

Course Overview

Classroom Facilities	1
Classroom Guidelines.....	1
Prerequisites	2
Prerequisite Certifications.....	2
Prerequisite Training.....	2
Course Synopsis	2
Objectives	3

Module 1 — Clustering Concepts

Objectives	1
The Need for Availability.....	2
Critical Application Classes.....	3
Availability Levels	4
Achieving High Availability	6
What Is a Cluster?	8
Cluster-Aware Applications.....	9
Cluster Models	10
Advantages of Clustering	11
High Availability of Resources.....	11
Scalability for Growth.....	12
Centralized Administration	12
Load Balancing	13
Planning for Storage Area Networks.....	14
Determining the Need for Clusters	14
Decision Points	15
Compaq Cluster Solutions.....	16
Learning Check	17

Module 2 — Hardware Planning

Objectives	1
Planning a Cluster	2
Planning a Server	3
Choosing a Server	4
Clustering a Server	5
Planning Shared Storage Subsystems	6
Storage Hardware Overview	6
Shared Storage Capacity Planning	7
Shared Storage Requirements	8
Shared Storage Subsystem Family	9
Shared Storage Subsystem Features	13
Shared SCSI Storage Subsystems	14
Planning SCSI IDs	14
Nonshared Disk Drives	15
Planning for Storage Area Networks	16
Multiple FC-AL Configuration	18
Redundant FC-AL Configuration	19
Fibre Channel Host Bus Adapters	20
Planning NCS Intracluster Connections	30
Planning Cluster Backup and Restore	31
Compaq StorageWorks Backup Sizing Tool	32
Enterprise Backup Components	33
Cluster Backup and Restore Design Considerations	36
Server-Based (Local) Backup	36
LAN-Based (Remote) Backup	36
SAN-Based Backup	37
Failure During Backup	37
Cluster Restoration	38
NetWare 6 Backup Simplification	38
NSS Mirroring with Novell Cluster Services	39
Learning Check	40

Module 3 — Compaq ProLiant Cluster Solutions for NetWare 6

Objectives	1
Overview	2
Features	3
Components	4
Compaq ProLiant DL380 G2 Packaged Cluster	6
Features	7
Server Integrated Smart Array 5i Controller.....	9
Shared SCSI Block.....	10
High-Availability Fibre Channel Solutions.....	11
ProLiant Cluster HA/N100	11
NetWare Cluster Basic Documentation Kit.....	11
ProLiant Cluster HA/N200	12
NetWare Cluster Redundant Path Kit.....	13
RA4100 SAN Solution	13
ProLiant Cluster HA/N500	14
NetWare Cluster Redundant Path Kit.....	15
Features Comparison.....	15
Software Components	16
From Novell	16
Novell NetWare 6 Operating System Software.....	16
From Compaq	17
Array Configuration Utility	17
CLI and SWCC.....	18
ACS.....	19
Learning Check	21

Module 4 — Shared Storage Systems for Novell Cluster Services

Objectives	1
Overview	2
Configuring Shared Storage	3
Creating the Cluster Services Partition	4
Configuring the Storage System	5
Applying Power	5
Configuring for High Availability	6
Identifying and Eliminating Single Points of Failure	6
Configuring Storage Arrays for High Availability	8
ProLiant DL380 G2 Packaged Cluster	8
Planning Highly Available Cluster Configurations for Fibre Channel	9
Redundant Controller Option	10
HBA Configurations Supported with the RA4100	11
SANworks Secure Path for NetWare	12
Learning Check	21

Module 5 — Network Configuration

Objectives	1
Network Adapters Supported	2
Configuring Communication Links for High Availability	3
Compaq Advanced Network Services	4
Network Adapter Teaming	5
Advantages of Teaming	6
Dual-Port NICs	7
Types of Teaming	7
Adapter Failover	8
Redundant Hubs and Switches	9
Supported Transport Protocols	10
TCP/IP	10
IP Planning	10
Network Capacity	11
Learning Check	12

Module 6 — Novell Cluster Services

Objectives	1
Novell Cluster Services	2
Typical Cluster Configuration.....	4
Master Node	5
Cluster Heartbeat.....	5
Failover and Failback.....	6
Fan-Out Failover™	7
Cluster Resources	8
Resource Objects.....	9
Load Scripts	11
Unload Scripts	13
Policies.....	15
Nodes	17
Resource Templates	18
Application Resources	19
Volume Resources	19
Cluster-Enabled Volumes.....	20
Regular Volumes	24
Cluster Nodes	25
Number + IP Address	25
NCP Server	25
Resource Priority	26
eDirectory Cluster Objects.....	27
ConsoleOne Properties.....	28
Quorum Membership.....	29
Protocol Settings.....	31
Cluster Management Port	34
System Architecture	35
Cluster Configuration Library.....	35
Group Interprocess Protocol	36
Cluster Resource Manager	37
Cluster Management Agent.....	39
Portal Cluster.....	39
Cluster Membership Monitor.....	40
Split Brain Detector.....	41
Restoring an SBD Partition	42
Globally Unique Identifiers.....	43
VIPL Extensions	44
Cluster System Services.....	45
Cluster Volume Broker	45

Installing or Upgrading NCS	46
Creating Storage Pools and Volumes	48
Hardware and Software Requirements	50
Upgrading from NetWare 5.1 to NetWare 6	52
Client Configuration Parameters	53
Learning Check	55

Module 7 — Cluster Monitoring and Management

Objectives	1
Monitoring Novell Cluster Solutions	2
Compaq Insight Manager 7	2
Cluster Monitor	4
Settings Menu	4
User Settings	5
Cluster Settings	6
Compaq Intelligent Services Link	6
Management Using ConsoleOne	7
Cluster State View	10
Identifying Cluster States	11
Console View	12
Partition and Replica View	13
Migrating Resources	14
NCS Command Prompt Commands	15
NetWare Remote Manager	16
Tuning for Improved Cluster Performance	17
Managing Network Clients	17
Load Balancing	17
Managing a Cluster Without Interrupting Cluster Services	18
Managing a Cluster in a Degraded Condition	18
Learning Check	19

Module 8 — Application Deployment

Objectives	1
Cluster- and Noncluster-Aware Applications	2
Cluster-Aware	2
Deploying DHCP with NCS	3
Deploying Novell GroupWise 6 with NCS	4
Noncluster-Aware	7
Application Sizing with Compaq ActiveAnswers	8
Application Pairing	8
Compaq Application Partner Software	9
ISV Partner Program	9
Directly Connected Devices	10
Learning Check	11

