

Dell™ PowerEdge™ Cluster
FE650Wi 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.

© 2006-2007 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*, and *Windows NT* are registered trademarks of Microsoft Corporation; *EMC* and *CLARiiON* are registered trademarks and *Navisphere*, *Navisphere Agent*, *Navisphere Manager*, *PowerPath*, *Access Logix*, *ControlCenter*, *MirrorView*, *SAN Copy*, and *SnapView* are trademarks of EMC Corporation.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

December 2007 Rev. A01

Contents

| | | |
|---|---|-----------|
| 1 | Introduction | 7 |
| | Overview | 7 |
| | Virtual Servers and Resource Groups | 7 |
| | Quorum Resource | 7 |
| | Shared Storage Systems | 8 |
| | Cluster Solution | 8 |
| | Operating System | 8 |
| | Storage System | 9 |
| | Storage Management Software | 10 |
| | iSCSI Components | 11 |
| | Cluster Configurations | 12 |
| | Direct-Attached Cluster | 12 |
| | iSCSI Switch-Attached Cluster | 14 |
| | Configuring Active and Passive Cluster Nodes | 14 |
| | Failover Policies | 16 |
| | Windows Server 2003 R2, Enterprise Edition Cluster Configurations | 16 |
| | System Requirements | 21 |
| | Cluster Nodes | 21 |
| | Cluster Storage | 22 |
| | Other Documents You May Need | 23 |
| 2 | Cabling Your Cluster Hardware | 25 |
| | Cabling the Mouse, Keyboard, and Monitor | 25 |
| | Cabling the Power Supplies | 25 |
| | Cabling Your Cluster For Public and Private Networks | 27 |
| | Cabling to the Public Network | 28 |
| | Cabling the Private Network | 28 |
| | NIC Teaming | 29 |

| | | |
|---|--|----|
| 3 | Preparing Your Systems for Clustering | 31 |
| | Before You Begin. | 31 |
| | Installation Overview | 31 |
| | Selecting a Domain Model | 32 |
| | Configuring the Nodes as Domain Controllers | 33 |
| | Configuring Internal Drives in the Cluster Nodes | 33 |
| | Installing and Configuring the Windows Operating System | 33 |
| | Assigning Static IP Addresses to Cluster Resources and Components | 34 |
| | Configuring IP Addresses for the Private Network. | 34 |
| | Verifying Communications Between Nodes | 36 |
| | Assigning and Preparing the Virtual Disks for Cluster nodes | 37 |
| | Assigning the Virtual Disks to a Cluster Node | 37 |
| | Using the Windows Dynamic Disks and Volumes | 37 |
| | Naming and Formatting Drives on the Shared Storage System | 38 |
| | Configuring Hard Drive Letters When Using Multiple Shared Storage Systems | 40 |
| | Installing and Configuring MSCS | 41 |
| | Verifying Cluster Readiness | 41 |
| | Installing Applications in the Cluster Group | 42 |
| | Installing the Quorum Resource | 42 |
| | Configuring Cluster Networks Running Windows Server 2003 R2, Enterprise Edition | 42 |
| | Configuring Microsoft Cluster Service (MSCS) | 43 |
| | Verifying MSCS Operation | 44 |
| | Verifying Cluster Functionality | 45 |
| | Verifying Cluster Resource Availability | 45 |
| | Troubleshooting Failed Resources | 45 |
| 4 | Installing Your Cluster in a Direct-Attached Environment | 47 |
| | Before You Begin. | 47 |
| | Connecting the Storage Systems to Your Cluster | 47 |
| | Direct-Attached Cluster | 47 |

| | | |
|----------|--|-----------|
| | Cabling One Cluster to a Dell EMC Storage System. | 49 |
| | Cabling Two Clusters to a Dell EMC Storage System | 54 |
| | Installing the Windows Operating System and Cluster Management Software. | 54 |
| 5 | Installing Your Cluster in an iSCSI Switch-Attached Environment. | 55 |
| | iSCSI Network Overview | 55 |
| | Connecting the Storage Systems to Your Cluster | 56 |
| | iSCSI Switch-Attached Cluster | 56 |
| | Cabling One iSCSI Switch-Attached Cluster to an AX150i Storage System | 58 |
| | Cabling Multiple iSCSI Switch-Attached Clusters to a Dell EMC Storage System | 63 |
| | Connecting a Dell PowerEdge Cluster to Multiple Storage Systems | 64 |
| | iSCSI Switch-Attached Cluster Configurations. | 65 |
| | Installing the Windows Operating System and Cluster Management Software. | 66 |
| 6 | Maintaining Your Cluster | 67 |
| | Adding a Network Adapter to a Cluster Node. | 67 |
| | Changing the IP Address of a Cluster Node on the Same IP Subnet | 68 |
| | Removing Nodes From Clusters | 68 |
| | Running chkdsk /f on a Quorum Resource | 69 |
| | Recovering From a Corrupt Quorum Disk | 69 |
| | Changing the MSCS Account Password | 70 |
| | Reformatting a Cluster Disk | 70 |
| 7 | Using MSCS | 73 |
| | Cluster Objects. | 73 |

| | |
|---|-----------|
| Cluster Networks | 73 |
| Preventing Network Failure. | 73 |
| Node-to-Node Communication | 73 |
| Network Interfaces | 74 |
| Cluster Nodes | 74 |
| Forming a New Cluster | 74 |
| Joining an Existing Cluster | 75 |
| Cluster Resources | 75 |
| Setting Resource Properties | 75 |
| Resource Dependencies | 76 |
| Setting Advanced Resource Properties | 76 |
| Resource Parameters. | 76 |
| Quorum Resource. | 77 |
| Resource Failure | 77 |
| Resource Dependencies | 78 |
| Creating a New Resource | 78 |
| Deleting a Resource | 79 |
| File Share Resource Type. | 79 |
| Failover and Failback Capabilities | 80 |
| Failover | 80 |
| Failback | 80 |
| | |
| 8 Troubleshooting | 83 |
| | |
| 9 Cluster Data Form | 87 |
| | |
| Index | 89 |

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 cluster nodes, or nodes) 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 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. See "Failover and Failback Capabilities" on page 80 for more information.

Virtual Servers and Resource Groups

In a cluster environment, users do not access a physical server; they access a virtual server, that 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*, that contains the cluster resources. Ownership of virtual servers and resource groups is transparent to users. See "Cluster Resources" on page 75 for more information on resource groups.

When MSCS detects a failed node or failed application, MSCS 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 Microsoft Windows NT® file system (NTFS)

See "Quorum Resource" in "Using MSCS" and the MSCS online documentation for more information.



NOTE: Dell™ PowerEdge™ FE 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 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 implements two-node to eight-node clustering and provides the following features:

- Gigabit Ethernet 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

This section provides an overview of the Windows operating system features. See your operating system documentation for a complete list of features. Windows Server 2003 R2, Enterprise Edition:


- Supports from two up to eight nodes per cluster using the AX150i (dual SP) storage system
- Supports up to 32 GB of RAM per node
- Cluster configuration and management using **Configure Your Server** and **Manage Your Server** wizards
- Provides Metadirectory Services

 **NOTE:** MSCS and Network Load Balancing (NLB) features cannot coexist on the same node, but can be used together in a multi-tiered 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 Dell | EMC AX150i-series storage systems in an Internet SCSI (iSCSI) Direct-attached or switch-attached environment. These storage systems work together with the following hardware components:

- Storage processor
- Uninterrupted Power Supply (UPS)
- Network Interface Cards (NICs) dedicated to iSCSI traffic
- Ethernet switches dedicated to iSCSI traffic

 **NOTE:** Throughout this document, "NICs dedicated to iSCSI traffic" and "iSCSI NICs" refer to the same technology.

Processor Enclosure

The processor enclosure is configured with storage processors (or SPs) 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 AX150i-series processor enclosures vary in configuration and application. Table 1-1 provides a list of features for each enclosure.

Table 1-1. Processor Enclosure Features

| AX150SCi Processor Enclosure | AX150i Processor Enclosure |
|--|---|
| One SP | Two SPs |
| Supports direct-attached clusters only | Supports direct-attached and switch-attached clusters |
| Two power supplies (second power supply is optional) | 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 and RAID 10 | Supports RAID 1, RAID 5 and RAID 10 |

UPS

A UPS is required for the dual storage system (AX150i). It supplies battery power to keep storage processor A (SP-A) running long enough to write the cached data to the disk during a complete power outage.

