN-Attached Cluster



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 (active^x) and active/passive (active^x/passive^x) configurations. The variable *x* indicates the number of nodes that are active or passive.

Cluster solutions running on the Windows Server 2003 operating system can support up to eight nodes in multiple configurations as shown in Table 1-4. Cluster solutions running Windows 2000 Advanced Server can support only two nodes and are limited to active² and active¹/passive¹ configurations because this solution supports only two nodes.

An active/active (active^x) 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. The 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 (active^x/passive^x) 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.

Table 1-3 provides a description of active/active configuration types.

Table 1-3. Active/Active Configuration Types

Configuration Type	Active Cluster Node(s)	Definition
Active ²	2	
Active ³	3	
Active ⁴	4	The active node(s) processes requests and provides failover for each other, depending on node resources and your configuration.
Active ⁵	5	
Active ⁶	6	NOTE: Configurations running Windows 2000 Advanced Server are limited to two-node, active/active (active ²) configurations.
Active ⁷	7	
Active ⁸	8	

Table 1-4 provides a description of some active/passive configuration types.

Table 1-4. Active/Passive Configuration Types

Configuration Type	Active Cluster Node(s)	Passive Cluster Node(s)	Description
Active ¹ /Passive ¹	1	1	
Active ² /Passive ¹	2	1	
Active ² /Passive ²	2	2	
Active ³ /Passive ¹	3	1	<p>The active node(s) processes requests while the passive node waits for the active node to fail.</p> <p>NOTE: Configurations running Windows 2000 Advanced Server are limited to two-node, active/passive (active¹/passive¹) configurations.</p>
Active ³ /Passive ²	3	2	
Active ⁴ /Passive ¹	4	1	
Active ⁴ /Passive ²	4	2	
Active ⁵ /Passive ¹	5	1	
Active ⁵ /Passive ²	5	2	
Active ⁶ /Passive ¹	6	1	
Active ⁶ /Passive ²	6	2	
Active ⁷ /Passive ¹	7	1	

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 2000 Advanced Server Cluster Configurations

Configurations running Windows 2000 Advanced Server include only two nodes. Therefore, failover to one node is the only option.

Windows Server 2003 Cluster Configurations

Cluster configurations running Windows Server 2003 provide the following failover policies:

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

Table 1-5 provides an overview of the failover policies implemented with Windows 2000 Advanced Server and Windows Server 2003. For more information, see the sections that follow this table.

Table 1-5. Windows Server 2003 Failover Policies

Failover Policy	Description	Advantage	Disadvantage(s)
$N + I$	One or more nodes provides backup for multiple servers.	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 fail over 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 1-4 shows an example of a 4 + 2 (N + I) failover configuration with four active nodes and two passive nodes. Table 1-6 provides an N + I failover matrix for Figure 1-4.

Figure 1-4. Example of an N + I Failover Configuration for a Six-Node Cluster




Table 1-6. Example of an N + I Failover Configuration for a Six-Node Cluster

Cluster Resource Group	Primary Node	AntiAffinityClassNamesValue
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 2003, 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 2003 to assign a public property (or attribute) to a dependency between groups to ensure that they fail over to similar or separate nodes. This property is called *group affinity*.

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

For example, in Table 1-6, the `AntiAffinityClassNamesValues` 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 will fail over to the next backup node (node 5). If node 2 fails, group B will skip node 5 and instead fail over to node 6 because the `AntiAffinityClassNamesValues` string value (`AString`) identifies group A and group B as conflicting groups.

To set the public property for the cluster groups shown in Table 1-6:

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

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

- 3 Repeat step 2 for the remaining cluster groups.

Use the "Cluster Data Form" in this document to specify group affinity in your $N + I$ cluster configuration.

Failover Pair

Failover pair is a policy in which each application can fail over 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, active/active configurations running Windows 2000 Advanced Server, and you want to migrate these applications to Windows Server 2003, 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 1-5 shows an example of a failover pair configuration. Table 1-7 provides a failover configuration for the cluster shown in Figure 1-5.

Figure 1-5. Example of a Failover Pair Configuration




Table 1-7. Example of a Failover Pair Configuration for a Four-Node Cluster

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

Multway Failover

Multway 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 1-6 shows an example of four-node multway failover configuration.

Table 1-8 shows a four-node multway failover configuration for the cluster shown in Figure 1-6. For each resource group, the failover order in the **Preferred Owners** list in Cluster Administrator outlines the order that you want that resource group to fail over. In this example, node 1 owns applications A, B, and C. If node 1 fails, applications A, B, and C fail over 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. For more information, see "Using Microsoft Cluster Service (MSCS)".

Figure 1-6. Example of a Four-Node Multiway Failover Configuration




Table 1-8. 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 fail over 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 1-7 shows an example of a failover ring configuration.

Figure 1-7. Example of a Four-Node Failover Ring Configuration



System Requirements

Your cluster requires the following components:

- Servers (nodes)
- Storage
- Interconnects (private network)
- Client network connections (public network)
- Operating system and storage management software

Cluster Nodes

Table 1-9 lists the hardware requirements for the cluster nodes.

Table 1-9. Cluster Node Requirements

Component	Minimum Requirement
Cluster nodes	<ul style="list-style-type: none">• Two supported PowerEdge systems running Windows 2000 Advanced Server connected to an AX100SC/150SC storage system.• Two supported PowerEdge systems running Windows Server 2003 connected to the AX100SC/150SC storage system.• Two to eight PowerEdge systems running Windows Server 2003 connected to the AX100/150 storage system.
Processors	At least two processors per node.
RAM	At least 256 MB of RAM installed per node.
HBAs	One Fibre Channel HBA per node for the AX100SC/150SC storage system. OR Two Fibre Channel HBAs per node for the AX100/150 storage system. Place the HBAs on separate PCI buses to improve availability and performance. See the <i>Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide</i> for information about supported systems, HBAs, and PCI slot configuration guidelines.
NICs	At least two NICs: one NIC for the public network and another NIC for the private network. NOTE: For best practice, ensure that the NICs on each public network are identical, and that the NICs on each private network are identical.

Table 1-9. Cluster Node Requirements (continued)

Component	Minimum Requirement
RAID controller (optional)	One controller connected to at least two internal hard drives for each node. Use any supported RAID controller or SCSI adapter. Two hard drives are required for mirroring (RAID 1) and at least three hard drives are required for disk striping with parity (RAID 5). NOTE: For best practice, use hardware-based RAID or software-based disk-fault tolerance for the internal drives.

Cluster Storage

Table 1-10 lists supported storage systems and the configuration requirements for the cluster nodes and stand-alone systems connected to the storage systems. Table 1-11 lists hardware requirements for the cluster storage.


Table 1-10. Cluster Storage Requirements

Hardware Components	Minimum Requirement
Supported storage systems	One AX100SC/150SC storage system or AX100/150 storage system. See Table 1-11 for specific storage system requirements.
Cluster nodes	All nodes must be attached to a single storage system or to multiple storage systems through a SAN.
Multiple clusters and stand-alone systems	Can share one or more supported storage systems using optional software that is available for your storage system. See "Storage Management Software."


Table 1-11. Cluster Storage System Requirements


Storage System	Storage Processors	Primary Storage	Secondary Storage	UPS
AX100SC/ 150SC	1	At least three internal hard drives	Not supported	Not supported
AX100/150	2	At least four internal hard drives	Not supported	One UPS

NOTE: The AX100/150 storage system must be connected to the UPS (that shipped with the storage) to enable write caching in the event of a power failure.

 **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.

Other Documents You May Need

 The *Product Information Guide* 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.

- The *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* provides information about the platforms that support your cluster configuration.
- 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 *Installation and Troubleshooting Guide* describes how to troubleshoot the system and install or replace system components.
- 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 about configuring and installing these options.
- The RAID documentation provides information for installing and configuring a RAID controller card.
- The documentation that came with your storage system.
- The EMC PowerPath documentation that came with your HBA kit(s).
- 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.

Cabling Your Cluster Hardware


 **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.

Cabling the Mouse, Keyboard, and Monitor

 **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.

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

Cabling the Power Supplies

 **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.


For best practice, use the following guidelines to protect your cluster solution from power-related failures:

- For nodes with multiple power supplies, plug each power supply into a separate AC circuit.
- Use uninterruptible power supplies (UPS).
- For some environments, consider having backup generators and power from separate electrical substations.

See your system documentation for more information about the specific power requirements for your cluster solution's components.


Figure 2-1, and Figure 2-2 illustrate the proper power cabling of cluster solution components consisting of two systems and two storage systems. To ensure redundancy, the primary power supplies of all the components are grouped onto one or two circuits and the redundant power supplies are grouped onto a different circuit.

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



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

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




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

Fibre Channel Cable Connectors

Your cluster configuration requires fiber optic connectors to connect the various Fibre Channel storage components to the cluster configuration. Your solution includes the duplex LC multimode fiber optic connector.

The duplex LC multimode fiber optic connector (see Figure 2-3) is used to connect a Fibre Channel switch to a Fibre Channel host bus adapter (HBA) on a cluster node or to a storage system. This type of connection requires fiber optic cables with duplex LC multimode fibre optic connectors.


Figure 2-3. Duplex LC Multimode Fiber Optic Connector



The LC connector consists of two individual fiber optic connectors. Each connector is indexed and must be inserted and aligned properly in the GBIC or SFP module connector.

When using duplex multimode fiber optic connectors, keep the covers on the connectors until you are ready to insert the connectors into the system. The covers protect the cables and connectors and prevent foreign particles from entering and damaging the connector.

Cabling Your Cluster for Public and Private Networks

 **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.


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

Table 2-1. Network Connections

Network Connection	Description
Public network	All connections to the client LAN. At least one public network must be configured for Mixed mode for private network failover.
Private network	A dedicated connection for sharing cluster health and status information only.

Figure 2-4 shows an example of cabling in which dedicated network adapters in each node are connected to each other (for the private network) and the remaining network adapters are connected to the public network.

Figure 2-4. Example of Network Cabling Connection



Cabling to the Public Network


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

Cabling the Private Network

The private network connection to the nodes is provided by a different network adapter in each node. This network is used for intracluster communications. Table 2-2 describes three possible private network configurations.