Module 9 — Cluster Maintenance and Troubleshooting

Objectives	1
Adding and Troubleshooting Shared Storage in an NCS Cluster	2
Adding Shared Storage to the RA4100	2
Installing the Hardware	2
Troubleshooting the RA4100 Shared Storage	3
Using FC-AL or FC-SW Connectivity	3
Adding Shared Storage to the MA8000 and EMA12000	4
Verifying Configuration	5
Verifying Details	6
Troubleshooting the MA8000 and EMA12000 Shared Storage	6
Verifying Connectivity to a Redundant FC-AL or FC-SW	6
Using FC-AL or FC-SW Connectivity	7
Connecting the MA8000 or EMA12000 to FC-AL or FC-SW Connection Paths	7
Connecting Server Nodes to FC-AL or FC-SW Connection Paths	8
Determining Why a Node Cannot Connect to the Shared Drives	9
Using Common Fibre Channel Shared Storage Troubleshooting Techniques	10
Verifying the Hardware	10
Verifying Interconnects	14
Adding or Replacing a Cluster Node	15
Installing NetWare on a Server to Be Added to an Existing Cluster	15
Adding a Node to a Cluster to Which it Had Been a Member	16
NetWare Troubleshooting Tools and Tips	17
For the Operating System	17
For NCS	17
For Licenses	18
Tools	19
Tips	19

Compaq Troubleshooting Tools and Techniques	20
Using Compaq QuickFind 2000	21
Using Fibre Channel	22
Using the Compaq Fibre Channel Fault Isolation Utility	22
Using SANworks Secure Path for NetWare	25
Using Device Drivers	28
Using NLMs	30
Comparing Software Components	31
Shared SCSI Solutions	31
Fibre Channel Solutions	31
Learning Check	32

Learning Check Answers

Appendices

Appendix A — Compaq ProLiant CL380 Packaged Cluster

Appendix B — Resources

Classroom Facilities

The instructor will provide detailed information concerning:

- Location of restrooms
- Class hours
 - Class start time
 - Scheduled breaks
 - Class stop time

Classroom Guidelines

Class attendees must adhere to the following guidelines:

- Do not interfere with other students' learning.
 - Be on time for class.
 - Turn all mobile phones and pagers to *off* or a silent setting.
 - Be professional in speech and actions.
 - Do not change or modify lab equipment passwords and software configuration.
- Do not smoke in the classroom.



Important

Class attendees may be removed from the classroom and not allowed to return if they fail to follow the classroom guidelines.

Prerequisites

This course is a Compaq Master Accredited Systems Engineer (MASE) level course. It is designed for people who have the required certifications or equivalent knowledge and experience.

Prerequisite Certifications

The following certifications are required prior to taking this class:

- Certified NetWare Engineer (CNE) for NetWare 5 or NetWare 6
- Compaq Accredited Systems Engineer (ASE) NetWare Specialist

Prerequisite Training

In addition to these certifications, the student is required to have completed the StorageWorks Full-Line Technical Web-Based Training (WBT) or have the equivalent knowledge and experience.



Important

This course builds upon knowledge gained in the prerequisite certifications and training. If a student does not meet the prerequisites, this course can be extremely difficult or impossible to complete. The course is written, and will be taught, as if all students have met the prerequisites.

Course Synopsis

This course provides information about Compaq ProLiant Cluster Solutions for NetWare 6, including:

- Planning
- Installing
- Configuring
- Managing
- Maintenance
- Troubleshooting

Objectives

Upon completion of this course, you should be able to:

- Explain clustering concepts.
- Describe hardware planning.
- Identify Compaq ProLiant cluster components.
- Describe shared storage systems for Novell Cluster Services (NCS).
- Describe network configuration guidelines for NCS.
- Describe the functions, installation, and configuration of NCS.
- Explain how to manage and administer NCS.
- Describe the deployment of applications with NCS.
- Discuss cluster maintenance and troubleshooting procedures for Compaq ProLiant Cluster Services for NetWare 6.

Objectives

After completing this module, you should be able to:

- Discuss the need for availability and how it affects client/server applications.
- List the four availability levels and the primary causes of downtime.
- List the methods of Intelligent Fault Resilience that help to achieve high availability.
- Define clusters, the two critical application classes, and the three cluster models.
- List the advantages of clustering.
- Discuss the factors involved in choosing a high-availability solution.
- Describe the cluster solutions available from Compaq.

The Need for Availability

Availability is the measure of how well a computer system can deliver services continuously, and is expressed as a percentage. It is measured against a period of one year. The best possible rating, 100%, indicates that the system is available all the time.

The rapid growth of the Internet as a medium for conducting business has spurred the development of the high-availability solutions industry. Online shopping, for example, often involves order and inventory databases available to the online consumer 24 hours a day, 365 days a year. Even a brief period of downtime can result in significant losses and customer dissatisfaction.

Example

A company has an e-commerce site that earns daily revenues of one million dollars, or an average of more than \$40,000 per hour. Assume that 80% of the transactions that are attempted during downtime are completed later, and 20% of the transactions are lost to competitors. This equates to more than \$8,000 of losses for each hour of downtime.

A company that makes \$30 million per day in e-commerce revenue will lose \$250,000 for each hour of downtime.

System availability must be increased to prevent losses. Achieving these higher levels of availability can be expensive, but it is easy to justify the extra cost when compared to the losses that can occur when the system is not available.



Note

There are several different types of clusters. Some are focused on performance while others are focused on high availability. The cluster concepts discussed in this module are focused on high availability.

Critical Application Classes

Critical applications can be classified as either:

- **Mission-critical** — These applications require 100% uptime. If one of these systems becomes unavailable something tragic can happen. Examples of mission-critical applications include:
 - Air traffic control
 - 911 emergency call center systems
 - Stock exchange trading floors
 - Aerospace mission control
- **Business-critical** — These applications can tolerate minimal interruptions. A business-critical application is extremely important to a company, but can have small amounts of downtime. These applications include:
 - Electronic transfers in banking
 - Company payroll systems
 - Human resources systems
 - Workgroup applications
 - Reservation systems
 - Cash machine systems
 - Transportation logistics systems
 - Messaging (email) systems
 - E-commerce sites

Example

It is important for the company payroll system to be available when it is time to pay the employees. However, if the system is down for a few minutes or hours on a non-payday, it will probably not have as large an impact.