NICs Dedicated to iSCSI

The NIC controlled by iSCSI Software Initiator acts as an I/O adapter to connect the system's expansion bus and the storage components. Cluster solutions that are configured with the AX150SCi storage system require one iSCSI NIC in each PowerEdge system to transfer data to the storage system. Cluster solutions that are configured with the AX150i storage system require two iSCSI NICs in each PowerEdge system to provide redundant paths and load balance the I/O data transfer to the storage system.

See the *Platform Guide* for a list of supported iSCSI Initiators and NICs.

Ethernet Switches Dedicated to iSCSI

The Gigabit switch for iSCSI access functions as a regular network switch that provides extension and dedicated interconnection between the node and the storage systems. See the *Platform Guide* for a list of supported switches.

Storage Management Software

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

Navisphere Storage System Initialization Utility

The Navisphere Storage System Initialization Utility provides a user interface to initialize your AX150i-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: It is recommended that the utility be installed on at least one node in the cluster.

Navisphere Server Utility

The Navisphere Server Utility registers the cluster node with the AX150i-series storage systems, allowing the nodes to access the cluster storage data. The tool is also used for procedures, such as:


- Viewing and configuring iSCSI connections on the cluster node
- Discovering iSCSI targets on attached Dell | EMC storage systems
- Configuring mutual Challenge Handshake Authentication Protocols (CHAP) on the cluster node for target authentication
- Managing snapshot tasks




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

PowerPath iSCSI

EMC PowerPath™ iSCSI is a server-resident software that enhances application availability by detecting and re-establishing a failed connection to a processor enclosure through automatic path fail-over. PowerPath iSCSI for Windows is specifically designed for use within the Microsoft Multipathing I/O (MPIO) framework.

 **NOTE:** PowerPath iSCSI does not provide dynamic load balancing across all data paths.

 **NOTE:** Because the AX150SCi 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 AX150i-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 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.

iSCSI Components

This section provides a brief overview of iSCSI protocol and the elements of an iSCSI network.

iSCSI Protocol

iSCSI is a protocol for the new generation of storage systems that natively uses TCP/IP by encapsulating SCSI data blocks on an IP network. The SCSI data blocks are sent between the Microsoft iSCSI Initiator that resides in the host and the iSCSI target, which is usually a storage device.

Implementing iSCSI in a cluster provides the following advantage.

- Connectivity — A single technology for connection of storage systems, servers and clients within existent LAN, WAN and storage network
- Low cost for Availability — Redundant connections provide multiple data paths that are available through inexpensive TCP/IP network components
- Geographic distribution — Wider coverage of Ethernet technology allows substantial geographic remoteness of storage systems

iSCSI Initiator

The iSCSI Initiator (also known as an iSCSI client) converts blocks of SCSI commands and data into a format suitable for an IP network. The iSCSI Initiator is configured with a unique iSCSI name or Internet Qualified Name (IQN) in a PowerEdge system, allowing an administrator to identify and

manage multiple initiators in the network. For iSCSI access, the iSCSI Initiator assigns the same iSCSI name to all NICs in the PowerEdge system so they appear as one single initiator. However, each NIC is still identified by a unique IP address.

 **NOTE:** The PowerEdge FE650Wi cluster solution supports NICs with Microsoft Software Initiator and does not support iSCSI HBAs.

iSCSI Target


An iSCSI target (usually a storage device) receives iSCSI commands from the initiator. Each iSCSI target is identified by a unique IQN and each port on the storage array controller is identified by an IP address.

iSNS

The Internet Storage Name Service (iSNS) is used in an iSCSI network for automated discovery, management, and iSCSI device configuration. iSNS includes the following components:

- iSNS server - Processes iSNS registrations, deregistrations, and queries using TCP/IP from iSNS clients, and maintains this information in a database. The iSNS server resides in the iSCSI network and is similar to the simple name service (SNS) server in a switched fabric.
- iSNS client - Registers its iSCSI port information with the iSNS server. The iSNS client resides on iSCSI storage systems and hosts.

See your Microsoft and EMC documentation for information about setting the iSNS service in your iSCSI storage network.

 iSNS is supported only on Windows platforms in a iSCSI network configuration.

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 NIC in each node.

Figure 1-1 shows a basic direct-attached, single-cluster configuration that is configured with the AX150SCi (one SP) storage system.

Figure 1-1. Direct-Attached, Single-Cluster Configuration Using the AX150SCi Storage System





Figure 1-2 shows a basic direct-attached, single channel NIC single-cluster configuration that is configured with the AX150i (two SPs) storage system. Because the AX150i 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 Channel NIC, Single Cluster Configuration Using the AX150i Storage System



PowerPath Limitations in a Direct-Attached Cluster


PowerPath provides failover capabilities and multiple path detection on the same storage processor. If your cluster solution is configured with the AX150SCi storage system, PowerPath iSCSI or PowerPath standard 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.

iSCSI Switch-Attached Cluster

In an iSCSI switch-attached cluster, all of the nodes are attached to a single storage system or to multiple storage systems through redundant LANs for high availability. iSCSI switch-attached clusters are superior to direct-attached clusters in configuration flexibility, expandability, and performance.

Figure 1-3 shows an iSCSI switch-attached cluster.

Figure 1-3. iSCSI Switch-Attached 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.

Configuring Active and Passive Cluster Nodes

Active nodes process application requests and provide client services. Passive nodes are backup nodes that ensure 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 the Windows Server 2003 Release 2 (R2), Enterprise Edition operating system can support up to eight nodes in multiple configurations, as shown in Table 1-3.

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. Those node(s) can run their own cluster-aware applications (or another instance of the

same application) while providing failover for the resources on node 1. For example, multiway failover is an active/active failover solution because running applications from a failed node can migrate to multiple active nodes in the cluster. However, you must ensure that adequate resources are available on each node to handle the increased load if one node fails.

In an active/passive (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 1-2 provides a description of active/active configuration types.

Table 1-2. Active/Active Configuration Types

| Configuration Type | Active Cluster Node(s) | Definition |
|---------------------------|-------------------------------|---|
| Active ² | 2 | The active node(s) processes requests and provides failover for each other, depending on node resources and your configuration. |
| Active ³ | 3 | |
| Active ⁴ | 4 | |
| Active ⁵ | 5 | |
| Active ⁶ | 6 | |
| Active ⁷ | 7 | |
| Active ⁸ | 8 | |

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

Table 1-3. Active/Passive Configuration Types

| Configuration Type | Active Cluster Node(s) | Passive Cluster Node(s) | Description |
|---|------------------------|-------------------------|---|
| Active ¹ /Passive ¹ | 1 | 1 | The active node(s) processes requests while the passive node waits for the active node to fail. |
| Active ² /Passive ¹ | 2 | 1 | |
| Active ² /Passive ² | 2 | 2 | |
| Active ³ /Passive ¹ | 3 | 1 | |
| 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 Server 2003 R2, Enterprise Edition Cluster Configurations

Cluster configurations running Windows Server 2003 R2, Enterprise Edition provide the following failover policies:

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

Table 1-4 provides an overview of the failover policies that are implemented with Windows Server 2003 R2, Enterprise Edition. For more information, see the sections that follow this table.

Table 1-4. Windows Server 2003 R2, Enterprise Edition 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 three active nodes and one passive node. Table 1-5 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 1-5. 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 R2, Enterprise Edition, 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 R2, Enterprise Edition to assign a public property (or attribute) to a dependency between groups to ensure that they fail over to similar or separate nodes. This public 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 1-5, 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 4).

To set the public property for the cluster groups shown in Table 1-5:

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

```
cluster group "A" /prop AntiAffinityClassNames="AString"
```
- 3 Repeat step 2 for the remaining cluster groups.

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.

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 1-6 provides a failover configuration for the cluster shown in Figure 1-5.