Table 2-2. Private Network Hardware Components and Connections

Method	Hardware Components	Connection
Network switch	Fast Ethernet or Gigabit Ethernet network adapters and switches	Connect <i>standard</i> Ethernet cables from the network adapters in the nodes to a Fast Ethernet or Gigabit Ethernet switch.
Point-to-Point Fast Ethernet (two-node clusters only)	Fast Ethernet network adapters	Connect a <i>crossover</i> Ethernet cable between the Fast Ethernet network adapters in both nodes.
Point-to-Point Gigabit Ethernet (two-node clusters only)	Copper Gigabit Ethernet network adapters	Connect a <i>standard</i> Ethernet cable between the Gigabit Ethernet network adapters in both nodes.

 **NOTE:** On certain Microsoft® Windows® 2000 Advanced Server configurations, using an Ethernet cable in a point-to-point connection can impact node-to-node communications. For more information, see the Microsoft Knowledge Base articles 239924, 242430, 254651, and 258750 located on the Microsoft Support website at support.microsoft.com. This issue has been corrected in Windows Server® 2003.

Using Dual-Port Network Adapters for the Private Network

You can configure your cluster to use the public network as a failover for private network communications. If dual-port network adapters are used, do not use both ports simultaneously to support both the public and private networks.

NIC Teaming

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

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

Preparing Your Systems for Clustering

Before You Begin

- 1 Ensure that your site can support the cluster's power requirements.
Contact your sales representative for information about your region's power requirements.



CAUTION: See your *Product Information Guide* for complete information about safety precautions, working inside the computer, and protecting against electrostatic discharge.

- 2 Ensure that the following components are installed in each Dell™ PowerEdge™ system:
 - Network adapters
 - Host bus adapters (HBAs)
 - Hard drives
 - Any additional peripheral components
- 3 Configure the storage system(s) as described in your storage system documentation.
- 4 Cable the system hardware.
See "Cabling Your Cluster Hardware" for more information.

Installation Overview

This section provides installation overview procedures for configuring a cluster running the Microsoft® Windows® 2000 Advanced Server or Windows Server® 2003 operating system.

- 1 Ensure that the cluster meets the requirements as described in "Before You Begin."

- 2 Select a domain model that is appropriate for the corporate network and operating system.

See "Selecting a Domain Model."

- 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 Microsoft Cluster Service (MSCS).

- 4 Configure the internal hard drives.

See "Configuring Internal Drives in the Cluster Nodes."

- 5 Install and configure the Windows operating system.

The Windows operating system must be installed on all of the 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."

- 6 Configure the hosts (cluster nodes) and storage as described in your storage system documentation.



NOTE: For installation and configuration of Dell | EMC AX-series storage system, see the documentation at www.emc.com/dell-ax100support for the AX100 storage system and www.emc.com/dell-ax150support for the AX150 storage system

To cable direct-attached or SAN-attached cluster nodes to the storage system, see "Installing Your Cluster in a Direct-Attached Environment" or "Installing Your Cluster in a SAN-Attached Environment".

- 7 Assign the virtual disks (configured on the storage system) to cluster nodes. See "Configuring the Virtual Disks for Cluster Nodes".
- 8 Install and configure the MSCS software.

See "Installing and Configuring MSCS."

- 9 Verify cluster functionality. Ensure that:
 - The cluster components are communicating properly.
 - MSCS is started.

See "Verifying Cluster Functionality."

- 10 Verify cluster resource availability.

Use Cluster Administrator to check the running state of each resource group. See "Verifying Cluster Resource Availability."

The following subsections provide detailed information for each step in the "Installation Overview" that is specific to the operating system.

Selecting a Domain Model

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

- All nodes are member servers in an Active Directory domain.
- All nodes are domain controllers in an Active Directory domain.
- One or more nodes are a domain controller in an Active Directory and the remaining nodes are member servers.

Configuring the Nodes as Domain Controllers

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. For more information about RAID configurations, see the RAID controller documentation.



NOTE: If you are not going to use a hardware-based RAID solution, use the Windows Disk Management tool or Dell OpenManage™ Array Manager to provide software-based redundancy.

Installing and Configuring the Windows Operating System



NOTICE: 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 "Before You Begin."
- 2 Cable the hardware.



NOTE: Do not connect the nodes to the shared storage systems yet. See "Cabling Your Cluster Hardware."

- 3 Install and configure the Windows 2000 Advanced Server operating system with the latest service pack or the Windows Server 2003, Enterprise Edition or the Windows Server 2003 R2, Enterprise Edition operating system with the latest service pack on each node.

See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for more information about the latest supported service pack in your cluster configuration.

- 4 If you are installing Windows Server 2003, go to step 5.

If you are installing Windows 2000 Advanced Server, select the option to install the Cluster Service files when you are prompted. You will configure the Cluster Service later.

- 5 Ensure that the network adapter drivers installed on each cluster node are the latest supported version.
- 6 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."

See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for information on required drivers in your cluster configuration.

Configuring Windows Networking

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

Assigning Static IP Addresses to Cluster Resources and Components

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 3-1.

Table 3-1. 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 static IP addresses are required when MSCS is configured with application programs that require IP addresses, such as file sharing.

Table 3-1. Applications and Hardware Requiring IP Address Assignments (continued)

Application/Hardware	Description
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 (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.
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>See "Cabling Your Cluster Hardware" for more information about cluster interconnect options.</p> <p>NOTE: To ensure operation during a DHCP server failure, use static IP addresses.</p> <p>NOTE: On certain Windows 2000 Advanced Server configurations, using an Ethernet cable in a point-to-point connection can impact node-to-node communications. For more information, see the Microsoft Knowledge Base articles 239924, 242430, 254651, and 258750 at support.microsoft.com. This issue has been corrected in Windows Server 2003.</p>

Configuring IP Addresses for the Private Network

Use the static IP address assignments for the network adapters used for the private network (cluster interconnect).


 **NOTE:** The IP addresses in Table 3-2 are used as examples only.

Table 3-2. 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


Table 3-2. Examples of IP Address Assignments (continued)

Usage	Cluster Node 1	Cluster Node 2
Public network subnet mask	255.255.255.0	255.255.255.0
Default gateway	192.168.1.1	192.168.1.1
WINS servers	Primary 192.168.1.11	Primary 192.168.1.11
	Secondary 192.168.1.12	Secondary 192.168.1.12
DNS servers	Primary 192.168.1.21	Primary 192.168.1.21
	Secondary 192.168.1.22	Secondary 192.168.1.22
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 you are running Windows 2000 Advanced Server or Windows Server 2003, disable NetBIOS on the private network.

If multiple cluster interconnect network adapters are connected to a network switch, ensure that all of the private network's network adapters have a unique address. You can continue the IP address scheme in Table 3-2 with 10.0.0.3, 10.0.0.4, and so on for the private network's network adapters or network adapter 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 network adapters for the cluster interconnect.

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

Creating Separate Subnets for the Public and Private Networks

The public and private network's network adapters 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.

Setting the Network Interface Binding Order for Clusters Running Windows 2000

- 1 On the Windows 2000 desktop, right-click **My Network Places**, and then click **Properties**.

The **Network and Dial-up Connections** window appears.

- 2 Click the **Advanced** menu and then click **Advanced Settings**.

The **Advanced Settings** window appears.

- 3 In the **Adapters and Bindings** tab in the **Connections** box, ensure that the **Public** connections designated for **Client access only** or **All communications** are at the top of the list.

To change the connection order:

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

Setting the Network Interface Binding Order for Clusters Running Windows Server 2003

- 1 Click the **Start** button, select **Control Panel**, and double-click **Network Connections**.
- 2 Click the **Advanced** menu, and then click **Advanced Settings**.

The **Advanced Settings** window appears.

- 3** In the **Adapters and Bindings** tab, ensure that the **Public** connection is at the top of the list and 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.

Dual-Port Network Adapters and Adapter Teams in the Private Network

Dual-port network adapters and network adapter teams 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. See "Troubleshooting" for more information.

Configuring the Virtual Disks for Cluster Nodes

Before you configure and assign the virtual disks for nodes, ensure that:

- The storage system is properly initialized.
- The cluster nodes are registered with the storage system.
- The virtual disks created in the disk pools are fully initialized.



NOTE: For best practice, have at least one virtual disk for each application. If multiple NTFS partitions are created on a single LUN or virtual disk, these partitions will not be able to fail over individually from node-to-node.

Assigning the Virtual Disks to a Cluster Node

To perform data I/O to the virtual disks, assign the virtual disks to a cluster node by performing the following steps:

- 1 Open a Web browser.
- 2 In the **Address** field, type the storage system IP address.
The **Navisphere Express** console appears.
- 3 In the **Manage** menu, click **Virtual Disks**.
- 4 In the **Manage Virtual Disks** window, select the virtual disks that you want to assign to the cluster node and click **Assign Server**.
- 5 In the **Virtual Disk - Assign Server** screen, select the cluster nodes that you want to assign to the virtual disk and click **Apply**.
- 6 Close the **Navisphere Express** window.
- 7 Verify that the PowerPath on the cluster nodes can access all paths to the virtual disks.
See your storage system documentation for more information about your PowerPath software.

Using the Windows Dynamic Disks and Volumes

The Windows 2000 Advanced Server operating system that is shipped from Microsoft does not support dynamic disks (upgraded disks) or volumes as shared storage in a cluster environment. If the shared cluster storage is configured as a dynamic disk, the Cluster Configuration wizard is not able to discover the disks, preventing the cluster and network clients from accessing the disks.

VERITAS® Volume Manager for Windows 2000 enables you to add dynamic disk features to a cluster. When you install Volume Manager on a cluster and configure Volume Manager Disk Group resources, VERITAS is the first point of support for cluster issues related to those resources.


For more information, see the Microsoft Knowledge Base Article 237853 located on the Microsoft Support site at support.microsoft.com.



NOTE: The Cluster FE550W and FE650W cluster solutions do not support dynamic disks.

Naming and Formatting Drives on the Shared Storage System


After the virtual disks are assigned to host systems, assign drive letters to the virtual disks and then format the drives as NTFS drives. Format the drives and assign volume labels from the first cluster node. When completed, the remaining nodes can access the file systems and volume labels.


 **NOTICE:** Accessing the hard drives from multiple cluster nodes may corrupt the file system.

Assigning Drive Letters and Mount Points

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

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

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

 **NOTICE:** If the disk letters are manually assigned from the remaining node(s), the shared disks are simultaneously accessible from both nodes. To ensure file system integrity and prevent possible data loss before you install the MSCS software, prevent any I/O activity to the shared drives by performing this procedure on one node at a time, and ensure that all other nodes are shut down.