Availability Levels

Four availability levels (AL) have been defined:

- AL1
 - Work stops.
 - Uncontrolled shutdown results.
 - Data integrity is ensured.
- AL2
 - User is interrupted, but can log on again.
 - User might have to rerun some transactions from a journal file.
 - Performance degradation can result.
- AL3
 - User stays online.
 - Current transaction might need restarting.
 - Possible performance degradation can result.
- AL4
 - Process is transparent to user.
 - No work is interrupted.
 - No transactions are lost.
 - No performance degradation results.

Stand-alone Compaq ProLiant servers are rated AL1. If the server is down, all work being performed by that server stops. ProLiant clusters can achieve AL2 or AL3. The user may notice a small interruption or degraded performance. Compaq also has systems at AL4, including the Compaq Tandem Himalaya NonStop Integrity systems.

By using redundant components and mechanisms that detect and recover from faults, clusters increase the availability of applications critical to business operations. In addition to the system hardware, the operating system and application software must be designed for availability.


Environmental factors such as air conditioning and power must be taken into account and configured to provide availability. Procedures, such as change control and disaster-recovery procedures, must be designed and followed to ensure that business practices do not negatively impact availability.

Causes of Downtime

Downtime, measured by the amount of time the system is unavailable, has several primary causes:

- **Hardware faults** — These can be caused by faulty or damaged equipment.
- **Software faults** — Software problems are often a result of interactions between the operating system and third-party applications or drivers.
- **Planned service events** — Upgrading hardware, such as adding memory, or upgrading software, such as installing a patch, can require the system to be shut down or restarted.
- **Operator errors** — Examples include:
 - Misconfiguration
 - Shutting down the wrong server
 - Omitting a procedural step
- **Environmental factors** — Natural disasters, air conditioning, and power grids are examples of environmental factors that can affect availability.
- **Security breaches** — An unauthorized user could break into a system and delete files, change configurations, or shut down the system and make it unavailable.

Achieving High Availability



Overcoming downtime requires thorough fault management. The key to achieving overall system high availability is to ensure that each server is maximized for high availability. Compaq ProLiant servers and storage systems offer a comprehensive set of features that contribute to overall system availability.

The goal is to prevent failures from occurring. If that is not possible, then provide fault tolerance mechanisms, such as redundant components, to automatically take over if a component failure were to occur.

If the system does fail, rapid recovery mechanisms bring the system back online as soon as possible.

Intelligent Fault Resilience, part of the Compaq Adaptive Infrastructure, is an approach to predict, diagnose, and rapidly respond to potential and actual fault conditions.



Note

Specific features may fall under multiple categories within the Adaptive Infrastructure framework.


Examples of the implementation of Intelligent Fault Resilience are:

- Server Fault Resilience and Predictive Management
 - Compaq Insight Manager 7
 - Prefailure warranty
 - Uninterruptible power supplies (UPSs)
 - Hot-plug components
 - ◆ Disks
 - ◆ Power supplies
 - ◆ Fans
 - ◆ PCI I/O cards
 - ◆ Memory
- Intelligent Fault Avoidance and Recovery
 - High-availability clustering
 - Data Replication Manager (DRM)
 - Automatic Server Recovery (ASR)
 - Tape backup
 - Remote Insight Lights-Out Edition (RILOE) technology
- Self-Healing Architectures
 - Online spare disk drives and memory
 - Error control and correcting (ECC) and RAID memory
 - RAID 1, 1+0, 4, 5, and Advanced Data Guarding (ADG)
 - Battery-backed cache

**Note**

High-end cluster solutions such as OpenVMS clusters are fault resilient, but entry-level and mid-range clusters are used for Intelligent Fault Recovery.

What Is a Cluster?



Basic High-Availability Cluster Configuration

A cluster is a set of loosely coupled servers used as a single, unified computing resource. This type of cluster contains four basic elements:

- Servers, also known as cluster nodes
- Inter-node communication paths, called interconnects or the private network
- Shared physical storage
- Cluster-enabling software

Clustering should be implemented only after fault-resilient and self-healing architectures have been implemented.

Clusters keep server-based applications highly available, regardless of individual component failures. They also ensure minimal interruption of operations and can increase processing capacity and I/O bandwidth.

The services running on any system in the cluster are available to all connected network users. A client interacts with a cluster as though it were a single server and therefore might not even know that it is a cluster.

A cluster offers a level of scalability and availability that exceeds that of stand-alone servers. To users, this can translate into increased performance and reliability. High availability is obtained by reducing single points of failure in the cluster. A single-point-of-failure is any one item whose failure would cause the service or application to become unavailable.

Example

One single-point-of-failure could be a device that is directly connected to only one node of a cluster, such as a printer or modem bank. If that node fails, the device would no longer be available.

Clustering traditionally has existed only in the realm of proprietary systems designed for mission-critical applications such as stock trading and aerospace mission control. Clustering technology has entered mainstream, industry-standard computing because of the increasing demand to keep business-critical applications available.

These industry-standard clustering solutions:

- Use industry-standard hardware and software.
- Provide essential clustering features and benefits at a price lower than proprietary clustering systems.
- Increase the usefulness and life span of software applications.

Cluster-Aware Applications

A cluster-aware application recognizes that it is being installed on a cluster and creates the necessary resources that it requires for clustering. This makes clustering the application easier than if the application is noncluster-aware. A noncluster-aware application must be manually configured for failover. In addition to being simple to install, a cluster-aware application is easier to manage and generally can recover from more faults than a noncluster-aware application.

Cluster Models

The three basic cluster models are:

- **Shared-nothing** — Nodes in this model have access to a shared physical storage system, but the nodes cannot access the same logical drives at the same time.
- **Shared-disk** — A distributed lock manager is used to allow more than one node to access the same logical disks at the same time.
- **Shared-everything** — The cluster nodes share memory, processors, and disks.



Note

A shared-everything cluster is also referred to as a single system image cluster. The administrator, clients, and operating system see the system as one unit.

Advantages of Clustering

Clustering provides the following advantages:

- High availability of resources
- Scalability for growth
- Centralized administration
- Load balancing

High Availability of Resources

Clustering ensures the high availability of resources. If one node in a cluster either fails or the administrator takes it offline, the resources can fail over to a surviving node. During failover the remaining servers automatically redistribute the tasks of the node that is down or failed. The process is transparent to clients in the network.