Figure 1-5. Example of a Failover Pair Configuration




Table 1-6. Example of a Failover Pair Configuration for a Four-Node Cluster

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

Multiway Failover

Multiway failover is an active/active policy where running applications from a failed node migrate to multiple nodes in the cluster. This solution provides automatic failover and load balancing. Ensure that the failover nodes have sufficient resources to handle the workload. Figure 1-6 shows an example of four-node multiway failover configuration.

Table 1-7 shows a four-node multiway 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 fallback to avoid performance degradation. See "Using MSCS" for more information.

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




Table 1-7. Example of a Four-Node Multiway Failover Configuration

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


Failover Ring

Failover ring is an active/active policy where all running applications migrate from the failed node to the next preassigned node in the Preferred Owners List. If the failing node is the last node in the list, the failed node's applications 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 1-8 lists the hardware requirements for the cluster nodes.

Table 1-8. Cluster Node Requirements

| Component | Minimum Requirement |
|---------------|--|
| Cluster nodes | <ul style="list-style-type: none"> • Two supported PowerEdge systems running Windows Server 2003 R2, Enterprise Edition directly connected to AX150SCi or AX150i storage systems. • Two to eight PowerEdge systems running Windows Server 2003 R2, Enterprise Edition connected through a switch to the AX150i storage system. |
| Processors | At least two processors per node. |
| RAM | At least 256 MB of RAM installed per node. |

Table 1-8. Cluster Node Requirements (continued)

| Component | Minimum Requirement |
|----------------------------|--|
| iSCSI Initiator | Complete installation of the iSCSI port driver, Initiator service and Software Initiator on each node. NOTE: Microsoft MPIO Multipathing support for iSCSI is NOT installed and PowerPath is used instead. |
| NICs for iSCSI access | One or two iSCSI NIC(s) per node. Place the NIC on separate PCI buses to improve availability and performance. See the <i>Platform Guide</i> for information about supported systems, iSCSI NICs and PCI slot configuration guidelines. |
| NICs (private and public) | At least two NICs: one NIC for the public network and another NIC for the private network. NOTE: To avoid cluster communication problems due to different transfer rates and performance issues, Dell recommends that the NICs on each public network be identical, and that the NICs on each private network be identical. |
| RAID controller (optional) | One controller connected to at least two internal hard drives for each node. Two hard drives are required for mirroring (RAID 1) and at least three are required for disk striping with parity (RAID 5). Use any supported RAID controller or SCSI adapter. NOTE: Dell recommends that you use hardware-based RAID or software-based disk-fault tolerance for the internal drives. Software RAID uses the server resources and may decrease its performance. |

Cluster Storage

Table 1-9 lists supported storage systems and the configuration requirements for the cluster nodes and stand-alone systems connected to the storage systems. Table 1-10 lists hardware requirements for the cluster storage.

Table 1-9. Cluster Storage Requirements


| Hardware Components | Minimum Requirement |
|---|---|
| Supported storage systems | One AX150SCi or AX150i storage system. See Table 1-10 for specific storage system requirements. |
| Cluster nodes | All nodes must be attached to a single storage system or to multiple storage systems through an iSCSI network. |
| 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 1-10. Cluster Storage System Requirements


| Storage System | Storage Processors | Primary Storage | Backup Battery |
|-----------------------|---------------------------|-------------------------------------|-----------------------|
| AXI50SCi | 1 | At least three internal hard drives | Internal |
| AXI50i | 2 | At least four internal hard drives | One external UPS |

NOTE: The AXI50i must be connected to the UPS (that shipped with the storage) to enable write caching in the event of a power failure.

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.

- The *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 NIC documentation provides installation instructions for the NICs.
- 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.
- The *Using Dell Blade Servers in a Dell PowerEdge High Availability Cluster* document to configure Dell blade server modules in a Dell PowerEdge cluster.

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

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


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 for cluster solutions that are configured with two nodes 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



➡ NOTICE: 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



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

Cabling Your Cluster For Public and Private Networks

The network adapters in the cluster nodes provide at least three 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. |

Table 2-1. Network Connections (continued)


| Network Connection | Description |
|--------------------|--|
| iSCSI network | One or two direct or switched iSCSI connection(s) from the cluster node to the storage system. |

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

NOTE: For iSCSI cabling and connections, see "Installing Your Cluster in a Direct-Attached Environment and Installing Your Cluster in an iSCSI Switch-Attached Environment."

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.

Figure 2-3. 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 intra-cluster 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. |

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 private and iSCSI networks.

When forming a NIC team, use identical NICs and teaming drivers. Otherwise, the cluster communications may not function properly and you may experience unpredictable results.

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 for public, private and iSCSI networks
 - RAID controllers
 - 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® Server 2003 R2, Enterprise Edition 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:
 - iSCSI network
 - Public network
 - Private network
 - Cluster virtual serverUse these IP addresses when you install the Windows operating system and 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 connecting initiators (cluster nodes) and the storage system as described in your storage system documentation.

 **NOTE:** For installing and configuring the Dell | EMC AX-series storage system, see the AX150i documentation located at <http://www.emc.com/dell-ax150support>.

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

7 Assign the virtual disks (configured on storage system) to cluster nodes. See "Assigning and Preparing 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 the 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 either added new internal hard drives to your system or you are setting up the RAID configuration for the first time, you must configure the RAID array using the RAID controller's BIOS configuration utility before installing the operating system.

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



NOTE: 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 Server 2003 R2, Enterprise Edition operating system with the latest service pack on each node.
See the *Platform Guide* for more information about the latest supported service pack.
- 4 Ensure that the network adapter drivers installed on each cluster node are the latest supported version.
See the *Platform Guide* for information on required drivers.
- 5 Configure the public and private network adapter interconnects in each node, and place the interconnects on separate IP subnetworks using static IP addresses.
- 6 Configure the network adapters for iSCSI in each node and place the redundant adapters on separate IP subnetworks using static IP addresses.
- 7 Install the latest Software Initiator package on each node.
See the *Platform Guide* for information about supported versions.

The following subsection introduces some principles and procedures for configuring and assigning IP addresses for public, private and iSCSI networks.

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 and iSCSI components 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. |
| 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, Microsoft SQL Server, Enterprise Edition requires at least one static IP address for the virtual server (Microsoft 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 | 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). See "Cabling Your Cluster Hardware" for more information about cluster interconnect options. NOTE: To ensure operation during a DHCP server failure, use static IP addresses. |
| iSCSI network adapters | For iSCSI operation, one or two (redundant paths) network adapter(s) are required. iSCSI Software Initiator uses these adapters to connect to the iSCSI network. |

Configuring IP Addresses for the Private Network

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

If multiple cluster interconnect network adapters are connected to a network switch, ensure that all of the network adapters in the private network are configured with unique IP addresses.

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

 **NOTE:** NIC teaming is supported on the public network, but not on the private and iSCSI networks.

Creating Separate Subnets for the Public, Private and iSCSI 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.

To improve iSCSI performance, use a dedicated gigabit physical network or LAN. Create separate subnets for redundant iSCSI paths to avoid communication problems within iSCSI network, as well as with cluster public and private networks.

An initiator (cluster node) with more than one iSCSI portal connection requires different subnet IP addresses for each port. The iSCSI ports on the storage system are configured accordingly to comply with the topology.

Table 3-2 shows an example of IP address assignments for a two-node cluster with redundant iSCSI NICs.

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

Table 3-2. IP Address Assignments for a Two-Node Cluster

| Usage | Cluster Node 1 | Cluster Node 2 |
|---|---------------------------|---------------------------|
| Public network static IP address (for client, management, and domain controller communications) | 192.168.1.101 | 192.168.1.102 |
| Public network subnet mask | 255.255.255.0 | 255.255.255.0 |
| Default Gateway | 192.168.1.1 | 192.168.1.1 |
| WINS server | Primary 192.168.1.11 | Primary 192.168.1.11 |
| | Secondary 192.168.1.12 | Secondary 192.168.1.12 |
| DNS server | 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 |

Table 3-2. IP Address Assignments for a Two-Node Cluster (continued)

| Usage | Cluster Node 1 | Cluster Node 2 |
|---|----------------|----------------|
| Private network subnet mask | 255.0.0.0 | 255.0.0.0 |
| iSCSI network subnet-1 static IP address | 172.31.1.101 | 172.31.1.102 |
| iSCSI network subnet-1 mask | 255.255.255.0 | 255.255.255.0 |
| iSCSI network subnet-2 static IP address (redundant path) | 172.31.2.101 | 172.31.2.102 |
| iSCSI network subnet-2 mask | 255.255.255.0 | 255.255.255.0 |

NOTE: The default gateway, NetBIOS, WINS and DNS should not be configured on the private network and may not be required on the iSCSI network.