The number of drive letters required by individual servers in a cluster may vary. For best practice, name the shared drives in reverse alphabetical order beginning with the letter z.

To assign drive letters, create mount points, and format the disks on the shared storage system:

- 1 With the remaining node(s) shut down, open **Disk Management** on node 1.
- 2 Allow Windows to enter a signature on all new physical or logical drives.



NOTE: Do not create dynamic disks on your hard drives.

- 3 Locate the icon for the first unnamed, unformatted drive on the shared storage system.
- 4 Right-click the icon and select **Create** from the submenu.

If the unformatted drives are not visible, verify the following:

- The HBA driver is installed.
- The storage system is properly cabled to the servers.
- The virtual disks have been assigned to the host.
- The PowerPath software installed on the server can access all data paths to the virtual disks.

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




NOTE: The MSCS software allows only one node to access a logical drive at a time. If a logical drive is partitioned into multiple disks, only one node is able to access all the partitions for that logical drive. If each node needs to access a separate disk, two or more logical drives must be present in the storage system.

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

To assign a drive letter to an NTFS volume:

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

To create a mount point:

- a** Click **Add**.
 - b** Click **Mount** in the following empty NTFS folder.
 - c** Type the path to an empty folder on an NTFS volume, or click **Browse** to locate it.
 - d** Click **OK**.
 - e** Go to step 9.
- 9** Click **Yes** to confirm the changes.
 - 10** Right-click the drive icon again and select **Format** from the submenu.
 - 11** Under **Volume Label**, enter a descriptive name for the new volume; for example, `Disk_Z` or `Email_Data`.
 - 12** In the dialog box, change the file system to **NTFS**, select **Quick Format**, and click **Start**.
-  **NOTE:** The NTFS file system is required for shared-disk resources under MSCS.
- 13** Click **OK** at the warning.
 - 14** Click **OK** to acknowledge that the format is complete.
 - 15** Click **Close** to close the dialog box.
 - 16** Repeat step 3 through step 15 for each remaining drive.
 - 17** Close **Disk Management**.
 - 18** Shut down node 1.
 - 19** Perform the following steps on the remaining node(s), one at a time:
 - a** Turn on the node.
 - b** Open **Disk Management**.
 - c** Assign the drive letters to the drives.
This procedure allows Windows to mount the volumes.
 - d** Reassign the drive letter, if necessary.
To reassign the drive letter, repeat step 7 through step 9.
 - e** Power down the node.

Creating a Virtual Disk for the Quorum Resource

For best practice, create a separate virtual disk—approximately 1 GB in size—for the quorum resource.

When you create the virtual disk for the quorum resource:

- Format the virtual disk with NTFS.
- Use the virtual disk exclusively for your quorum logs.
- Do not store any application or user data on the quorum resource.
- To easily identify the quorum resource, Dell recommends that you assign the drive letter "Q" to the quorum resource.



NOTE: The **Majority Node Set Quorum** types for Windows Server 2003 are not supported.

Configuring Hard Drive Letters When Using Multiple Shared Storage Systems

Before installing MSCS, 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, you can assign drive letters to more than 22 logical drive letters.

The maximum amount of logical drive letters varies, depending on the operating system that is running in the cluster:

- Windows 2000 — Up to 22 logical drive letters (E through Z).
- Windows Server 2003 — Up to 22 logical drive letters plus additional volumes using volume mount points.

See "Assigning Drive Letters and Mount Points" for more information.



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 of the storage devices using Windows 2000 or Windows Server 2003 Disk Management or Dell OpenManage™ Array Manager.

Windows 2000 Disk Management displays all of the accessible disks from the first HBA, followed by all of the accessible disks from the second HBA. The disks attached to a port with a lower port number on the Fibre Channel switch are displayed first, followed by those with a higher port number.

- 2** To maintain proper drive letter assignments, ensure the first HBA detected by each node is connected to the first switch or SP-A and the second detected HBA is connected to the second switch or SP-B.

See Figure 5-4 in "Installing Your Cluster in a Direct-Attached Environment" for the location of SP-A and SP-B on the AX100/150 storage system.

- 3** Go to "Formatting and Assigning Drive Letters and Volume Labels to the Disks."

Formatting and Assigning Drive Letters and Volume Labels to the Disks

- 1** Shut down all the cluster nodes except node 1.
- 2** Format the disks, assign the drive letters and volume labels on node 1 by using the Windows Disk Management utility or Array Manager.
For example, create volumes labeled "Volume Y" for disk Y and "Volume Z" for disk Z.
- 3** Shut down node 1 and perform the following steps on the remaining node(s), one at a time:
 - a** Turn on the node.
 - b** Open **Disk Management**.
 - c** Assign the drive letters for the drives.
This procedure allows Windows to mount the volumes.
 - d** 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 **Edit**, select the letter you want to assign the drive (for example, Z), and then click **OK**.
 - Click **Yes** to confirm the changes.
- e Shut down the node.

If the cables are connected properly, the drive order is the same as on each node, and the drive letter assignments of all the cluster nodes follow the same order as on node 1. The volume labels can also be used to double-check the drive order by ensuring that the disk labeled "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.

Installing and Configuring MSCS

MSCS is a software component in Windows 2000 and an integrated service in Windows Server 2003. MSCS performs the basic cluster functionality, which includes membership, communication, and failover management. When MSCS is installed properly, the service starts on each node and responds automatically in the event that one of the nodes fails or goes offline. To provide application failover for the cluster, the MSCS software must be installed on each cluster node. For more information, see "Using Microsoft Cluster Service (MSCS)".

Verifying Cluster Readiness

To ensure that your server and storage systems are ready for MSCS installation, ensure that these systems are functioning correctly, and then verify the following:

- All cluster nodes can log on to the same domain.
- The shared disks are partitioned and formatted, and the same drive letters that refer to logical drives on the shared storage system are used on each node.
- All IP addresses and network names for each cluster node are communicating with each other and the public network.

Installing Applications in the Cluster Group

The Cluster Group contains a network name and IP address resource, which are used to manage the cluster. The Cluster Group is dedicated to cluster management and therefore for best cluster performance, do not install applications in this group.

Installing the Quorum Resource

When you install a Windows 2000 Advanced Server cluster, the installation wizard prompts you for a location to install the quorum resource. When you install a Windows Server 2003 cluster, the installation wizard automatically selects an NTFS disk as the quorum resource, which you can modify later. When you complete the procedures in the wizard, you can select another disk for the quorum using **Cluster Administrator**. To prevent quorum resource corruption, Dell and Microsoft recommend that you do not place applications or data on the disk.

Installing and Configuring Windows 2000 Cluster Service (MSCS)

After the Windows 2000 setup files are copied to the cluster node, Windows 2000 reboots and enters the second phase of installation.

During this phase, the Windows 2000 setup prompts you for the product key, licensing mode, system name, and additional information before starting the Windows 2000 Components program. From this dialog box, you can select the Windows 2000 core components for installation. Selecting **Cluster Service** copies the required files to the system.

If your operating system was preinstalled by Dell or you did not select Cluster Service when you installed the operating system, you can install the Cluster Service later by running the **Add/Remove Components** from the **Control Panel**.

After the Windows 2000 setup completes and all clustering prerequisites are met, run the Cluster Configuration Wizard to complete the installation, add additional node(s), and configure cluster resources.

To access the Cluster Configuration wizard:

- 1 Click the **Start** button and select **Run**.
- 2 In the **Run** box, type the following:
`C:\winnt\cluster\cluscfg.exe`

- 3 Click OK.
- 4 Follow the directions in the **Cluster Configuration Wizard** window.

Configuring the Public and Private Networks

When you install and configure a Windows 2000 Advanced Server cluster, the software installation wizard prompts you to identify the public and private network segments connected to your cluster nodes. For best practice, configure your cluster by performing the following steps. This procedure configures your cluster with added fault tolerance for the private network.

- 1 Set the private network (cluster interconnect) to **Use for internal communications only**.
- 2 Name this network **Private**.
- 3 Set the client public network segment(s) to **All communications**.
This setting provides a redundant path for the cluster-to-cluster communication in the event the private network fails.
- 4 For all subsequent NICs, set each NIC to **Client use only or All communications**.
- 5 Set the priority of the networks so that the network you designated as **Private** has the highest priority for internal communications.

You can set the priority of the networks when you install MSCS or when using Cluster Administrator software.

Configuring Cluster Networks Running Windows Server 2003

When you install and configure a cluster running Windows Server 2003, the software installation wizard automatically assigns and configures the public and private networks for your cluster. You can rename a network, allow or disallow the cluster to use a particular network, or modify the network role using **Cluster Administrator**. For best practice, configure at least one network for the cluster interconnect (private network) and one network for all communications.

Configuring Windows Server 2003 Cluster Service (MSCS)

The cluster setup files are automatically installed on the system disk. To create a new cluster:

- 1** Click the **Start** button, select **Programs**→**Administrative Tools**→**Cluster Administrator**.
- 2** From the **File** menu, select **Open Connection**.
- 3** In the **Action** box of the **Open Connection to Cluster**, select **Create new cluster**.

The **New Server Cluster Wizard** window appears.

- 4** Click **Next** to continue.
- 5** Follow the procedures in the wizard, and then click **Finish**.
- 6** Add the additional node(s) to the cluster.

- a** Turn on the remaining node(s).
- b** Click the **Start** button, select **Programs**→**Administrative Tools**, and then double-click **Cluster Administrator**.
- c** From the **File** menu, select **Open Connection**.
- d** In the **Action** box of the **Open Connection to Cluster**, select **Add nodes to cluster**.
- e** In the **Cluster or server name** box, type the name of the cluster or click **Browse** to select an available cluster from the list, and then click **OK**.

The **Add Nodes Wizard** window appears.

If the **Add Nodes Wizard** *does not* generate a cluster feasibility error, go to step f.

If the **Add Nodes Wizard** generates a cluster feasibility error, go to "Adding Cluster Nodes Using the Advanced Configuration Option."

- f** Click **Next** to continue.
- g** Follow the procedures in the wizard and click **Finish**.

Adding Cluster Nodes Using the Advanced Configuration Option

If you are adding additional nodes to the cluster using the Add Nodes wizard and the nodes are not configured with identical internal storage devices, the wizard may generate one or more errors while checking cluster feasibility in the **Analyzing Configuration** menu. If this situation occurs, select **Advanced Configuration Option** in the Add Nodes wizard to add the nodes to the cluster.