When a node that failed or was taken offline is ready, the resources can fail back to the original node. This process can be done manually, or it can be configured to happen automatically.

Scalability for Growth

In addition to providing high availability, clusters can be highly scalable. Through clustering, the following resources can be incrementally and efficiently expanded:

- Processor
- I/O
- Storage
- Applications

Clustering provides reliable access to system resources and data as well as investment protection for hardware and software resources.

Clusters are also scalable in terms of investment. Protect investments in both hardware and software by clustering existing hardware with new computers. Instead of replacing a stand-alone computer with a new one of twice the capacity, add another computer of equal capacity.

Example

The number of clients using an application on a server increases and performance degrades. The application is written so that the processing can be split across multiple nodes. Clustering this server with another server will improve performance and increase availability of the application.

Centralized Administration

In a typical server environment, various administrative tools identify the servers on the network and monitor their contents and activities. However, in the cluster environment, the administration of applications and services is centralized.

Load Balancing

Clusters not only provide high availability, scalability, and centralized administration, but also can provide load balancing.

Example

A system administrator discovers that too many applications are running on one server while another server is barely being used. Ordinarily, the system administrator would have to shut down and reconfigure both systems.

However, if the servers are part of a cluster group with application failover capabilities, the system administrator can manually fail over the applications to another server and balance the workload without shutting down the server.

During normal operation, most clustering software allows manual or automatic load balancing of resources. Novell Cluster Services support only manual load balancing.

Environments that are not resource-intensive (web servers) are perfect for server farms (multiple, independent servers). Each server can handle simultaneous connections. These servers distribute incoming requests using a load balancer, which can be either hardware or software.


Hardware load balancers have the following benefits:

- Greater scalability than software load balancers
- Easier to manage
- No processor load on the application server farms

Software load balancers have the following characteristics:

- Up to 32 systems can be included in a load-balancing cluster.
- All servers must be on the same network segment.
- Load balancing is performed on the same server as the application.
- No external hardware is needed, so the cost is lower than hardware load balancer.

Planning for Storage Area Networks



Selecting the appropriate high-availability solution is important. There is a direct relationship between the degree of availability and the cost of the solution. The greater the degree of availability required, the higher the cost of the solution needed to fulfill those requirements. Therefore, it is important to understand environments and the incremental cost of adding levels of availability.

Determining the Need for Clusters

Research shows that the leading causes of downtime can be ranked in the following order:

1. Infrastructure problems (building, power, and network)
2. Software failures (operating system, application, tools, and drivers)
3. Operational and administrative activities (procedures, personnel, and maintenance activities)
4. Hardware failures (disk, power supply, and memory)

The least likely component to fail in an environment is the hardware. Downtime usually is related to events and activities outside the server. Issues such as poorly trained personnel, building outages (power or air conditioning), or flawed backup and restore procedures typically account for more downtime than events such as hard drive failures.

Decision Points

Each organization must determine its leading causes of downtime before deploying a cluster solution. The following decision points will assist in determining whether clustering is right for the organization.

What Are the Availability Requirements?

Many organizations can tolerate hours or even days of downtime for their environment. Others need 99.999% availability. Current technology might not meet the requirements of some organizations.

Can Availability Requirements Be Achieved by Investment in Other Areas?

Availability of 99% and 99.9% can be attained without additional investment in clustering technology. Compaq provides many hardware and off-the-shelf technologies, such as RAID disk arrays and redundant power supplies and fans, that can increase availability without clustering. Personnel training or procedural redesign also can alleviate problems in the areas that are causing downtime.

Does Clustering Solve the Leading Causes of Downtime?

After evaluating the leading causes of downtime in an environment, determine whether clustering would address these issues.

Clustering support includes protection from hardware failures and increased availability. If the organization's leading downtime causes are operational or administrative issues, software failures, or infrastructure problems, clustering technology might not reduce downtime. Because hardware failures are the least frequent cause of downtime, implementing clustering as a means of protection from hardware failures cause disappointment.

Can the Environment Tolerate the Increased Complexity that Clustering Introduces?

To administer a cluster, be familiar with clustering concepts and cluster administration tools. Procedures for failover and failback operations must also be developed.

Some applications are very complex; clustering those applications increases the total complexity drastically.

Is the Investment in Clustering Justified by the Return?

If the environment is mission-critical, it might warrant investment in every tool available to increase availability. Decide whether investments in other areas make more sense or if the additional availability achieved through clustering is justified.

Compaq Cluster Solutions

Compaq has offered high-availability solutions across all platforms since 1995. Early solutions increased data availability for multiple Compaq servers and allowed the servers to recover from hardware errors. The latest products are designed to meet customer requirements for industry-leading clustered solutions.

Compaq ProLiant clusters support the following platforms:

- Microsoft Windows 2000 Cluster service
- Novell Cluster Services
- LifeKeeper for Linux
- UnixWare NonStop Clusters
- Compaq Parallel Database Cluster (PDC) for Oracle *9i* Real Application Cluster (RAC)

In addition, Compaq also offers the following high-end cluster solutions that run on Compaq AlphaServers or Himalaya systems:

- Tru64 UNIX TruCluster Server
- OpenVMS Clusters
- Tandem Himalaya NonStop Integrity systems

This course focuses on ProLiant clusters that use the Novell Cluster Services 1.6 service. The following Compaq courses provide more information on other ProLiant cluster solutions:

- Implementing Windows 2000 on Compaq ProLiant Clusters
- LifeKeeper for Linux
- Compaq High Availability Full Line Technical WBT

INTERNET

Find technical information on all industry-standard, high-availability solutions at:

<http://www.compaq.com/highavailability>

Learning Check

1. Which availability level (AL) allows users to stay online although current transactions might need restarting and there is a possibility of performance degradation?
.....
2. Name the two classes of critical applications.
.....
.....
3. What methods of high availability should be attempted first before clustering should be implemented?
.....
4. Name the three basic cluster models.
.....
.....
.....
5. A company has two servers. One server is busy running all the applications that the company uses for business. The other server has only a few, rarely used applications. If the servers are clustered with application failover capabilities, the system administrator can manually fail over part of the applications to the other server to redistribute the workload. This is an example of what type of cluster advantage?
.....

6. Which component is the least likely to fail in a network environment?

.....

.....

.....