NOTE: The iSCSI network in the above example is a class-B network (172.31.x.x) with two sub-nets of 172.31.1.x and 172.31.2.x

Setting the Network Interface Binding Order for Clusters Running Windows Server 2003 R2, Enterprise Edition

- 1 Click **Start**, 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 and then the **iSCSI** connections.
To change the connection order:
 - a Click **Public**, **Private**, or **iSCSI** connections.
 - 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 Teaming

Dual-port network adapters are supported on public and iSCSI networks, but are not supported in the private network. Network adapter teaming or NIC teaming is only supported 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 by node name (DNS assigned name) as well as by IP address. 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.

Assigning and Preparing the Virtual Disks for Cluster nodes

Before you assign the virtual disks to cluster nodes, ensure that:

- The storage system is properly initialized.
- The iSCSI initiators have established connection with target (storage).
- The cluster nodes are registered with the storage system.
- The virtual disks created in 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 failover individually from node-to-node.

Assigning the Virtual Disks to a Cluster Node

To perform data I/O to the virtual disks, you must assign the virtual disks to a cluster node.

- 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 EMC® PowerPath™ iSCSI on the cluster nodes identifies all paths to the virtual disks.
For more information on PowerPath iSCSI, see your storage system documentation.

Using the Windows Dynamic Disks and Volumes

The Windows Server 2003 R2, Enterprise Edition 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 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.

See Microsoft Knowledge Base Article 237853 located on the Microsoft Support site at support.microsoft.com for more information.

 **NOTE:** The PowerEdge Cluster FE650Wi solution does 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 will see 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 drives functions the same as a normal drive, but is assigned a label or name instead of a drive letter. Using mount points, a cluster can support more shared disks than the number of available drive letters.

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


 **NOTE:** 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 storage system is properly cabled to the servers.
- The iSCSI Initiator target connections are active.
- The virtual disks have been assigned to the hosts.

- PowerPath iSCSI on the server identifies all 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, it is recommended 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.

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 Server 2003 Disk Management or Dell OpenManage™ Array Manager.
- 2** To maintain proper drive letter assignments, ensure the first NIC (for iSCSI access) detected by each node is connected to the first switch or SP-A and the second detected NIC is connected to the second switch or SP-B.
See Figure 4-3 in "Installing Your Cluster in a Direct-Attached Environment" for the location of SP-A and SP-B on the AXI50i 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 with volume label "Volume Z" is assigned to drive letter Z and so on for each disk on each node. Assign drive letters on each of the shared disks, even if the disk displays the drive letter correctly.

Installing and Configuring MSCS

MSCS is an integrated service in Windows Server 2003 R2, Enterprise Edition. 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. See "Using MSCS" for more information.

Verifying Cluster Readiness

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

- All cluster servers are able to log on to the same domain.
- The shared disks are partitioned and formatted, and the same drive letters that reference 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 is used to manage the cluster. Because the Cluster Group is dedicated to cluster management and for best cluster performance, it is recommended that you do not install applications in this group.

Installing the Quorum Resource

When you install a Windows Server 2003 cluster, the installation wizard automatically selects an NTFS disk as the quorum resource for you, 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.

Configuring the Public, Private, and iSCSI Networks

When you install and configure Windows operating system, the software installation wizard prompts you to identify the public and private network segments that are connected to your cluster nodes. For best cluster performance, configure the private network as Private and the client public network segments to All communications. This configuration provides 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 Disable the iSCSI network by unchecking the option **Enable this network for cluster use**.
- 5 For all subsequent NICs, set each NIC to **Client use only** or **All communications**.
- 6 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 R2, Enterprise Edition

When you install and configure a cluster running Windows Server 2003 R2, Enterprise Edition, 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**. Configure at least one network for the cluster interconnect (private network) and one network for all communications.

Configuring Microsoft Cluster Service (MSCS)

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

- 1** Click **Start** and 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** window, 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 **Start** and 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 then 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

To verify that the MSCS service is operating properly:

- 1 Click **Start** 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 for configuring the Cluster Service.

See "Using MSCS" for more information on the Cluster Service.

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 consist 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 **Start** 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 Server 2003 R2, Enterprise Edition for information about troubleshooting resource failures.

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



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.

A direct-attached cluster configuration consists of one or two NICs (for iSCSI access) cabled directly to a Dell | EMC AX150SCi or AX150i storage system. Direct-attach configurations are self-contained and do not share any physical resources with other server or storage systems outside of the cluster.

Figure 4-1 shows an example of a direct-attached, single cluster configuration that is connected to an AX150SCi storage system with one iSCSI NIC installed in each cluster node.

Figure 4-1. Direct-Attached Cluster Configuration Using the AX150SCi Single SP Storage System

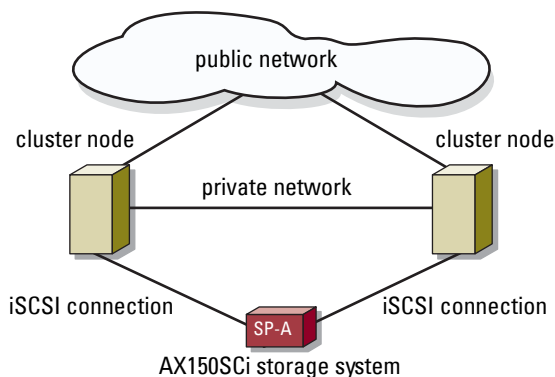



Figure 4-2 shows an example of a direct-attached, single cluster configuration that is connected to an AX150i storage system with two iSCSI NICs installed in each cluster node.

Figure 4-2. Direct-Attached Cluster Configuration Using the AX150i Dual SP Storage System



NOTE: The first three disks (0-2) in AX150SCi and first four disks (0-3) in AX150i storage system are preloaded with storage system software and are used to boot the storage. These disks must remain in their original slots.

Cabling One Cluster to a Dell | EMC Storage System

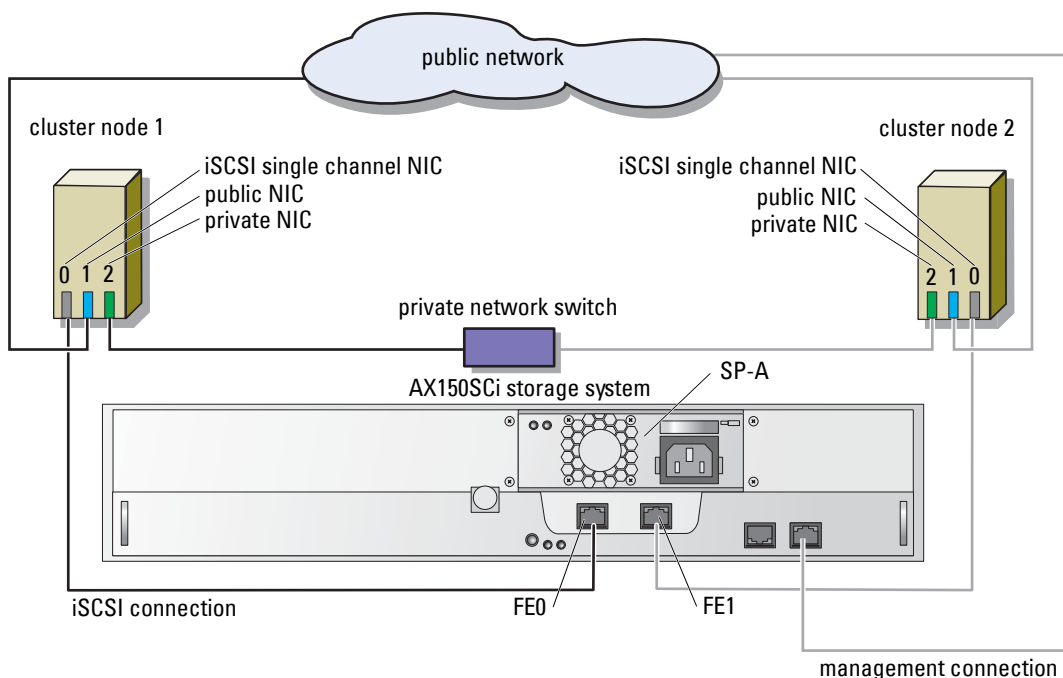
Each cluster node attaches to the storage system using CAT5e or CAT6 LAN cables with RJ45 connectors that attach to Gigabit Ethernet (GbE) NICs in the cluster nodes and the Gigabit iSCSI SP ports in the Dell | EMC storage system.