To add the nodes using the **Advanced Configuration Option**:

- 1 From the **File** menu in Cluster Administrator, select **Open Connection**.
- 2 In the **Action** box of the **Open Connection to Cluster**, select **Add nodes to cluster** and then click **OK**.

The **Add Nodes Wizard** window appears.

- 3 Click **Next**.
- 4 In the **Select Computers** menu, click **Browse**.
- 5 In the **Enter the object names to select (examples)**, type the names of one to seven systems to add to the cluster, with each system name separated by a semicolon.


- 6 Click **Check Names**.

The **Add Nodes Wizard** verifies and underlines each valid system name.

- 7 Click **OK**.
- 8 In the **Select Computers** menu, click **Add**.
- 9 In the **Advanced Configuration Options** window, click **Advanced (minimum) configuration**, and then click **OK**.
- 10 In the **Add Nodes** window, click **Next**.
- 11 In the **Analyzing Configuration** menu, Cluster Administrator analyzes the cluster configuration.

If Cluster Administrator discovers a problem with the cluster configuration, a warning icon appears in the **Checking cluster feasibility** window. Click the plus ("+") sign to review any warnings, if needed.

- 12 Click **Next** to continue.

- 13** In the **Password** field of the **Cluster Service Account** menu, type the password for the account used to run the Cluster Service, and click **Next**.
The **Proposed Cluster Configuration** menu appears with a summary with the configuration settings for your cluster.
- 14** Click **Next** to continue.
The new systems (hosts) are added to the cluster. When completed, **Tasks completed** appears in the **Adding Nodes to the Cluster** menu.
 **NOTE:** This process may take several minutes to complete.
- 15** Click **Next** to continue.
- 16** In the **Completing the Add Nodes Wizard** window, click **Finish**.

Verifying MSCS Operation

After you install MSCS, verify that the service is operating properly.

If you selected **Cluster Service** when you installed the operating system, see "Obtaining More Information."

If you did not select **Cluster Service** when you installed the operating system:

- 1** Click the **Start** button and select **Programs**→**Administrative Tools**, and then select **Services**.
- 2** In the **Services** window, verify the following:
 - In the **Name** column, **Cluster Service** appears.
 - In the **Status** column, **Cluster Service** is set to **Started**.
 - In the **Startup Type** column, **Cluster Service** is set to **Automatic**.

Obtaining More Information

See Microsoft's online help to configure the Cluster Service.

For more information about the Cluster Service, see "Using Microsoft Cluster Service (MSCS)".

Verifying Cluster Functionality

To verify cluster functionality, monitor the cluster network communications to ensure that your cluster components are communicating properly with each other. Also, verify that MSCS is running on the cluster nodes.

Verifying Cluster Resource Availability

In the context of clustering, a resource is a basic unit of failover management. Application programs are made up of resources that are grouped together for recovery purposes. All recovery groups, and therefore the resources that comprise the recovery groups, must be online (or in a ready state) for the cluster to function properly.

To verify that the cluster resources are online:

- 1 Start **Cluster Administrator** on the monitoring node.
- 2 Click the **Start** button and select **Programs**→**Administrative Tools (Common)**→**Cluster Administrator**.
- 3 Open a connection to the cluster and observe the running state of each resource group. If a group has failed, one or more of its resources might be offline.

Troubleshooting Failed Resources

Troubleshooting the failed resources is beyond the scope of this document, but examining the properties of each resource and ensuring that the specified parameters are correct are the first two steps in this process. In general, if a resource is offline, it can be brought online by right-clicking the resource and selecting **Bring Online** from the drop-down menu.

See the documentation and online help for Windows 2000 Advanced Server or Windows Server 2003 for information about troubleshooting resource failures.

Upgrading to a Cluster Configuration

Before You Begin

Before you upgrade your nonclustered 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."
- Verify that your hardware and storage systems are installed and configured as explained in the following sections:
 - "Cabling Your Cluster Hardware"
 - "Preparing Your Systems for Clustering"
 - "Installing Your Cluster in a Direct-Attached Environment" (for direct-attached configurations)
 - "Installing Your Cluster in a SAN-Attached Environment" (for SAN-attached cluster configurations)

Supported Cluster Configurations

Dell™ certifies and supports only solutions that are configured with the Dell products described in this guide. For a description of the cluster components, see your *Dell PowerEdge™ Cluster FE550W/FE650W Systems Platform Guide*.

Completing the Upgrade

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

The final phase for upgrading to a cluster solution is to install and configure the Microsoft® Windows® 2000 Advanced Server, Microsoft Windows Server® 2003, Enterprise Edition, or Microsoft Windows Server 2003 R2, Enterprise Edition operating system with the Microsoft Cluster Service (MSCS).

See Table 4-1 for information about installing the operating system and related software components for your cluster configuration.

Table 4-1. Installing the Operating System

Configuration	Installation Procedures
Direct-attached	See "Installing the Windows Operating System and Cluster Management Software" in "Installing Your Cluster in a Direct-Attached Environment."
SAN-attached	See "Installing the Windows Operating System and Cluster Management Software" in "Installing Your Cluster in a SAN-Attached Environment."

Upgrading Your Operating System

You can upgrade your Windows 2000 Advanced Server cluster to a Windows Server 2003 cluster using one of the following methods.



NOTE: "Windows Server 2003" refers to either Windows Server 2003, Enterprise Edition or Windows Server 2003 R2, Enterprise Edition. See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for additional information on the supported operating system versions and service packs.

- Standard upgrade — Upgrade the operating system on each cluster node while all cluster nodes are disconnected from the client network. This procedure requires you to re-create your cluster configuration.

See your Windows operating system documentation for performing a standard upgrade.

- Rolling upgrade — Upgrade the operating system on each cluster node while the remaining cluster nodes are connected to the client network and available to handle client requests. This procedure *does not* require you to re-create your cluster configuration. However, each cluster node must be configured with the appropriate resources to run all virtual servers and services for the entire cluster while you upgrade the remaining node.

The following section explains how to perform a rolling upgrade on a two-node cluster.

Performing a Rolling Upgrade

Before you perform a rolling upgrade:

- Ensure that your cluster nodes are running Windows 2000 Advanced Server.
- Back up your data and system states.
- Run the Check System Compatibility Wizard to determine if your cluster nodes are configured with the appropriate resources to run Windows Server 2003.
- Ensure that your cluster service account is an explicit member of the local administrator's group and that its user-right privileges are set to **Act as part of operating system**.



NOTE: Some cluster resources do not support a rolling upgrade. See the Microsoft Support website located at support.microsoft.com for more information.

Upgrading Node 1

- 1 Click Start and select **Programs→Administrative Tools→Cluster Administrator**.
- 2 In the **Cluster Administrator** window, right-click a node and then click **Pause Node**.

The following steps refer to the node you selected as node 1.

- 3 Right-click a cluster group and then click **Move Group**.
The cluster group is moved and restarted on node 2.
- 4 Repeat step 3 for the remaining cluster groups.

- 5 Uninstall the EMC® PowerPath™ software.
See your EMC PowerPath software documentation for more information.
- 6 Insert the *Microsoft Windows Server 2003* CD into the CD drive.
- 7 Double-click **Install Windows Server 2003**.
The **Windows Setup Wizard** window appears.
- 8 Follow the instructions in the **Windows Setup Wizard** to upgrade your operating system.



NOTE: If you are running the IIS World Wide Web Publishing service on your cluster node, this service is disabled during the upgrade to protect your system.

- 9 Reinstall the host bus adapter (HBA) drivers, PowerPath software, and any other required software applications.
See your EMC PowerPath software documentation for more information.
- 10 Verify that the upgraded node is functioning correctly.
 - a Click **Start** and select **Programs**→**Administrative Tools**→**Cluster Administrator**.
 - b Move one or more cluster resource groups from node 2.
 - c Verify that the resource group(s) can be brought online.
 - d Close the **Cluster Administrator** window.
 - e Remove the CD from the CD drive.


- 11 Go to "Upgrading Node 2."



NOTE: After you upgrade node 1, your cluster is running two separate operating systems. For best practice, do not modify your cluster configuration—such as adding or removing cluster nodes or resources—until you upgrade both cluster nodes.

Upgrading Node 2

- 1 On node 2, click **Start** and select **Programs**→**Administrative Tools**→**Cluster Administrator**.
- 2 In the **Cluster Administrator** window, right-click node 1 and then click **Resume Node**.

- 3 Right-click node 2 and then click **Pause Node**.
- 4 Right-click a cluster group and then click **Move Group**.
The cluster group is moved and restarted on node 1.
- 5 Repeat step 4 for the remaining cluster groups.
- 6 Uninstall the PowerPath software.
See your EMC PowerPath software documentation for more information.
- 7 Insert the *Microsoft Windows Server 2003* CD into the CD drive.
- 8 Double-click **Install Windows Server 2003**.
The **Windows Setup Wizard** window appears.
- 9 Follow the instructions in the **Windows Setup Wizard** to upgrade your operating system.
 **NOTE:** If you are running IIS World Wide Web Publishing service on your cluster node, this service is disabled during the upgrade to protect your system.
- 10 Reinstall the HBA drivers, PowerPath software, and any other required software applications.
See your EMC PowerPath software documentation for more information.
- 11 Verify that the upgraded node is functioning correctly.
 - a Click **Start** and select **Programs**→**Administrative Tools**→**Cluster Administrator**.
 - b Move one or more cluster resource groups from node 1.
 - c Verify that the resource group(s) can be brought online.
 - d Remove the CD from the CD drive.
- 12 In the **Cluster Administrator** window, redistribute the cluster groups to the appropriate cluster nodes.
- 13 Close the **Cluster Administrator** window.

Installing Your Cluster in a Direct-Attached Environment


Before You Begin

Before you install the operating system, you must make specific low-level software configurations to the cluster.

Before you install the Microsoft® Windows® operating system on your cluster nodes, verify that your cluster hardware and storage systems meet the following requirements:

- The minimum system requirements for your cluster solution as described in "System Requirements"
- Hardware and storage systems are installed and configured properly as explained in the following sections:
 - "Cabling Your Cluster Hardware"
 - "Preparing Your Systems for Clustering"

After you have completed these preliminary procedures, you can begin installing and configuring the operating system on your cluster nodes.

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

Connecting the Storage Systems to Your Cluster

This section provides information for connecting your cluster to one or more storage systems using direct-attached cable connections.

Direct-Attached Cluster

A direct-attached cluster configuration consists of one or two Fibre Channel host bus adapters (HBAs) cabled directly to a Dell | EMC AX100SC/150SC or AX100/150 storage system, respectively. Direct-attach configurations are self-contained and do not share any physical resources with other server or storage systems outside of the cluster.