7. Name three platforms supported by Compaq ProLiant Clusters.

.....

.....

.....

8. Where is technical information found on all industry-standard, high-availability solutions for Compaq clusters?

.....

Objectives

After completing this module, you should be able to:

- List hardware requirements for Novell Cluster Services (NCS).
- Plan NCS intraccluster connections.
- Plan an NCS cluster backup and restore strategy.

Planning a Cluster

When planning an NCS cluster, consider the types and configurations of the servers, the storage system, and the interconnection between servers used in the cluster. In many cases, hardware decisions are driven primarily by the existing hardware. Other decisions depend on the specific uses of the cluster.

Planning begins by taking an inventory of what hardware currently exists and identifying the new equipment required for the cluster configuration.

Consider the following topics when planning a cluster:

- Type of business using the cluster
- Business function provided by the cluster and the impact on the business of any individual server failing
- Number of users accessing the system and whether this number will increase or decrease in the future
- Storage capacity required
- Physical layout of the cluster
- Percentage of downtime users can accept and still perform core business functions
- Primary tasks of the servers, such as administration, storage, or messaging
- Aggregate resource requirements for the servers
- Workload failover requirements (can the user separate the workload that relies on failover from the workload that can tolerate failure) and load balancing capabilities
- Enterprise backup and restore options



Note

If clustering is to be implemented with existing hardware, perform a full backup of the data **before** making changes to the system.

Planning a Server

Servers are the primary components of any cluster. Compaq ProLiant servers were used during the development of NCS, and the partnership between Novell and Compaq has ensured that the Compaq servers certified for NCS satisfy all NetWare clustering solution requirements.


During the planning stage of a cluster solution, decisions must be based on the identification and prioritization of business-critical applications.

To choose the appropriate servers for clustering solutions, consider the following criteria:

- Business function of the servers and whether they will be used as file/print servers, database servers, web servers, or for a combination of uses
- Impact on the business if server functionality is lost
- Percentage of downtime users can accept and still adequately perform core business functions
- Required uptime for the server and whether this is greater than what is currently provided
- Number of users that need support and the predicted increase of this number
- Expansion requirements for future business growth

These considerations determine requirements such as the speed, quantity, and size of servers needed.

Choosing a Server




The following ProLiant servers are currently certified to support the clustering solutions available for NCS:

- **ProLiant DL380 Generation 2 (G2) and the ProLiant DL380 G2 Packaged Cluster**
 - For Internet services, corporate data centers, or remote sites
 - For multipurpose server environments running business-critical applications to eliminate the need to balance I/O
 - For 7 x 24 multiserver environments
- **ProLiant DL530, ProLiant DL580, and ProLiant DL760** — For data centers using external storage and backup solutions
- **ProLiant ML570 and ProLiant ML750** — For entry-level applications or high-volume file services

INTERNET

For up-to-date lists of cluster-certified servers, review the Certification Matrix for the Compaq ProLiant Cluster Solutions for NetWare available at: <http://www.compaq.com/highavailability>

Clustering a Server



A ProLiant cluster requires at least two ProLiant servers. Compaq ProLiant Cluster Solutions for NetWare supports 2- to 12-node configurations.

Certified ProLiant servers offer some or all of the following high-availability features as part of the Compaq Intelligent Fault Resilience:

- Hot-pluggable hard drives
- PCI Hot Plug slots
- Error checking and correcting (ECC) and RAID memory
- Online spare memory
- Redundant power supplies
- Redundant processor power modules
- Redundant cooling fans
- Easy installation through SmartStart
- Manageability through Compaq Insight Manager (all versions)

Planning Shared Storage Subsystems

Shared storage is necessary element of a cluster solution. To increase the availability of specified resources, critical data and applications often must be shared among servers. Shared storage provides continued access to this critical information if one server in the cluster fails. The remaining servers restore availability by mounting the logical drives from the shared storage subsystem previously used by the failed server.


Storage Hardware Overview

The shared storage can be configured with either shared SCSI or Fibre Channel.

Compaq shared SCSI storage solutions consist of a SCSI host bus adapter (HBA) in each cluster node and a SCSI-based storage system. A SCSI cable connects the HBA to the I/O module of the storage system.

Compaq Fibre Channel storage solutions use components to create storage area networks (SANs). Servers in the cluster have PCI-to-optical fibre HBAs that connect to Fibre Channel interconnect devices in the SAN through fiber optic cable. The interconnect devices also connect to one or more StorageWorks RAID or Modular Arrays.

Shared Storage Capacity Planning



File Locations in a Cluster

Before selecting the shared storage subsystem to build, determine the amount of storage necessary to support the applications and data on the clustered servers.

In the preceding graphic, the software running on each cluster node is divided into three categories:

- Operating system
- Clustered applications and services
- Nonclustered applications and services

All of these resources, their dependencies, and the desired level of protection must be considered when making storage decisions.

Shared Storage Requirements

Two factors help determine the required amount of shared storage disk space:

- The amount of space required for all clustered applications and their dependencies
- The level of data protection required for each type of data used by each clustered application, which includes:
 - Performance required for each drive volume
 - Recovery time required for each drive volume

Physical Space Requirement

Each node is connected to the shared storage subsystem that stores data files of clustered applications and services.

In some instances, it might be more logical to store the clustered application program files on shared storage. If the information within the program files is customized, placing the program files in the shared storage allows the secondary node to launch the application with the same customizations that exist on the primary node when a failover event occurs.

Level of Protection — Hardware RAID

RAID technology improves data availability. The array controller implements hardware RAID and is a key component of Compaq storage systems.