Cabling a Two-Node Cluster to an AX150SCi Storage System using single Channel Network adapters

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

Figure 4-3 illustrates how to cable a two-node direct-attached cluster to an AX150SCi storage system. Each cluster node is configured with one NIC that is connected to a single SP in the storage system. This cluster configuration does not provide data redundancy.

Figure 4-3. Cabling the Cluster Nodes to the AX150SCi Storage System



The following table shows an example of IP address assignment for a Direct-Attached Two node cluster with iSCSI single channel NICs and an AX150SCi storage system.

Table 4-1. Sample IP Address for Direct-Attached Two Node Cluster With iSCSI Single Channel NICs and AX150SCi

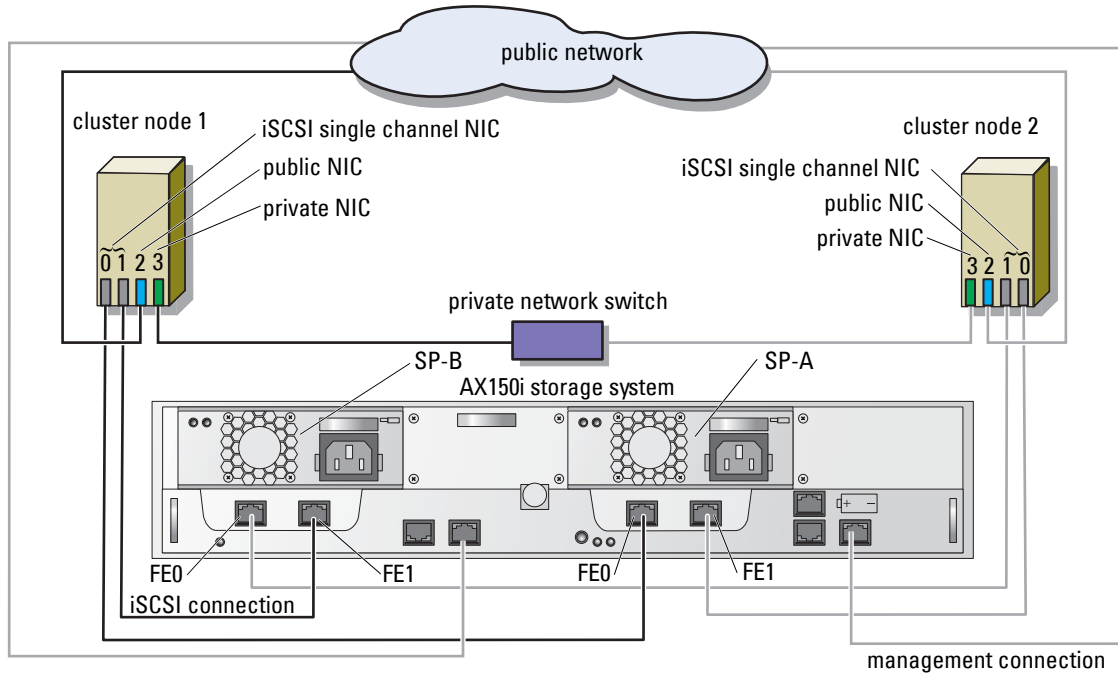
| Port | IP address | Subnet mask | Gateway |
|------------------------------|---------------|---------------|-------------|
| SP-A port FE 0 | 172.31.1.50 | 255.255.255.0 | 0.0.0.0 |
| SP-A port FE 1 | 172.31.1.51 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 iSCSI NIC-0 | 172.31.1.101 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 public NIC-1 | 192.168.1.101 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 1 private NIC-2 | 10.0.0.101 | 255.0.0.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-0 | 172.31.1.102 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 2 public NIC-1 | 192.168.1.102 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 2 private NIC-2 | 10.0.0.102 | 255.0.0.0 | 0.0.0.0 |

Cabling a Two-Node Cluster to an AX150i Storage System using single Channel Network Adapters

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

Figure 4-4 illustrates how to cable a two-node direct-attached cluster to an AX150i storage system. Each cluster node is configured with two single channel NICs 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 4-4. Cabling the Cluster Nodes to the AX150i Storage System using iSCSI Single Channel NICs



The following table shows an example of IP address assignment for a Direct-Attached Two node cluster with iSCSI single channel NICs and AX150i.

Table 4-2. Sample IP Address for Direct-Attached Two Node Cluster With iSCSI Single Channel NICs and AX150i

| Port | IP address | Subnet mask | Gateway |
|------------------------------|---------------|---------------|-------------|
| SP-A port FE 0 | 172.31.1.50 | 255.255.255.0 | 0.0.0.0 |
| SP-A port FE 1 | 172.31.1.51 | 255.255.255.0 | 0.0.0.0 |
| SP-B port FE 0 | 172.31.2.50 | 255.255.255.0 | 0.0.0.0 |
| SP-B port FE 1 | 172.31.2.51 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 iSCSI NIC-0 | 172.31.1.101 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 iSCSI NIC-1 | 172.31.2.101 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 public NIC-2 | 192.168.1.101 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 1 private NIC-3 | 10.0.0.101 | 255.0.0.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-0 | 172.31.1.102 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-1 | 172.31.2.102 | 255.255.255.0 | 0.0.0.0 |

Table 4-2. Sample IP Address for Direct-Attached Two Node Cluster With iSCSI Single Channel NICs and AX150i

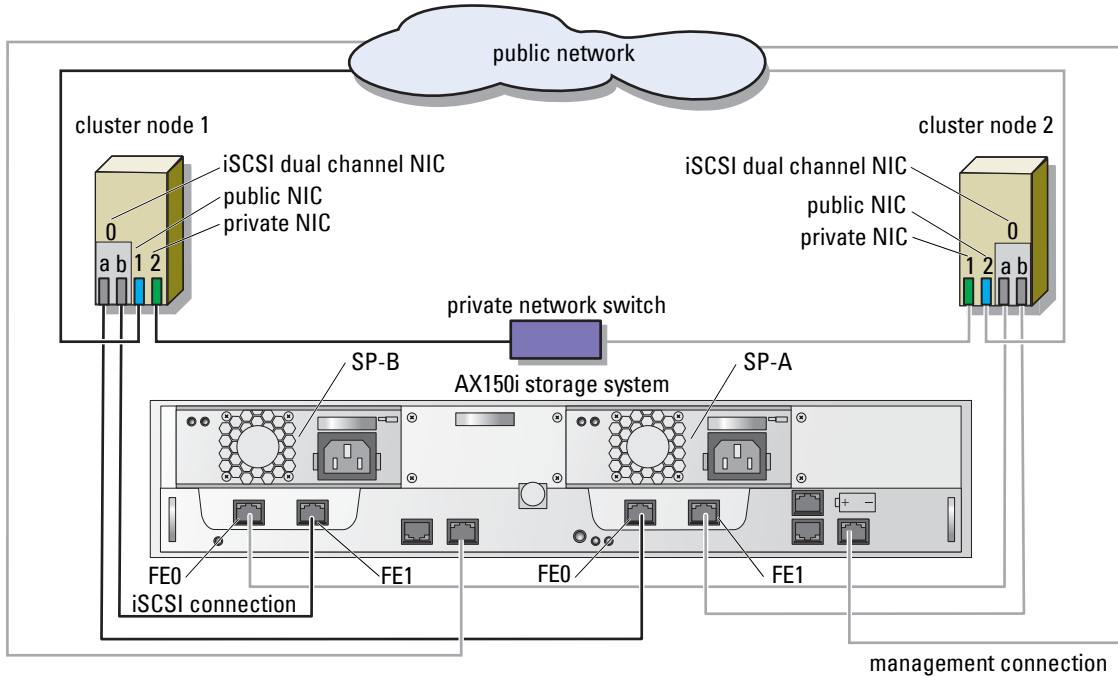
| Port | IP address | Subnet mask | Gateway |
|------------------------------|---------------|---------------|-------------|
| Cluster node 2 public NIC-2 | 192.168.1.102 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 2 private NIC-3 | 10.0.0.102 | 255.0.0.0 | 0.0.0.0 |

Cabling a Two-Node Cluster to an AX150i Storage System Using iSCSI Dual Channel Network Adapters

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

Figure 4-5 illustrates how to cable a two-node direct-attached cluster to an AX150i storage system. Each cluster node is configured with one iSCSI dual channel NIC connected to dual SPs in the storage system. This cluster configuration provides data redundancy and eliminates link failure on the cluster node similar to the previous configuration but it does not protect from hardware failure of the network adapter.