Figure 5-1 shows an example of a direct-attached, single cluster configuration that is connected to an AX100SC/150SC storage system with one HBA installed in each cluster node.

Figure 5-1. Direct-Attached Cluster Configuration Using the AX100SC/150SC Single SP Storage System





Figure 5-2 shows an example of a direct-attached, single cluster configuration that is connected to an AX100/150 storage system with two HBAs installed in each cluster node.

Figure 5-2. Direct-Attached Cluster Configuration Using the AX100/150 Dual SP Storage System



Cabling One Cluster to a Dell | EMC Storage System

Each cluster node attaches to the storage system using fiber optic cables with duplex LC multimode connectors that attach to the HBAs in the cluster nodes and the SP ports in the Dell | EMC storage system. These connectors consist of two individual fibre optic connectors with indexed tabs that must be aligned properly into the HBAs and SP ports.

➡ NOTICE: Do not remove the connector covers until you are ready to insert the connectors into the system.


The duplex LC multimode connectors attach to the SP ports (SP-A and SP-B [if applicable]) on the back of the storage system.

Cabling a Two-Node Cluster to an AX100SC/150SC Storage System

- 1** Connect cluster node 1 to the storage system by installing a cable from cluster node 1 HBA0 to SP-A port FE 0.
- 2** Connect cluster node 2 to the storage system by installing a cable from cluster node 2 HBA0 to SP-A port FE 1.

Figure 5-3 illustrates how to cable a two-node direct-attached cluster to an AX100SC/150SC storage system. Each cluster node is configured with one HBA that is connected to a single SP in the storage system. This cluster configuration does not provide data redundancy.

Figure 5-3. Cabling the Cluster Nodes to the AX100SC/150SC Storage System




Cabling a Two-Node Cluster to an AX100/150 Storage System

- 1** Connect cluster node 1 to the storage system.
 - a** Connect a cable from cluster node 1 HBA0 to SP-A port FE 0.
 - b** Connect a cable from cluster node 1 HBA1 to SP-B port FE 1.
- 2** Connect cluster node 2 to the storage system.
 - a** Connect a cable from cluster node 2 HBA0 to SP-B port FE 0.
 - b** Connect a cable from cluster node 2 HBA1 to SP-B port FE 1.

Figure 5-3 illustrates how to cable a two-node direct-attached cluster to an AX100/150 storage system. Each cluster node is configured with two HBAs that are connected to dual SPs in the storage system. This cluster configuration provides data redundancy and eliminates a single point of failure on the cluster nodes and the storage system.

Figure 5-4. Cabling the Cluster Nodes to the AX100/150 Storage System



Cabling Two Clusters to a Dell | EMC Storage System

The AX100/150 storage system does not support more than one direct-attached 2-node cluster.

Installing the Windows Operating System and Cluster Management Software

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.

See "Installation Overview" in "Preparing Your Systems for Clustering" for information about installing the Windows operating system and the cluster management software on your cluster nodes.

Installing Your Cluster in a SAN-Attached Environment

SAN Overview


A SAN is a high-performance network that is used to move data between server and storage resources, providing scalable storage in a networked environment. Using Fibre Channel technology and fiber optic cabling, a SAN provides the high-speed connection between the cluster nodes and the external storage systems.

Table 6-1 lists the key components in a SAN.

Table 6-1. Key SAN Components

Component	Description
Cluster nodes	Dell™ PowerEdge™ systems used as cluster nodes in the cluster. Each system must have dual host bus adapters (HBAs) installed in the server to manage the I/O data transfer from the server's internal bus to the storage system.
Interconnects	The connections between the cluster nodes and the storage systems. These connections include switches and cabling.
Storage system	Provides external data storage for the host systems (cluster nodes). A SAN cluster only supports the Dell EMC AX100/150 storage system.
Fabrics	Private connections of one or more Fibre Channel switches that provide connections between the cluster nodes and storage systems. NOTE: The Cluster FE550W and FE650W solution supports only one switch in a fabric.

Connecting the Storage Systems to Your Cluster

 **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.

This section provides information for connecting your cluster to one or more storage systems using a SAN (Fibre Channel switch fabric).

SAN-Attached Cluster

A SAN-attached cluster is a cluster configuration where all cluster nodes are attached to a single storage system or to multiple storage systems through a SAN using a redundant switch fabric.

SAN-attached cluster configurations provide more flexibility, expandability, and performance than direct-attached configurations.

For more information about Fibre Channel switch fabrics, see "Fibre Channel Switch Fabric".

Figure 6-1 shows an example of a two node, SAN-attached cluster running Microsoft® Windows® 2000 Advanced Server. Each cluster node is configured with two HBAs that are connected to two Fibre Channel switches. The storage system includes two SPs that provide redundant connections to the Fibre Channel switches. This configuration provides redundancy and eliminates a single point of failure on the host or storage system.

Figure 6-1. SAN-Attached Cluster Running Windows 2000 Advanced Server





Figure 6-2 shows an example of an eight-node, SAN-attached cluster running Windows Server[®] 2003. This configuration is similar to a SAN-attached cluster running Windows 2000 Advanced Server, but also provides scalability, additional memory support, and easier cluster management using Configure Your Server and Manage Your Server wizards.

Figure 6-2. SAN-Attached Cluster Running Windows Server 2003



Cabling One SAN-Attached Cluster to an AX100/150 Storage System

The AX100/150 storage system is configured with one uninterrupted power supply (UPS) that provides backup power for the write cache in the event of a power failure. The cluster nodes attach to the storage system using a redundant switch fabric and fiber optic cables with duplex LC multimode connectors.

The switches, the HBAs in the cluster nodes, and the SP ports in the storage system use duplex LC multimode connectors. The connectors consist of two individual fiber optic connectors with indexed tabs that must be inserted and

aligned properly in the small form-factor pluggable (SFP) module connectors on the Fibre Channel switches and the connectors on the cluster nodes and storage systems.

For more information about the duplex LC multimode fiber optic connector, see "Fibre Channel Cable Connectors".

Each HBA is cabled to a port on a Fibre Channel switch. Two cables connect from outgoing ports on the Fibre Channel switch to a storage processor (SP) on the AX100/150 storage system.

Table 6-2 provides information for cabling your storage system to the Fibre Channel switch.

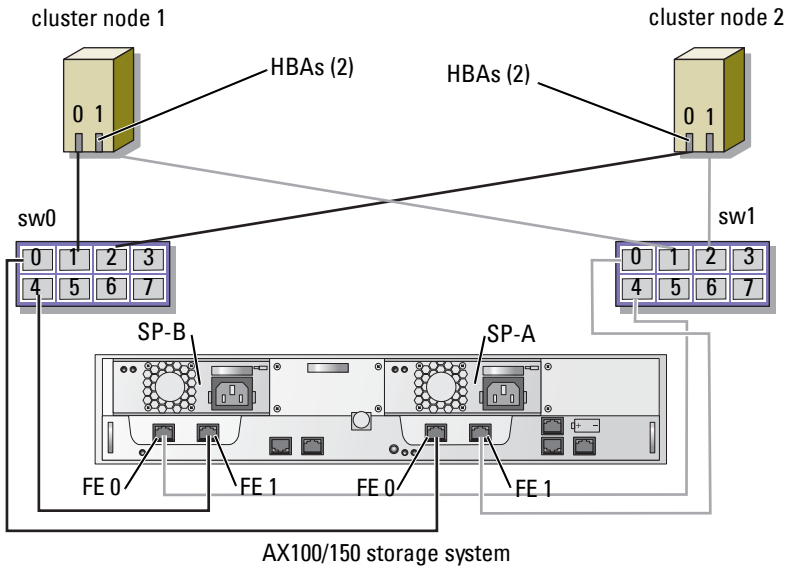
Table 6-2. Storage System Cabling Description

Storage System	SP Ports	Fiber Optic Cables Required	Cabling Description
AX100/150	Two ports per SP	4	Attach one cable from each SP to a port on each Fibre Channel switch.

Figure 6-3 illustrates how to cable a SAN-attached cluster to the AX100/150 storage system.

Customer-installable switches with hard-zones are pre-configured to use port 0 and port 4 for SP connections. Ports 1-3 and ports 5-7 (in 8-port switch) and ports 1-3 and ports 5-15 (in 16-port switch) are used for HBA connections. For high-availability, the FE ports on SP-A to switch port 0 or the FE port on SP-B to switch port 4.

Figure 6-3. Cabling a SAN-Attached Cluster to and Dell | EMC AX100/150 DPE



NOTE: Figure 6-3 shows an 8-port Fibre Channel switch cable configuration. The 16-port Fibre Channel switch cable configuration is identical.

- 1 Connect cluster node 1 to the SAN. See Figure 6-3.
 - a Connect a cable from HBA0 to Fibre Channel switch 0 (sw0) port 1.
 - b Connect a cable from HBA1 to Fibre Channel switch 1 (sw1) port 1.
- 2 Connect cluster node 2 to the SAN. See Figure 6-3.
 - a Connect a cable from HBA0 to Fibre Channel switch 0 (sw0) port 2.
 - b Connect a cable from HBA1 to Fibre Channel switch 1 (sw1) port 2.
- 3 Repeat step 2 for each remaining cluster node, replacing port 2 with port 3, 5, and so on. See Figure 6-3.

NOTE: For best practice, connect the cables to identical switch ports on both switches in ascending order. For example, after you perform step 2, connect the HBAs on node 3 to port 3 on sw0 and sw1, which are the next available ports on each switch.

- 4 Connect the storage system to the SAN. See Figure 6-3.
 - a Connect a cable from Fibre Channel switch 0 (sw0) port 0 to SP-A port FE 0.
 - b Connect a cable from Fibre Channel switch 0 (sw0) port 4 to SP-B port FE 1.
 - c Connect a cable from Fibre Channel switch 1 (sw1) port 0 to SP-A port FE 1.
 - d Connect a cable from Fibre Channel switch 1 (sw1) port 4 to SP-B port FE 0.

Cabling Multiple SAN-Attached Clusters to a Dell | EMC Storage System

To cable two or three clusters to the storage system, connect each cluster node to the appropriate Fibre Channel switch, and then connect the Fibre Channel switches to the appropriate SPs on the processor enclosure.

See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for rules and guidelines for SAN-attached clusters.



NOTE: The following procedures use Figure 6-3 as an example for cabling additional clusters.