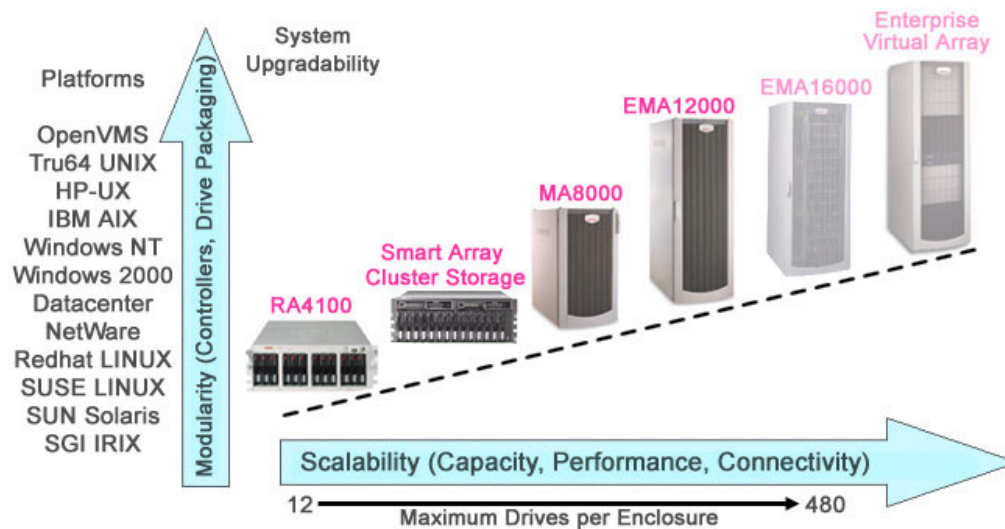
Compaq strongly recommends that:

- Hardware RAID be implemented on all shared disk drives.
- An online disk spare be used.

If RAID is not used, shared disk drive failure will disrupt all clustered applications and services that are dependent on the drive. Failover of a cluster node will not resolve this failure because neither server can read from a failed drive.

For each physical disk resource, determine the capacity and level of protection required for the data to be stored.

Shared Storage Subsystem Family



The Compaq StorageWorks RAID Array systems place the RAID controller in an external drive enclosure.



Note

Compaq ProLiant Cluster Solutions for NetWare do not support the MSA1000, EMA16000, or Enterprise Virtual Array.

The result is an external storage system that is ideal for:

- ProLiant clusters
- Databases
- Data marts
- Web servers
- Server consolidation
- Email
- Other applications that require scalability and bandwidth

Choosing RAID Levels

When choosing which RAID level is best suited to the environment, consider the following factors:

- Capacity usage and budget
- Availability and uptime requirements
- Performance requirements

The following table shows the current RAID levels, purposes, and limitations.

RAID Level	Purpose	Limitation
RAID 1+0	Mirroring: Identical data is stored on multiple drives, providing high fault tolerance and improved performance.	This RAID level requires 50% of disk capacity to be dedicated to fault protection.
RAID 5	Distributed Data Guarding: Parity data is distributed across all drives protecting against the failure of any one drive in an array. It provides improved performance at a minimum cost.	This RAID level is limited to 14-drive volumes. It also provides for continuous availability in the event that only one drive fails.
RAID ADG	Advanced Data Guarding: Two sets of parity data are distributed across all drives. This configuration provides for continuous availability in the event of the simultaneous failure of any two drives in an array. It also provides high fault tolerance at a minimum implementation cost.	This RAID level has lower performance than other RAID levels.

Refer to the following table for help in choosing a RAID method.

Most Important	Secondary Importance	RAID Level Choice
Cost effectiveness (cost per usable capacity)	Fault tolerance	RAID ADG
	Performance	RAID 5 (RAID 0 if fault tolerance is not required)
Fault tolerance	Cost effectiveness	RAID ADG
	Performance	RAID 1+0
Performance	Cost effectiveness	RAID 5 (RAID 0 if fault tolerance is not required)
	Fault tolerance	RAID 1+0



Note

In the past, Compaq sometimes referred to RAID 1+0 as RAID 0+1.

Fibre Channel Shared Storage

The Fibre Channel shared storage subsystems includes:

- Throughput of up to 100MB/s data bandwidth on one Fibre Channel Arbitrated Loop (FC-AL) or Fibre Channel Switched Fabric (FC-SW)
- Shortwave Gigabit Interface Converters (GBICs) and multimode fiber optic cables.
- Longwave GBICs and single-mode fiber optic cables (interconnect distance up to 10km per segment)
- More reliable disk link schemes critical to cluster operation
- Dynamic storage attachment and scalability (up to 12 devices on the FC-AL)
- Increased connectivity and ease of use
 - Smaller cables and connectors
 - No bus termination required
 - No need for splitters or signal enhancers
- Data path redundancy
- Reliable data transmission in a multi-initiator (shared) configuration
 - Immunity to electrical noise
 - Packet protocol with cyclic redundancy check (CRC)
 - Fault-tolerant support with RAID Levels 1, 1+0, 4, and 5
- Battery backup ECC cache
 - High levels of performance
 - Reduction of power failures
- Manageability through Compaq Insight Manager (all versions)
 - Thorough hardware status reporting
 - Automatic failure notification

Fibre Channel shared storage systems consist of:

- One or more Fibre Channel storage units (rack-mountable or tower models)
- Fibre Channel storage hubs or switches
- Fibre Channel HBAs
- Fibre Channel array controllers
- Multimode or single-mode fiber optic cables
- GBICs

In shared storage Fibre Channel architecture, servers share access to a common set of hard drives through FC-ALs or Fibre Channel switches or hubs. Servers in the cluster use Fibre Channel HBAs that connect to Fibre Channel interconnects in the SAN through fiber optic cable. The hubs connect to one or more StorageWorks RAID Arrays with fiber optic cable.


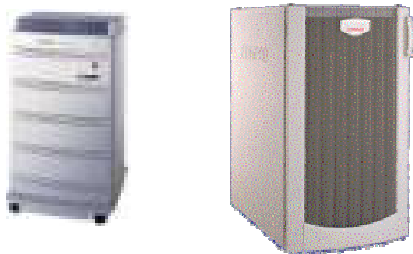




Note

Although both multimode and single-mode fiber optic cables are supported, a single FC-AL should not exceed 25Km.

Shared Storage Subsystem Features

Features of the shared storage enclosures are listed in the following table.