Figure 4-5. Cabling the Cluster Nodes to the AX150i Storage System using iSCSI Dual Channel NIC



The following table shows an example of IP address assignment for a Direct-Attached Two node cluster with iSCSI dual channel NICs and AX150i

Table 4-3. Sample IP Address for Direct-Attached Two Node Cluster With iSCSI Dual Channel NICs and AX150i

| Port | IP address | Subnet mask | Gateway |
|------------------------------|---------------|---------------|-------------|
| SP-A port FE 0 | 172.31.1.50 | 255.255.255.0 | 0.0.0.0 |
| SP-A port FE 1 | 172.31.1.51 | 255.255.255.0 | 0.0.0.0 |
| SP-B port FE 0 | 172.31.2.50 | 255.255.255.0 | 0.0.0.0 |
| SP-B port FE 1 | 172.31.2.51 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 iSCSI NIC-0a | 172.31.1.101 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 iSCSI NIC-0b | 172.31.2.101 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 public NIC-1 | 192.168.1.101 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 1 private NIC-2 | 10.0.0.101 | 255.0.0.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-0a | 172.31.1.102 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-0b | 172.31.2.102 | 255.255.255.0 | 0.0.0.0 |

Table 4-3. Sample IP Address for Direct-Attached Two Node Cluster With iSCSI Dual Channel NICs and AX150i

| Port | IP address | Subnet mask | Gateway |
|------------------------------|-------------------|--------------------|----------------|
| Cluster node 2 public NIC-1 | 192.168.1.102 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 2 private NIC-2 | 10.0.0.102 | 255.0.0.0 | 0.0.0.0 |

Cabling Two Clusters to a Dell | EMC Storage System

The AX150i storage system does not support more than one direct-attached 2-node cluster.

For Direct-Attached cluster cable setup, see "Cabling One Cluster to a Dell | EMC Storage System".

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.

Installing Your Cluster in an iSCSI Switch-Attached Environment

iSCSI Network Overview

An Internet Small Computer System Interface (iSCSI) network provides a low cost, high-performance network that transmits block data between Dell™ PowerEdge™ systems and Dell | EMC storage systems without the associated constraints of direct attached storage. iSCSI networks provide data communications over an existing IP infrastructure while providing additional benefits, such as data storage security, disaster recovery strategies, and minimized interoperability barriers.

Table 5-1 lists the key components in an iSCSI network.

Table 5-1. Key iSCSI Switch-Attached Components

| Component | Description |
|------------------------------------|---|
| Cluster nodes/iSCSI Initiators | PowerEdge systems used as cluster nodes in the cluster. Each system must have an iSCSI Initiator installed to manage the iSCSI data transfer from the node'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/iSCSI target | An iSCSI target completes request from iSCSI Initiator and provides access to external data storage for host systems (cluster nodes). |
| iSNS | An iSCSI Storage Name Service (iSNS) is used for automated discovery and management of iSCSI devices, such as Initiators and targets. |
| Dedicated iSCSI networks (IP SANs) | Native IP SANs consists entirely of components that transmit the SCSI protocol over TCP/IP. Heterogeneous IP SANs transmit SCSI over TCP/IP and FC. PowerEdge Cluster FE650Wi configurations only include native IP SANs. Heterogeneous IP SANs are beyond the scope of this document. |

Table 5-1. Key iSCSI Switch-Attached Components

| Component | Description |
|-----------------------------|--|
| Storage management software | <p>Provides centralized control of the storage system for easier management. Your cluster solution uses EMC[®] storage management software utilities, which include:</p> <ul style="list-style-type: none">• EMC Navisphere™ Express• EMC PowerPath iSCSI• EMC Shared Storage Control (optional)• EMC Snapshot Management (optional) <p>See "Storage Management Software" for more information on the storage management software tools.</p> |

Connecting the Storage Systems to Your Cluster

This section provides information about connecting your cluster to one or more storage systems using an iSCSI network.

iSCSI Switch-Attached Cluster

In an iSCSI switch-attached cluster configuration, all cluster nodes are attached to a single storage system or to multiple storage systems through an iSCSI network using a redundant switch fabric. Switch-attached cluster configurations provide more flexibility, expandability, and performance than direct-attached configurations.

Figure 5-1 shows an example of a two node, an iSCSI switch-attached cluster running Microsoft[®] Windows[®] 2003 R2, Enterprise Edition. Each cluster node is configured with two iSCSI single channel NICs or one iSCSI dual channel NIC that is connected to two Gigabit Ethernet switches. The storage system includes two SPs that provide redundant connections to the Ethernet switches. This configuration provides data redundancy and eliminates a single point of failure on the host or storage system.

Figure 5-1. Two node iSCSI Switch-Attached Cluster





Figure 5-2 shows an example of a eight-node, iSCSI switch-attached cluster running Windows Server 2003, Enterprise Edition.

Figure 5-2. Eight node iSCSI Switch-Attached Cluster



Cabling One iSCSI Switch-Attached Cluster to an AX150i Storage System

The Dell | EMC AX150i storage system is configured with one 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 redundant Ethernet switches and LAN cables with RJ45 connectors.

Each NIC is cabled to a port on a Gigabit Ethernet switch. Two cables connect from the outgoing ports of the Gigabit Ethernet switch to a storage processor on the AX150i storage system.

NOTE: The PowerEdge Cluster FE650Wi solution limits AX150SCi storage system use to a direct attach environment.

Table 5-2 provides information for cabling your storage system to the Gigabit Ethernet switch.


Table 5-2. Storage System Cabling Description

| Storage System | SP Ports | LAN Cables Required | Cabling Description |
|-----------------------|---------------------------------|----------------------------|--|
| AX150i | Two ports per storage processor | 4 | Attach one cable from each storage processor to port to the Gigabit Ethernet switch. |

Cabling an iSCSI Switch-Attached Cluster to a Dell | EMC AX150i Storage System Using iSCSI Single Channel NIC

- 1** Connect cluster node 1 to the iSCSI network. See Figure 5-3.
 - a** Connect a LAN cable from NIC-0 to Ethernet switch 0 (sw0).
 - b** Connect a LAN cable from NIC-1 to Ethernet switch 1 (sw1).
- 2** Repeat step 1 for each cluster node.
- 3** Connect the storage system to the iSCSI network. See Figure 5-3.
 - a** Connect a LAN cable from Ethernet switch 0 (sw0) to SP-A port FE 0.
 - b** Connect a LAN cable from Ethernet switch 0 (sw0) to SP-B port FE 1.
 - c** Connect a LAN cable from Ethernet switch 1 (sw1) to SP-A port FE 1.
 - d** Connect a LAN cable from Ethernet switch 1 (sw1) to SP-B port FE 0.

Figure 5-3. Cabling an iSCSI Switch-Attached Cluster to the Dell | EMC DPE Using iSCSI Single Channel NICs



The following table shows an example of IP address assignment for switch-attached cluster nodes with iSCSI single channel NICs and an AX150i storage system.