Cabling Multiple SAN-Attached Clusters to the AX100/150 Storage System

- 1 In the first cluster, connect cluster node 1 to the SAN.
 - a Connect a cable from HBA0 to Fibre Channel switch 0 (sw0) port 1.
 - b Connect a cable from HBA1 to Fibre Channel switch 1 (sw1) port 1.
- 2 In the first cluster, repeat step 1 for each remaining cluster nodes, replacing port 1 with the next available port.
- 3 For each additional cluster, repeat step 1 and 2, replacing used port numbers with the next available ports on the switches.

- 4 Connect the storage system to the SAN.
 - a Connect a cable from Fibre Channel switch 0 (sw0) port 0 to SP-A port FE 0.
 - b Connect a cable from Fibre Channel switch 0 (sw0) port 4 to SP-B port FE 1.
 - c Connect a cable from Fibre Channel switch 1 (sw1) port 0 to SP-A port FE 1.
 - d Connect a cable from Fibre Channel switch 1 (sw1) port 4 to SP-B port FE 0.

Connecting a PowerEdge Cluster to Multiple Storage Systems

You can increase your cluster storage capacity by attaching multiple storage systems to your cluster using a redundant switch fabric. PowerEdge cluster systems can support configurations with multiple storage units attached to clustered servers. In this scenario, the Microsoft Cluster Service (MSCS) software can fail over disk drives in any cluster-attached shared storage array between the cluster nodes.


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

- A maximum of four storage systems per cluster.
- The shared storage systems and firmware must be identical. Using dissimilar storage systems and firmware for your shared storage is not supported.
- MSCS is limited to 22 drive letters. Because drive letters A through D are reserved for local disks, a maximum of 22 drive letters (E to Z) can be used for your storage system disks.
- Windows Server 2003 supports mount points, allowing greater than 22 drives per cluster

See "Assigning Drive Letters and Mount Points" for more information.

Figure 6-4 provides an example of cabling the cluster nodes to two Dell | EMC storage systems.

Figure 6-4. PowerEdge Cluster Nodes Cabled to Two Storage Systems



Implementing Port Zoning on a Fibre Channel Switched Fabric

Fibre Channel switches provide you with the ability to set up barriers between different devices and operating environments. These barriers create logical fabric subsets with minimal software and hardware intervention. Similar to subnets in the client/server network, logical fabric subsets divide a fabric into similar groups of components, regardless of their proximity to one another. The logical subsets that form these barriers are called *zones*.

Zoning automatically and transparently enforces access of information to the zone devices. More than one PowerEdge cluster configuration can share Dell | EMC storage system(s) in a switched fabric using Fibre Channel switch zoning and Shared Storage Control. By using Fibre Channel switches to implement zoning, you can segment the SANs to isolate heterogeneous servers and storage systems from each other.

Zoning can be implemented based on ports or a worldwide name (WWN). A WWN is a unique numeric identifier assigned to Fibre Channel interfaces, such as HBA ports, SP ports, and Fibre Channel to SCSI bridges or SNCs. Port zoning or hard zoning method restricts cluster communications between the nodes and the storage system to specific ports on the Fibre Channel

switch. Unlike the WWN "soft zoning" method where zones are enforced by the name server, port zoning restricts access to any device that is not connected to a specific switch port, thereby maximizing network security and unauthorized data access.

Table 6-3 provides a comparison between port zoning and WWN zoning.

Table 6-3. Port Zoning and WWN Zoning Descriptions

Port Zoning	WWN Zoning
Restricts access to physical switch ports in the fabric	Restricts access to WWNs in the fabric
Unauthorized users cannot access	Unauthorized users can access a device if they have the device's IEEE adapter address

Using Port Zoning

The hard-zoned Fibre Channel switches are preconfigured by Dell with single initiator / single target port zoning. Port 0 and 4 on the Fibre Channel switches are configured for the storage system (target); the remaining ports are configured for the HBAs (or initiators). This configuration scheme allows each HBA to be zoned to the storage ports. For example, port 1 (initiator port) can be zoned to port 0 and port 4 (storage ports).

Using Worldwide Name zoning

A WWN consists of an 8-byte hexadecimal number with each byte separated by a colon. For example, 10:00:00:60:69:00:8a is a valid WWN. Using WWN port name zoning allows you to move cables between switch ports within the fabric without having to update the zones.

Table 3-3 provides a list of WWN identifiers that you can find in the Dell | EMC cluster environment.

Table 6-4. Port Worldwide Names in a SAN Environment

Identifier	Description
xx:xx:00:60:69:xx:xx:xx	Dell EMC or Brocade switch
50:06:01:6x:xx:xx:xx:xx	Dell EMC storage processor

Table 6-4. Port Worldwide Names in a SAN Environment (continued)

Identifier	Description
xx:xx:00:E0:8B:xx:xx:xx	QLogic HBA ports (non-embedded)
xx:xx:xx:60:45:xx:xx:xx	PowerVault 132T and 136T tape libraries
xx:xx:xx:E0:02:xx:xx:xx	PowerVault 128T tape autoloader
xx:xx:xx:C0:01:xx:xx:xx	PowerVault 160T tape library and Fibre Channel tape drives



NOTE: When you replace a Fibre Channel HBA in a PowerEdge system, reconfigure your zones to provide continuous client data access. Additionally, when you replace a switch module, reconfigure your zones to prevent data loss or data corruption.

Single Initiator Zoning

Each host HBA port in a SAN must be configured in a separate zone on the switch with the appropriate storage ports. This zoning configuration, known as *single initiator zoning*, prevents different hosts from communicating with each other, thereby ensuring that Fibre Channel communications between the HBAs and their target storage systems do not affect one another.

When you create your single-initiator zones, follow these guidelines:

- Create a zone for each HBA port and its target storage devices.
- AX100 and AX150 storage systems can be connected to a maximum of 16 and 20 HBA ports respectively in a SAN-attached environment.
- Each host can be connected to a maximum of four storage systems.
- The integrated bridge or SNC on a tape library can be added to any zone.



NOTICE: To prevent data loss, data corruption, or data unavailability, configure your zones before you configure the LUNs and storage groups.

SAN-Attached Cluster Configurations

Certain rules and requirements apply when you configure your SAN-attached cluster. See the *Dell PowerEdge Cluster FE550W/FE650W Systems Platform Guide* for information about supported servers, specific HBA models supported by the platform, and PCI slot configuration guidelines.

Table 6-5 provides a list of the documentation you will need to configure a SAN-attached cluster.

Table 6-5. SAN-Attached Cluster Configuration Documentation


Information	Documentation	Location
General rules and guidelines for SAN-attached cluster configurations.	This document	Documentation included with your PowerEdge cluster or at the Dell Support website located at support.dell.com .
Rules and requirements for cluster consolidation configurations.	Your <i>Platform Guide</i> and this document	
Using Windows 2000 Advanced Server or Windows Server 2003 for cluster products and components.	This document	Documentation included with your PowerEdge cluster or at the Dell Support website located at support.dell.com .
Initializing and configuring the EMC® storage system.	EMC storage system documentation	Documentation included with your storage system or at the EMC Support website located at: www.emc.com/dell-ax100support for AX100 storage system www.emc.com/dell-ax150support for AX150 storage system
Latest firmware and software revision requirements.	EMC CLARiiON® AX100/150 Series storage system Support Matrix	EMC Technical Library located at www.emc.com
	Your <i>Platform Guide</i>	Documentation included with your PowerEdge cluster, or at the Dell Support website located at support.dell.com .

Installing the Windows Operating System and Cluster Management Software

See "Installation Overview" in "Preparing Your Systems for Clustering" for information about installing the Windows operating system and the cluster management software on your cluster nodes.

Maintaining Your Cluster

Adding a Network Adapter to a Cluster Node

 **NOTE:** Microsoft® Windows® 2000 Advanced Server or Windows Server® 2003 (including the latest service packs) and Microsoft Cluster Service (MSCS) must be installed on both the nodes to perform this procedure.

1 Move all resources from the node you are upgrading to another node. See the MSCS documentation for information about moving cluster resources to a specific node.


2 Shut down the node you are upgrading.

3 Install the additional network adapter.

See the system's *Installation and Troubleshooting Guide* for expansion card installation instructions.

4 Turn on the node and allow the Windows operating system to boot.

Windows detects the new adapter and installs the appropriate drivers.

 **NOTE:** If Windows *does not* detect the new network adapter, the network adapter is not supported.

5 Update the network adapter drivers (if required).


6 Configure the network adapter addresses:

a Click **Start**, select **Control Panel**, and then double-click **Network Connections**.

b In the **Connections** box, locate the new adapter that you installed in the system.

c Right-click the new adapter and select **Properties**.

d Assign a unique static IP address, subnet mask, and gateway.

 **NOTE:** Ensure that the network IDs of the first network adapter and the new network adapter are unique. For example, if the first network adapter's IP address is 192.168.1.101 with a subnet mask of 255.255.255.0, assign the second adapter's IP address to 192.168.2.102 with a subnet mask of 255.255.255.0.

- 7 Click **OK** and exit the network adapter properties.
- 8 Click **Start** and select **Programs**→**Administrative Tools**→**Cluster Administrator**.
- 9 Click the **Network** tab.
- 10 Verify that a new resource labeled "New Cluster Network" appears in the window.
To rename the new resource, right-click the resource and enter a new name.
- 11 Move all cluster resources back to the original node.
- 12 Repeat step 2 through step 11 on each node.



NOTE: For each node, ensure that you assign the IP address on the same subnet as you did on the first node.

If the installation and IP address are configured correctly, all of the new network adapter resources appear online and respond successfully to ping commands.

Changing the IP Address of a Cluster Node on the Same IP Subnet



NOTE: If you are migrating your nodes to a different subnet, take all cluster resources offline and then migrate all nodes together to the new subnet.

- 1 Open **Cluster Administrator**.
- 2 Stop **MSCS** on the node.
The Cluster Administrator utility running on the second node indicates that the first node is down by displaying a red icon in the **Cluster Service** window.
- 3 Reassign the IP address.
- 4 If you are running **DNS**, verify that the **DNS** entries are correct (if required).
- 5 Restart **MSCS** on the node.
The nodes re-establish their connection and Cluster Administrator changes the node icon back to blue to show that the node is back online.