Component	Features
	RA4100 <ul style="list-style-type: none"> Up to twelve 1-inch drives Tower or rack configurations Support for Compaq Universal Hard Drives Redundant Compaq Smart Fibre Array controller (RA4100 controller) option Redundant power supply option Controller modes: Active/standby RAID Levels 0, 1, 1+0, 4, 5 64MB battery-backed cache per controller
	RA8000 and MA8000 <ul style="list-style-type: none"> Up to 72-drive configuration, 14 drives per SCSI bus Storage Building Blocks (SBBs) for RA8000 Compaq Universal Hard Drives (MA8000) Redundant HSG80 controller option Up to three 4x00 series disk enclosures per model ACS 8.6 controller software Redundant hot-pluggable components Controller modes: multibus or transparent failover RAID levels 0, 1, 3/5 Up to 512MB battery-backed cache per controller Compaq Rack Series 9000 (for MA8000 only)
	ESA12000 and EMA12000 <ul style="list-style-type: none"> Up to 72-drive configuration, 14 drives per SCSI bus SBBs for RA8000 Compaq Universal Hard Drives (MA8000) Redundant HSG80 controller option ACS 8.6 controller software Redundant hot-pluggable components Controller modes: multibus or transparent failover RAID levels 0, 1, 3/5 Up to 512MB battery-backed cache per controller Compaq Rack Series 9000 universal cabinet (for EMA12000 only) Up to six 4300 series disk enclosures
	Smart Array Cluster Storage <ul style="list-style-type: none"> 14 hot-plug drive bays, up to 1TB storage 4U rack-mount SCSI storage, dual bus, Integrated U3 Smart Array controller Battery-backed 128MB cache, upgradable to 256MB Hot-pluggable redundant power supplies and fans Firmware cloning capability

Shared SCSI Storage Subsystems



Shared SCSI storage subsystems consist of two server nodes with embedded Ultra 3 SCSI controllers on each server node, and a shared storage cabinet.

Planning SCSI IDs

SCSI ID	Priority	SCSI ID	Priority
7	1 (highest)	15	9
6	2	14	10
5	3	13	11
4	4	12	12
3	5	11	13
2	6	10	14
1	7	9	15
0	8	8	16 (lowest)

SCSI priorities are based on SCSI IDs. For best performance, the array controllers that are assigned the highest priority should manage the data that is used the most often.

Priorities, from highest to lowest, should be ranked as follows:

1. Database files
2. Temporary files
3. Log files
4. NetWare Loadable Modules (NLMs) or executable files

Nonshared Disk Drives

If the program files of a clustered application or service will reside on local (nonshared) storage, determine the amount of local storage that will be needed on each node.




Note

Any part of a clustered application that requires files to be placed on the local (nonshared) disk must be installed on the local (nonshared) disk of every server that will potentially host that application.

Nonshared disk drives operate the same way in a cluster as they do in a single-server environment. These drives can be located in the server drive bays or in an external storage cabinet as long as they are not accessible by both nodes.

Typically, some form of RAID is used to protect the drives and aid in restoration of a failed drive.

Planning for Storage Area Networks



A SAN uses Fibre Channel interconnect technology to connect one or more server systems to storage devices. SANs should be configured to eliminate single points of failure. Multiple Fibre Channel hubs and switches provide highly available storage for NetWare clusters.

Fibre Channel storage systems have two distinct data paths separated by the Fibre Channel interconnect:

- The first data path runs from the Fibre Channel HBA in the server to the Fibre Channel interconnect device.
- The second data path runs from the Fibre Channel interconnect device to the Fibre Channel array.


If the first data path fails, a failover of all applications results. For instance, if one server can no longer access the interconnect (and by extension the shared storage), all the cluster groups with dependencies on the shared storage fail over to the second server. The cost of failure is relatively minor.

However, if the second data path fails, there are more severe implications. All clustered applications become inoperable. Attempting to fail over the applications to another cluster node will not gain access to the Fibre Channel array.

Without access to shared storage, clustered applications cannot reach their data or log files.

The data, however, is unharmed, and remains safely stored on the physical disks inside the Fibre Channel array. Therefore, planning multiple or redundant FC-ALs or Fibre Channel switches becomes an important consideration for configuring the shared storage subsystem for the highest availability necessary for the business environment.


Multiple FC-AL Configuration



Multiple FC-AL Configuration

A NetWare cluster configuration with multiple (nonredundant) Fiber Channel Arbitrated Loops (FC-ALs) is highly scalable and highly available.


Redundant FC-AL Configuration



Redundant FC-AL Configuration

Configurations with redundant FC-ALs can fail over from one FC-AL to the paired, redundant FC-AL without interrupting service or losing access to volumes. With redundant FC-ALs, a cluster can be configured to eliminate all significant points of failure in the storage subsystem.

Fibre Channel Host Bus Adapters



Fibre Channel HBAs, and their GBICs translate electrical signals from the server to optical signals that can be sent over multimode fiber optic cable. Each server can have multiple HBAs. The number of HBAs that can be configured in each node limits the maximum number of storage arrays.

The 64-bit/66MHz PCI-to-Fibre Channel HBA is the link between the ProLiant server and the shared Fibre Channel storage subsystem, as follows:

- Transfers 528MB/s between the system board and the HBA
- Uses PCI bus mastering data transfer
- Conforms to current PCI local bus specification
- Is compatible with 32-bit PCI buses
- Supports PCI Hot Plug
- Supports FC-AL and FC-SW topologies

Gigabit Interface Converters



GBICs are industry-standard connection devices that hot plug into the storage hubs, HBAs, and controllers at each end of the fiber optic cables. GBICs convert serial electrical signals into optical signals for transmission across the Fibre Channel media and convert received optical signals back to electrical signals.


There are three types of GBICs:

- Shortwave GBIC
 - Transmission rates of 100MB/s in each direction
 - Transmission distances up to 500m per segment using multimode fiber optic cables
- Longwave GBIC
 - Transmission rates of 100MB/s
 - Transmission distances up to 10km per segment using singlemode fiber optic cables
- Very Long Distance GBIC
 - Transmission rates of 1062.5Mb/s
 - Transmission distances up to 100km per segment

GBICs ship with storage arrays and Fibre Channel HBAs. Additional GBICs and cables can be purchased separately.

Very Long Distance GBICs are not currently supported with Compaq ProLiant Cluster Solutions for NetWare.

Fibre Channel-Based Shared Storage Interconnection Technologies



NCS solutions use a shared storage architecture in which servers share access to a common set of hard drives. For all Compaq NetWare-based cluster solutions (with the exception of ProLiant packaged clusters), shared storage is accomplished through a Fibre Channel SAN.

For Fibre Channel solutions, servers connect to shared storage arrays through:

- FC-AL hubs
- FC-AL switches
- Fibre Channel SAN fabric switches

The Fibre Channel technology configuration used determines the amount of hardware necessary.

As a network topology, a hub or a switch is used as a concentrator. The hub or switch can usually determine when to insert or remove a device. Thus, a failed device or broken fiber will not keep the whole network down.

Port Assignments

Fibre Channel interconnect devices connect together and to the larger SAN using a variety of port assignments.