Table 5-3. Sample IP Addresses for Switch-Attached Cluster Nodes with iSCSI Single Channel NICs and AX150i

| Port | IP address | Subnet mask | Gateway |
|----------------------------|--------------|---------------|---------|
| SP-A port FE 0 | 172.31.1.50 | 255.255.255.0 | 0.0.0.0 |
| SP-A port FE 1 | 172.31.2.50 | 255.255.255.0 | 0.0.0.0 |
| SP-B port FE 0 | 172.31.1.51 | 255.255.255.0 | 0.0.0.0 |
| SP-B port FE 1 | 172.31.2.51 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 iSCSI NIC-0 | 172.31.1.101 | 255.255.255.0 | 0.0.0.0 |

Table 5-3. Sample IP Addresses for Switch-Attached Cluster Nodes with iSCSI Single Channel NICs and AX150i


| Port | IP address | Subnet mask | Gateway |
|------------------------------|---------------|---------------|-------------|
| Cluster node 1 iSCSI NIC-1 | 172.31.2.101 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 public NIC-2 | 192.168.1.101 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 1 private NIC-3 | 10.0.0.101 | 255.0.0.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-0 | 172.31.1.102 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-1 | 172.31.2.102 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 2 public NIC-2 | 192.168.1.102 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 2 private NIC-3 | 10.0.0.102 | 255.0.0.0 | 0.0.0.0 |

NOTE: A maximum of eight nodes in a switch-attached cluster is supported. Continue the IP address assignments for the remaining nodes, if any.

Cabling an iSCSI Switch-Attached Cluster to AX150i using iSCSI Dual Channel NICs

- 1** Connect cluster node 1 to the iSCSI network. See Figure 5-4.
 - a** Connect a LAN cable from NIC-0a to Ethernet switch 0 (sw0).
 - b** Connect a LAN cable from NIC-0b to Ethernet switch 1 (sw1).
- 2** Repeat step 1 for each cluster node.
- 3** Connect the storage system to the iSCSI network.
 - a** Connect a LAN cable from Ethernet switch 0 (sw0) to SP-A port FE 0.
 - b** Connect a LAN cable from Ethernet switch 0 (sw0) to SP-B port FE 1.
 - c** Connect a LAN cable from Ethernet switch 1 (sw1) to SP-A port FE 1.
 - d** Connect a LAN cable from Ethernet switch 1 (sw1) to SP-B port FE 0.

Figure 5-4. Cabling an iSCSI Switch-Attached Cluster to Dell | EMC AX150i Using iSCSI Dual Channel NICs



The following table shows an example of IP address assignment for switch-attached cluster nodes with iSCSI dual channel NICs and an AX150i storage system.

Table 5-4. Sample IP Addresses for Switch-Attached Cluster Nodes With iSCSI Dual Channel NICs and AX150i

| Port | IP address | Subnet mask | Gateway |
|-----------------------------|--------------|---------------|---------|
| SP-A port FE 0 | 172.31.1.50 | 255.255.255.0 | 0.0.0.0 |
| SP-A port FE 1 | 172.31.2.50 | 255.255.255.0 | 0.0.0.0 |
| SP-B port FE 0 | 172.31.1.51 | 255.255.255.0 | 0.0.0.0 |
| SP-B port FE 1 | 172.31.2.51 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 iSCSI NIC-0a | 172.31.1.101 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 1 iSCSI NIC-0b | 172.31.2.101 | 255.255.255.0 | 0.0.0.0 |

Table 5-4. Sample IP Addresses for Switch-Attached Cluster Nodes With iSCSI Dual Channel NICs and AX150i

| Port | IP address | Subnet mask | Gateway |
|------------------------------|---------------|---------------|-------------|
| Cluster node 1 public NIC-1 | 192.168.1.101 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 1 private NIC-2 | 10.0.0.101 | 255.0.0.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-0a | 172.31.1.102 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 2 iSCSI NIC-0b | 172.31.2.102 | 255.255.255.0 | 0.0.0.0 |
| Cluster node 2 public NIC-1 | 192.168.1.102 | 255.255.255.0 | 192.168.1.1 |
| Cluster node 2 private NIC-2 | 10.0.0.102 | 255.0.0.0 | 0.0.0.0 |


NOTE: There can be a maximum of eight nodes in a switch-attached cluster. Continue the IP address assignment for the rest of the nodes, if any.

Cabling Multiple iSCSI Switch-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 Gigabit Ethernet switch and then connect the switches to the appropriate storage processors on the processor enclosure.


See the *Platform Guide* for rules and guidelines for iSCSI switch-attached clusters.

Cabling Multiple Switch-Attached Clusters to the AX150i Storage System Using Single Channel NICs

 **NOTE:** The following procedure uses Figure 5-3 as an example for cabling additional clusters.

- 1 In the first cluster, connect cluster node 1 to the iSCSI network.
 - a Connect a LAN cable from NIC-0 to Ethernet switch 0 (sw0).
 - b Connect a LAN cable from NIC-1 to Ethernet switch 1 (sw1).
- 2 In the first cluster, repeat step 1 for each node.
- 3 For each additional clusters, repeat step 1 and step 2.
- 4 Connect the storage system to the iSCSI network.
 - a Connect a LAN cable from Ethernet switch 0 (sw0) to SP-A port FE 0.
 - b Connect a LAN cable from Ethernet switch 0 (sw0) to SP-B port FE 1.
 - c Connect a LAN cable from Ethernet switch 1 (sw1) to SP-A port FE 1.
 - d Connect a LAN cable from Ethernet switch 1 (sw1) to SP-B port FE 0.

Cabling Multiple iSCSI Switch-Attached Clusters to AX150i Using Dual Channel NICs

 **NOTE:** The following procedure uses Figure 5-4 as an example for cabling additional clusters.

- 1 In the first cluster, connect cluster node 1 to the iSCSI network. See Figure 5-4
 - a Connect a LAN cable from NIC-0a to Ethernet switch 0 (sw0).
 - b Connect a LAN cable from NIC-0b to Ethernet switch 1 (sw1).

- 2 In the first cluster, repeat step 1 for each cluster node.
- 3 For additional clusters, repeat step 1 and step 2.
- 4 Connect the storage system to the iSCSI network.
 - a Connect a LAN cable from Ethernet switch 0 (sw0) to SP-A port FE 0.
 - b Connect a LAN cable from Ethernet switch 0 (sw0) to SP-B port FE 1.
 - c Connect a LAN cable from Ethernet switch 1 (sw1) to SP-A port FE 1.
 - d Connect a LAN cable from Ethernet switch 1 (sw1) to SP-B port FE 0.

Connecting a Dell 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 support configurations with multiple storage systems attached to clustered servers. In this scenario, the MSCS software can fail over disk drives in any cluster-attached shared storage system 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 R2, Enterprise Edition supports mount points, allowing greater than 22 drives per cluster
See "Assigning Drive Letters and Mount Points" for more information.

Figure 5-5 provides an example of cabling the cluster nodes to four Dell | EMC storage systems.

Figure 5-5. PowerEdge Cluster Nodes Cabled to Two Storage Systems



iSCSI Switch-Attached Cluster Configurations

Certain rules and requirements apply when you configure your iSCSI switch-attached cluster. See the *Platform Guide* for information about supported servers, specific iSCSI initiators and NICs supported by the platform, and PCI slot configuration guidelines.

Table 5-5 provides a list of the documentation you will need to configure an iSCSI switch-attached cluster.

Table 5-5. iSCSI Switch-Attached Cluster Configuration Documentation

| Information | Documentation | Location |
|---|---|---|
| General rules and guidelines for switch-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. | <i>Platform Guide</i> and this document | |
| Using Windows Server 2003 R2, Enterprise Edition for cluster products and components. | This document | |
| Initializing and configuring the EMC® storage system. | EMC storage system documentation | Documentation included with your storage system or at the EMC Support website for AX150i storage system located at www.emc.com/dell-ax150support . |

Table 5-5. iSCSI Switch-Attached Cluster Configuration Documentation (continued)

| Information | Documentation | Location |
|---|---|--|
| Latest firmware and software revision requirements. | EMC CLARiiON® AXI50i Series Support Matrix | EMC Technical Library located at www.emc.com |
| | <i>Platform Guide</i> | Documentation included with your PowerEdge cluster, or at the Dell Support website located at support.dell.com . |



 **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.

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


- 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 Microsoft® 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 ID portion of the new network adapter's IP address is different from that of the first network adapter. For example, if the first network adapter in the node had an address of 192.168.1.101 with a subnet mask of 255.255.255.0, for the second network adapter you might assign the IP address 192.168.2.102 and the subnet mask 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 assignments were configured correctly, all 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, move 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 indicate that the node is back online.