Uninstalling MSCS From Clusters Running Windows 2000 Advanced Server

- 1 Take all resource groups offline or move them to another cluster node.
- 2 Stop MSCS on the node that you want to uninstall.
- 3 Click **Start** and select **Settings**→**Control Panel**→**Add/Remove Programs**.
- 4 Select **Add/Remove Windows Components**.
- 5 Deselect the check box for MSCS, click **Next**, and then click **Finish**.
- 6 From the remaining node, click **Start** and select **Programs**→**Administrative Tools**→**Cluster Administrator**.
- 7 In the **Cluster Administrator** window, right-click the node icon and select **Options**→**Evict Node**.
- 8 Close **Cluster Administrator**.


Removing Nodes From Clusters Running Windows Server 2003

- 1 Move all resource groups to another cluster node.
- 2 Click **Start**, select **Programs**→**Administrative Tools**, and then double-click **Cluster Administrator**.
- 3 In the **Cluster Administrator** window, right-click the icon of the node you want to uninstall and then select **Stop Cluster Service**.
- 4 In the **Cluster Administrator** window, right-click the icon of the node you want to uninstall and then select **Evict Node**.

If you cannot evict the node or the node is the last node in the cluster:

- a Open a command prompt.
 - b Type `cluster node <node_name> /force`
where `<node_name>` is the cluster node you are evicting from the cluster.
- 5 Close the **Cluster Administrator** window.

Running `chkdsk /f` on a Quorum Resource

 **NOTE:** You cannot run the `chkdsk` command with the `/f` (fix) option on a device that has an open file handle active. Because MSCS maintains an open file handle on the quorum resource, you cannot run `chkdsk /f` on the hard drive that contains the quorum resource.

- 1 Move the quorum resource temporarily to another drive.
- 2 Right-click the cluster name and select **Properties**.
- 3 Click the **Quorum** tab.
- 4 Select another disk as the quorum resource and press <Enter>.
- 5 Run `chkdsk /f` on the drive that previously stored the quorum resource.
- 6 Move the quorum resource back to the original drive.

Recovering From a Corrupt Quorum Disk

The quorum disk maintains the configuration data necessary for recovery when a node fails. If the quorum disk resource is unable to come online, the cluster does not start and all of the shared drives are unavailable. If this situation occurs and you must run `chkdsk` on the quorum disk, start the cluster manually from the command line.

To start the cluster manually from a command line prompt:

- 1 Open a command line window.
- 2 Select the cluster directory by typing one of the following:
`cd \winnt\cluster` (for Windows 2000 Advanced Server)
or
`cd \windows\cluster` (for Windows Server 2003)
- 3 Start MSCS in manual mode (on one node only) with no quorum logging by typing the following:
`Clussvc -debug -noquorumlogging`
MSCS starts.

- 4 Run `chkdsk /f` on the disk designated as the quorum resource:
 - a Open a second command line window.
 - b Type `chkdsk /f`.
- 5 After the `chkdsk` utility completes, stop MSCS by pressing `<Ctrl><c>` in the first command line window.
- 6 Restart MSCS from the Services console:
 - a Click **Start** and select **Programs**→**Administrative Tools**→**Services**.
 - b In the **Services** window, right-click **Cluster Service**.
 - c In the drop-down menu, click **Start**.
 - d At the command line prompt in either window, type `Net Start Clussvc`.
The Cluster Service restarts.

For more information about recovering from a corrupt quorum disk, see the Microsoft Knowledge Base article 258078 located on the Microsoft Support website at support.microsoft.com.

Changing the MSCS Account Password in Windows Server 2003

To change the service account password for all nodes running Windows Server 2003, type the following at a command line prompt:

```
Cluster /cluster:[cluster_name] /changePASS
```

where `cluster_name` is the name of your cluster


For help changing the password, type:

```
cluster /changePASS /help
```



NOTE: Windows Server 2003 does not accept blank passwords for MSCS accounts.

Reformatting a Cluster Disk

 **NOTICE:** Ensure that all client systems are disconnected from the cluster disk before you perform this procedure.

- 1 Click **Start** and select **Programs**→**Administrative Tools**→**Cluster Administrator**.
- 2 In the **Cluster Administrator** left window, expand the **Groups** directory.
- 3 In the **Groups** directory, right-click the cluster resource group that contains the disk to be reformatted and select **Take Offline**.
- 4 In the **Cluster Administrator** right window, right-click the physical disk you are reformatting and select **Bring Online**.
- 5 In the **Cluster Administrator** right window, right-click the physical disk you are reformatting and select **Properties**.

The **Properties** window appears.

- 6 Click the **Advanced** tab.
- 7 In the "Looks Alive" **poll interval** box, select **Specify value**.
- 8 In the **Specify value** field, type:

6000000

where 6000000 equals 6,000,000 milliseconds (100 minutes)

- 9 Click **Apply**.
- 10 On the Windows desktop, right-click the **My Computer** icon and select **Manage**.

The **Computer Management** window appears.

- 11 In the **Computer Management** left pane, click **Disk Management**.
The physical disk information appears in the right pane.

- 12 Right-click the disk you want to reformat and select **Format**.
Disk Management reformats the disk.

- 13** In the **File** menu, select **Exit**.
- 14** In the "**Looks Alive**" poll interval box, select **Use value from resource type** and click **OK**.
- 15** In the **Cluster Administrator** left pane, right-click the cluster group that contains the reformatted disk and select **Bring Online**.
- 16** In the **File** menu, select **Exit**.

Using Microsoft Cluster Service (MSCS)

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 the Microsoft® Cluster Service (MSCS)

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 both networks.

Preventing Network Failure

When you install MSCS, identify the public and private network segments connected to your cluster nodes. To ensure cluster failover and noninterrupted 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 network priorities and roles when you install MSCS or when you use the Microsoft Cluster Administrator software.

Node-to-Node Communication

The Cluster Service does not use public-only networks for internal communication. If the networks are configured as private and all communications fail, the nodes cannot communicate with each other, even if the private network is still functioning. As a result, one of the nodes terminates MSCS.

Network Interfaces

You can use Cluster Administrator or another cluster management application 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 MSCS.

Each node in a cluster:

- Attaches to one or more cluster storage devices that store all of the cluster's configuration and resource data; nodes have access to all cluster configuration data
- 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 8-1 defines states of a node during cluster operation

Table 8-1. 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.

Table 8-1. Node States and Definitions (continued)

State	Definition
Up	The node is actively participating in all cluster operations, including hosting cluster groups.
Unknown	The node state cannot be determined.

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

Forming a New Cluster

MSCS 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 quorum resource and form 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. MSCS 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

When MSCS makes a resource request through a dynamic link library (DLL), the Resource Monitor checks and controls the resource's state.

Setting Resource Properties

Using the resource **Properties** dialog box, you can perform the following tasks:

- View or change the resource name, description, and possible owners.
- Assign a separate resource memory space.
- View the resource type, group ownership, and resource state.
- View which node currently owns the resource.
- View pre-existing dependencies and modify resource dependencies.
- Restart a resource and configure the resource settings (if required).
- Check the online state of the resource by configuring the **Looks Alive** (general check of the resource) and **Is Alive** (detailed check of the resource) polling intervals in MSCS.
- Specify the time requirement for resolving a resource in a pending state (**Online Pending** or **Offline Pending**) before MSCS places the resource in **Offline** or **Failed** status.
- Set specific resource parameters.

The **General**, **Dependencies**, and **Advanced** tabs are the same for every resource; however, some resource types support additional tabs.



NOTE: Do not update cluster object properties on multiple nodes simultaneously. See the MSCS online documentation for more information.

Resource Dependencies

MSCS 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 8-2 shows resources and their dependencies.



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

Table 8-2. 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

Setting Advanced Resource Properties

By using the **Advanced** tab in the **Properties** dialog box, you can perform the following tasks:

- Restart a resource or allow the resource to fail.
See "Adjusting the Threshold and Period Values" for more information.
- Adjust the **Looks Alive** or **Is Alive** parameters.
- Select the default number for the resource type.
- Specify the time parameter for a resource in a pending state.

Resource Parameters

The **Parameters** tab in the **Properties** dialog box is available for most resources. Table 8-3 shows each resource and its configurable parameters.

Table 8-3. Resources and Configurable Parameters

Resource	Configurable Parameters
File share	Share permissions and number of simultaneous users Share name (clients systems detect the name in their browse or explore lists) Share comment Shared file path

Table 8-3. Resources and Configurable Parameters (continued)


Resource	Configurable Parameters
IP address	IP address Subnet mask Network parameters for the IP address resource (specify the correct network)
Network name	Cluster name or virtual server
Physical disk	Hard drive for the physical disk resource (cannot be changed after the resource is created)

Quorum Resource

Normally, the quorum resource is a common cluster resource that is accessible by all of the nodes. The quorum resource—typically a physical disk on a shared storage system—maintains data integrity, cluster unity, and cluster operations.

When the cluster is formed or when the nodes fail to communicate, the quorum resource guarantees that only one set of active, communicating nodes is allowed to form a cluster. If a node fails and the node containing the quorum resource is unable to communicate with the remaining nodes, MSCS shuts down the node that does not control the quorum resource. If a node fails, the configuration database helps the cluster recover a failed resource or recreates the cluster in its current configuration.

The shared physical disk is the only resource supported by the solution that can act as a quorum resource.

 **NOTE:** The Majority Node Set Quorum resource type is not supported.

Additionally, the quorum resource ensures cluster integrity. MSCS uses the quorum resource's recovery logs to update the private copy of the cluster database in each node, thereby maintaining the correct version of the cluster database and ensuring that the cluster is intact.

The operating system uses the quorum resource to ensure that only one set of active, communicating nodes is allowed to operate as a cluster. A node can form a cluster only if the node can gain control of the quorum resource. A node can join a cluster or remain in an existing cluster only if it can communicate with the node that controls the quorum resource.

Resource Failure

MSCS periodically launches the Resource Monitor to check if a resource is functioning properly. Configure the **Looks Alive** and **Is Alive** polls to check for failed resources. The **Is Alive** poll interval is typically longer than the **Looks Alive** poll interval because MSCS requests a more thorough check of the resource's state.



NOTE: Do not adjust the **Looks Alive** and **Is Alive** settings unless instructed to do so by technical support.

Adjusting the Threshold and Period Values

The **Threshold** value determines the number of attempts to restart the resource before the resource fails over. The **Period** value assigns a time requirement for the **Threshold** value to restart the resource.

If MSCS exceeds the maximum number of restart attempts within the specified time period and the failed resource has not been restarted, MSCS considers the resource to be failed.