Removing Nodes From Clusters

1 Move all resource groups to another cluster node.

2 Click **Start**, select **Programs**→**Administrative Tools**, and then double-click **Cluster Administrator**.

3 In Cluster Administrator, right-click the icon of the node you want to uninstall and then select **Stop Cluster Service**.

4 In Cluster Administrator, 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 Cluster Administrator.

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 the following command:

```
cd \windows\cluster
```
- 3 Start MSCS in manual mode (on one node only) with no quorum logging by typing the following command:

```
Clussvc -debug -noquorumlogging
```

MSCS starts.
- 4 Run **chkdsk /f** on the disk that is 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 Clusvc`.
The Cluster Service restarts.

See Microsoft Knowledge Base article 258078 located on the Microsoft Support website at support.microsoft.com for more information on recovering from a corrupt quorum disk.

Changing the MSCS Account Password

To change the service account password for all nodes running Windows Server 2003 R2, Enterprise Edition, 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: The Windows Server 2003 R2, Enterprise Edition operating system 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 pane, 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 pane, right-click the physical disk you are reformatting and select **Bring Online**.
- 5 In the **Cluster Administrator** right pane, 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 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 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 the above (public-and-private network).

Preventing Network Failure

When you install MSCS, identify the public and private network segments connected to your cluster nodes. To ensure cluster failover and non-interrupted communications, perform the following procedures:

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

You can set priorities and roles of the networks when you install MSCS or when you use 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 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 7-1 defines states of a node during cluster operation

Table 7-1. Node States and Definitions

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

When 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 7-2 shows resources and their dependencies.


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

Table 7-2. 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 7-3 shows each resource and its configurable parameters.

Table 7-3. 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 |
| IP address | IP address Subnet mask Network parameters for the IP address resource (specify the correct network) |
| Network name | Cluster name or virtual server |

Table 7-3. Resources and Configurable Parameters (continued)

| Resource | Configurable Parameters |
|---------------|---|
| 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 7-4 describes resource dependencies.

Table 7-4. 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. |
| | One of the cables is faulty. | Replace the faulty cable. |
| | The CHAP password entered is wrong. | Enter correct user-name and password for CHAP, if used. |
| 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 <code>ping</code> command from each node to the other node. Try both the host name and IP address when using the <code>ping</code> command. |

Table A-1. General Cluster Troubleshooting (continued)

| Problem | Probable Cause | Corrective Action |
|---|--|--|
| Attempts to connect to a cluster using Cluster Administrator fail. | <p>The Cluster Service has not been started.</p> <p>A cluster has not been formed on the system.</p> <p>The system has just been booted and services are still starting.</p> | <p>Verify that the Cluster Service is running and that a cluster has been formed. Use the Event Viewer and look for the following events logged by the Cluster Service:</p> <p>Microsoft Cluster Service successfully formed a cluster on this node. OR Microsoft Cluster Service successfully joined the cluster.</p> <p>If these events do not appear in Event Viewer, see the Microsoft Cluster Service Administrator's Guide for instructions on setting up the cluster on your system and starting the Cluster Service.</p> |
| You are prompted to configure one network instead of two during MSCS installation. | <p>The TCP/IP configuration is incorrect.</p> <p>The private (point-to-point) network is disconnected.</p> | <p>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.</p> <p>Ensure that all systems are powered on so that the NICs in the private network are available.</p> |
| Using Microsoft® Windows NT® 4.0 to remotely administer a Windows Server 2003 cluster generates error messages. | Normal. Some resources in Windows® Server 2003 R2, Enterprise Edition are not supported in Windows NT 4.0. | Dell strongly recommends that you use Windows Server 2003 for remote administration of a cluster running Windows Server 2003 R2, Enterprise Edition. |

Table A-1. General Cluster Troubleshooting (continued)

| Problem | Probable Cause | Corrective Action |
|---|--|--|
| 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 the CHAP user-name and password, if used.• Use the "Advanced" with "Minimum" option |
| 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, 14

C

cable configurations
cluster interconnect, 28
for client networks, 28
for mouse, keyboard, and
monitor, 25
for power supplies, 25

chkdsk/f
running, 69

cluster
configurations, 12
forming a new cluster, 74
joining an existing cluster, 75
verifying functionality, 45
verifying readiness, 41
verifying resource
availability, 45

cluster configurations
active/active, 14
active/passive, 14
connecting to multiple shared
storage systems, 64
connecting to one shared
storage system, 56
direct-attached, 12, 47
switch-attached, 14, 65

cluster group

installing applications, 42
cluster networks
configuring Windows 2000
cluster networks, 42
configuring Windows Server
2003 cluster networks, 42

cluster nodes
about, 74
states and definitions, 74

cluster objects
about, 73

cluster resources
configurable parameters, 76
resource dependencies, 78
resource failure, 77
setting resource properties, 75

cluster storage
requirements, 22

clustering
overview, 7

D

Dell | EMC AX100
about, 9
cabling the cluster nodes, 49
cabling the cluster nodes in a
switch-attached
environment, 58
cabling to one cluster, 49
cabling to one switch-attached
cluster, 58

cabling to two clusters, 54
direct-attached cluster
about, 47
cabling, 49

domain model
selecting, 32

drive letters
assigning to shared storage
storage systems, 38

dynamic disks
using, 37

F

failback
about, 80

failover
configuring, 78
modifying failover policy, 80
policies, 16

failover configurations
for Windows Server 2003,
Enterprise Edition, 16

failover policies, 16
failover pair, 19
failover ring, 20
for Windows Server 2003,
Enterprise Edition, 16
multiway failover, 20
N+I failover, 17

file share resource type, 79

G

group affinity
about, 18
configuring, 18

H

high availability
about, 7

I

IP address
assigning to cluster resources
and components, 34

K

keyboard
cabling, 25

M

monitor
cabling, 25

mouse
cabling, 25

MSCS
installing and configuring, 41
verifying operation, 44

multiway failover, 20

N

N+I failover

configuring group affinity, 17

Navisphere Express
about, 11

Navisphere Server Utility
about, 10

Navisphere Storage System
Initialization Utility
about, 10

network adapters
cabling the private network, 28
cabling the public network, 28
using dual-port for the private
network, 29, 36

network failure
preventing, 73

network interfaces, 74

NIC
about, 10

O

operating system
installing, 31, 33

P

period values
adjusting, 77

power supplies
cabling, 25

PowerPath
about, 10

private network
cabling, 27-28
configuring IP addresses, 34

creating separate subnets, 35
hardware components, 29
hardware components and
connections, 28
using dual-port network
adapters, 36

public network
cabling, 27
creating separate subnets, 35

Q

quorum resource
about, 7, 77
creating a virtual disk, 40
definition, 7
installing, 42
running chkdisk, 69

R

resource
creating, 78
deleting, 79
resource dependencies, 76, 78
resource groups, 7
definition, 7
resource properties, 76

S

shared storage
assigning drive letters, 38
naming and formatting
drives, 38

Snapshot Management

- about, 11
- SnapView
 - about, 8
- software
 - for storage management, 10
- storage management software
 - Navisphere Express, 11
 - Navisphere Server Utility, 10
 - Navisphere Storage System Initialization Utility, 10
 - PowerPath, 10
 - Snapshot Management, 11
- storage system
 - configuring drives on multiple shared storage systems, 40
 - connecting to your cluster, 47, 56
 - using dynamic disks and volumes, 37
- subnets
 - creating, 35
- switch-attached cluster
 - about, 56
 - configurations, 56
- switches
 - about, 10
- system requirements, 21

T

- threshold
 - adjusting, 77
- troubleshooting
 - connecting to a cluster, 84
 - failed resources, 45
 - shared storage subsystem, 83

V

- virtual servers, 7
 - definition, 7
- volumes
 - using, 37

W

- warranty, 23
- Windows 2000 Advanced Server
 - installing, 31, 66
- Windows Server 2003, Enterprise Edition
 - cluster configurations, 17, 19-20
 - installing, 31, 66