NOTE: See "Setting Advanced Resource Properties" to configure the **Looks Alive**, **Is Alive**, **Threshold**, and **Period** values for a particular resource.



NOTE: Do not adjust the **Threshold** and **Period** settings unless instructed by technical support.

Configuring Failover

You can configure a resource to affect the group and fail over an entire group to another node when a resource fails in that group. If the number of failover attempts exceeds the group's threshold and the resource is still in a failed state, MSCS attempts to restart the resource after a period of time specified by the resource's **Retry Period On Failure** property.



NOTE: Do not adjust the **Retry Period On Failure** settings unless instructed by technical support.

When you configure **Retry Period On Failure**, use the following guidelines:

- Select a unit value of minutes rather than milliseconds (the default value is milliseconds).
- Select a value that is greater than or equal to the value of the resource's restart period property.

Resource Dependencies

A dependent resource requires another resource to operate. Table 8-4 describes resource dependencies.

Table 8-4. 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 New Resource

Before you add a resource to your cluster solution, verify that the following conditions exist in your cluster:

- The type of resource is either a standard resource provided with MSCS or a custom resource provided by Microsoft or a third party vendor.
- A group that will contain the resource already exists within your cluster.
- All dependent resources have been created.
- A separate Resource Monitor exists (recommended for any resource that has caused problems in the past).

To create a new resource:

- 1 Click **Start** and select **Programs**→**Administrative Tools**→**Cluster Administrator**.

The **Cluster Administrator** window appears.

- 2 In the console tree, double-click the **Groups** folder.
- 3 Select the group to which you want the resource to belong.

- 4** On the **File** menu, point to **New** and click **Resource**.
- 5** In the New Resource wizard, type the appropriate information in the **Name** and **Description** fields and select the appropriate **Resource type** and **Group** for the new resource.
- 6** Click **Next**.
- 7** Add or remove possible owners of the resource and click **Next**.
The **New Resource** window appears with **Available resources** and **Resource dependencies** selections.
 - To add dependencies, under **Available resources**, select a resource, and then click **Add**.
 - To remove dependencies, under **Resource dependencies**, select a resource, and then click **Remove**.
- 8** Repeat step 7 for all resource dependencies and click **Finish**.
- 9** Set the resource properties.
For more information about setting resource properties, see the MSCS online help.

Deleting a Resource

- 1** Click **Start** and select **Programs**→**Administrative Tools**→**Cluster Administrator**.
The **Cluster Administrator** window appears.
- 2** In the console tree, double-click the **Resources** folder.
- 3** In the details pane, select the resource that you want to remove.
- 4** In the **File** menu, click **Offline**.
The resource must be taken offline before it can be deleted.
- 5** In the **File** menu, click **Delete**.

When you delete a resource, Cluster Administrator deletes all of the resources that are dependent on the deleted resource.

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.
- Distributed File System (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.

Failover and Failback Capabilities

Failover

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

See "Setting Advanced Resource Properties" for more information.

Cluster resources are placed in a group so that MSCS 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, MSCS 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 MSCS 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 "Setting Advanced Resource Properties" for more information.

- Specify a Possible Owner List in Microsoft Cluster Administrator for cluster resources. The Possible Owner List for a resource controls which nodes are allowed to host the resource.

See the Cluster Administrator documentation for more information.

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. See "Cabling Your Cluster for Public and Private Networks" for more information.
	The length of the interface cables exceeds the maximum allowable length.	Ensure that the fiber optic cables do not exceed 300 m (multimode).
	One of the cables is faulty.	Replace the faulty cable.
	The cluster is in a SAN, and one or more zones are not configured correctly.	Verify the following: <ul style="list-style-type: none"> • There is only one initiator (HBA) in each zone. • Each zone contains the correct initiator and the correct storage port(s).

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
One of the nodes takes a long time to join the cluster.	The node-to-node network has failed due to a cabling or hardware failure.	Check the network cabling. Ensure that the node-to-node interconnection and the public network are connected to the correct NICs.
	Long delays in node-to-node communications may be normal.	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.
Attempts to connect to a cluster using Cluster Administrator fail.	The Cluster Service has not been started.	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.
	A cluster has not been formed on the system. The system has just been booted and services are still starting.	
		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.

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
You are prompted to configure one network instead of two during MSCS 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" 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.
Using Microsoft® Windows NT® 4.0 to remotely administer a Windows® 2000 Advanced Server or Windows Server™ 2003 cluster generates error messages.	Normal. Some resources in Windows 2000, Advanced Server and Windows Server 2003 are not supported in Windows NT 4.0.	Dell recommends that you use Windows 2000 Professional, Server, or Advanced server for remote administration of a cluster running Windows 2000 Advanced Server. Similarly, use Windows XP Professional or Windows Server 2003 for remote administration of a cluster running Windows Server 2003.
Unable to add a node to the cluster.	<p>The new node cannot access the shared disks.</p> <p>The shared disks are enumerated by the operating system differently on the cluster nodes.</p>	<ul style="list-style-type: none">• 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:• Check all cable connections• Check all zone configurations• Use the "Advanced" with "Minimum" option

Table A-1. General Cluster Troubleshooting (continued)

Problem	Probable Cause	Corrective Action
The disks on the shared cluster storage appear unreadable or uninitialized in Windows Disk Administration	This situation is normal if you stopped the Cluster Service. If you are running Windows Server 2003, this situation is normal if the cluster node does not own the cluster disk.	No action required.

Cluster Data Form

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

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

Node Name	Service Tag	Public IP Address	Private IP Address

Additional Networks

Array	Array xPE Type	Array Service Tag or World Wide Name Seed	Number of Attached DAEs
1			
2			

Index

A

active/active
about, 19

C

cable configurations
cluster interconnect, 37
for client networks, 37
for mouse, keyboard, and
monitor, 33
for power supplies, 33

chkdsk/f
running, 90

cluster
configurations, 17
forming a new cluster, 97
joining an existing cluster, 97
verifying functionality, 59
verifying readiness, 54
verifying resource availability, 60

cluster configurations
active/active, 19
active/passive, 19
connecting to multiple shared
storage systems, 80
connecting to one shared storage
system, 74

direct-attached, 17, 67
SAN-attached, 18, 83
supported configurations, 61

cluster group
installing applications, 55

cluster networks
configuring Windows 2000 cluster
networks, 56
configuring Windows Server 2003
cluster networks, 56

cluster nodes
about, 96
states and definitions, 96

cluster objects
about, 95

cluster resources
configurable parameters, 99
resource dependencies, 102
resource failure, 101
setting resource properties, 98

Cluster Service
uninstalling, 89

cluster storage
requirements, 29

clustering
overview, 9

connectors
about, 35
duplex LC multimode, 35

D

- Dell | EMC AX100
 - cabling the cluster nodes, 69
 - cabling the cluster nodes in a SAN-attached environment, 76
 - cabling to one cluster, 69
 - cabling to one SAN-attached cluster, 76
 - cabling to two clusters, 71
- Dell | EMC AX100/150
 - about, 12
- direct-attached cluster
 - about, 67
 - cabling, 69
- domain model
 - selecting, 41
- drive letters
 - assigning to shared storage storage systems, 49
- dynamic disks
 - using, 48

E

- EZ Switch Setup Utility
 - about, 15

F

- failback
 - about, 105
- failover
 - configuring, 101
 - modifying failover policy, 105
 - policies, 21
- failover configurations
 - for Windows 2000 Advanced Server, 21
 - for Windows Server 2003, 22
- failover policies, 21
 - failover pair, 24
 - failover ring, 26
 - for Windows 2000 Advanced Server cluster configurations, 21
 - for Windows Server 2003, 22
 - multiway failover, 25
 - N+I failover, 23
- Fibre Channel
 - about, 16
- file share resource type, 104

G

- group affinity
 - about, 24
 - configuring, 24

H

high availability
about, 9

I

IP address
assigning to cluster resources and
components, 43
example configuration, 44

K

keyboard
cabling, 33

M

monitor
cabling, 33

mouse
cabling, 33

MSCS
installing and configuring, 54
uninstalling, 89
verifying operation, 59

multiway failover, 25

N

N+I failover
configuring group affinity, 23

Navisphere Express
about, 15

Navisphere Server Utility
about, 14

Navisphere Storage System
Initialization Utility
about, 14

network adapters
cabling the private network, 37
cabling the public network, 37
using dual-port for the private
network, 38, 47

network failure
preventing, 95

network interfaces, 96

networking
configuring Windows, 43

O

operating system
installing, 39, 42
rolling upgrade, 63
standard upgrade, 62
upgrading, 62

P

period values
adjusting, 101

power supplies
cabling, 33

PowerPath

- about, 15

private network

- cabling, 36-37
- configuring IP addresses, 44
- creating separate subnets, 46
- hardware components, 38
- hardware components and connections, 37
- using dual-port network adapters, 47

public network

- cabling, 36
- creating separate subnets, 46

Q

quorum resource

- about, 10, 100
- creating a virtual disk, 52
- definition, 10
- installing, 55
- running chkdsk, 90

R

resource

- creating, 102
- deleting, 103

resource dependencies, 98, 102

resource groups, 9

- definition, 9

resource properties, 99

S

SAN

- about, 16

SAN-attached cluster

- about, 74
- configurations, 74

shared storage

- assigning drive letters, 49
- naming and formatting drives, 49

Snapshot Management

- about, 15

SnapView

- about, 10

software

- for storage management, 14

storage management software

- EZ Switch Setup Utility, 15
- Navisphere Express, 15
- Navisphere Server Utility, 14
- Navisphere Storage System Initialization Utility, 14
- PowerPath, 15
- snapshot management, 15

storage system

- configuring drives on multiple shared storage systems, 52
- connecting to your cluster, 67, 74
- using dynamic disks and volumes, 48

subnets

- creating, 46

switches
 about, 14
system requirements, 27

T

threshold
 adjusting, 101
troubleshooting
 connecting to a cluster, 108
 failed resources, 60
 shared storage subsystem, 107

U

upgrading
 operating system, 62
upgrading to a cluster solution
 completing the upgrade, 62

V

virtual servers, 9
 definition, 9
volumes
 using, 48

W

warranty, 30
Windows 2000 Advanced Server
 cluster configurations, 21
 installing, 39, 85
Windows Server 2003
 cluster configurations, 23
Windows Server 2003,
 Enterprise Edition
 cluster configurations, 24-26
 installing, 39, 85

Z

zones
 about, 17
 implementing port zoning, 81