Specific port types are available on each of the Fibre Channel interconnect devices.

- **G_port** — A general port that has nothing connected to it




Note

It also comes up as a G_port if the other end is down.

- **F_port** — A fabric port with one host or storage port connected to it
- **E_port** — An expansion port connected to another switch
- **L_port** — Loop mode only
- **FL_port** — Loop-to-fabric cascading
- **NL_port** — Node-to-loop connectivity
- **U_port** — Universal port, which can be a G, F, E, or FL_port

FC-AL Hubs



Data Route Using a Hub

Hubs use a single continuous loop, which requires all nodes to act as repeaters and share the same 100MB/s bandwidth with all other nodes. Potentially, this can create a bandwidth bottleneck and cause difficult fault isolation.

FC-AL hubs facilitate loop implementation by connecting loop ports through a physical star configuration. FC-AL hubs typically provide 7 to 12 ports, share bandwidth, and employ bypass circuitry at each port to keep dysfunctional nodes from disrupting loop traffic.

FC-AL is not a token-passing scheme. When a device is ready to transmit data, it first must arbitrate and gain control of the loop. Unlike token-passing schemes, there is no limit on how long a device can retain control of the loop.


In general, the total number of cluster nodes and storage devices cannot exceed the number of available Fibre Channel ports on the storage hubs.



Note

The cascading of storage hubs is not supported; that is, they cannot be interconnected with other storage hubs.

FC-AL Switches




Data Route Using an FC-AL

FC-AL switches are only supported for the RA4100. They make multiple nonshared bandwidth segments available. This means that each node does not have to act as a repeater and fault isolation is much easier than with a hub. Switches typically provide eight to 16 ports, with full gigabit speeds available at each port.

A Fibre Channel loop-switch port might be configured to support a single node or a shared segment of multiple nodes (such as a loop). A loop-switch requires more processing power, memory, and microcode at each port to properly route frames. Each device attached to a loop-switch communicates in a full duplex point-to-point fashion to each device it is communicating with, without sharing bandwidth with other devices for a potential bandwidth of 200MB/s for every connected device.

Used in producing a switched fabric solution, each port of an FC-AL switch delivers up to 100MB/s, full-duplex data transfers (200MB/s). Unlike hub-based FC-AL solutions—which reduce performance as devices are added—the switched fabric performance increases as additional switches are interconnected.

Fibre Channel SAN Fabric Switches



Data Route Within a Fabric Switch

The Fibre Channel SAN Fabric Switch is an alternative to the FC-AL topology.

The integration of the Fibre Channel and network creates an intelligent interconnection scheme, called a fabric, to connect devices. Fibre Channel ports manage a point-to-point connection between the port and the fabric. The fabric makes routing decisions and carries out the transmission regardless of the protocol used. A fabric cannot exist without switch technology.

Fibre Channel SAN fabric switches are used to create the I/O paths between cluster nodes and shared storage subsystems in clusters that use the Fibre Channel SAN fabric switch topology.

Cost per Port



For many customers, the cost per port or overall unit price is one of the first considerations when choosing an interconnect device. The following table shows the relative price comparisons for Compaq Fibre Channel hubs and switches.



Interconnect Device	Part Number	Estimated Price (USD)	Cost Per Port (USD)
StorageWorks Fibre Channel Storage Hub 7	234453-001	\$1,108	\$158
StorageWorks Fibre Channel Storage Hub 12	295573-001	\$6,022	\$502
StorageWorks Fibre Channel Arbitrated Loop Switch 8	177862-B21 177862-291 (Japan)	\$5,398	\$675
StorageWorks Fibre Channel Arbitrated Loop Switch 3-Port Expansion Module	177863-B22	\$1,705	\$568
StorageWorks Fibre Channel SAN Switch 8-EL	176219-B21	\$6,750	\$844
StorageWorks Fibre Channel SAN Switch 8	158222-B21	\$16,000	\$2,000
StorageWorks Fibre Channel SAN Switch 16	158223-B21	\$35,000	\$2,188
StorageWorks Fibre Channel SAN Switch 16-EL	212776-B21	\$18,375	\$1,148

The cost per port should not be the only consideration when selecting the appropriate solution. Determine the total number of ports that will be used in addition to the cost per port for the entire device.


In other words, consider the cost per port of any ports that are not being used. Each of the interconnect devices also has other capabilities and limitations that apply to the overall value of the product.

Comparing Fibre Channel Interconnects

Component	Features
	Fibre Channel Storage Hub 7/12 <ul style="list-style-type: none"> Seven or 12 removable GBIC ports Optional management module with 12-port hub Automatic insertion of operational loop devices without impacting existing FC-AL configuration Port mode: Seven or 12 NL_ports
	Compaq FC-AL Switch <ul style="list-style-type: none"> Eight-port removable GBICs and three-port removable GBIC on Port Expansion Module (PEM) Circuit switched Standard licenses N/A Port mode: Eight NL_ports (base) and three NL_ports (PEM) 10/100MB/s Ethernet port Manageable with StorageWorks Command Console (SWCC) Cascade up to two (master and slave) or connect to fabric using FL_port as public loop device (logs in using FLOGI) Selective Storage Presentation (SSP) support Only supported with RA4100 platforms
	Compaq SAN Switch 8-EL (Entry-Level) <ul style="list-style-type: none"> Eight-port, seven-port fixed switch transceivers and one removable GBIC Packet switched, cut-through routing Standard licenses: Fabric, Zoning, Web Optional Licenses: Multiple E-port connectivity software Port mode: Seven F_ports and one F/E_port, more E_ports with license 10/100MB/s Ethernet port Manageable with SWCC SSP support
	Compaq SAN Switch 8 (Second-Generation) <ul style="list-style-type: none"> Eight-port removable GBICs Packet switched, cut-through routing Standard licenses: fabric, zoning, web, and QuickLoop Port mode: Eight U_ports (G, F, E, or FL) 10/100MB/s Ethernet port Manageable with SWCC SSP support

Component	Features
	Compaq SAN Switch 16 (Second-Generation) <ul style="list-style-type: none">• 16-port removable GBICs• Packet switched, cut-through routing• Standard licenses: fabric, zoning, web, and QuickLoop• Port mode: 16 U_ports (G, F, E, or FL)• Front panel display 10/100MB/s Ethernet port• Manageable with SWCC• SSP support
	Compaq SAN Switch 16-EL (Entry-Level) <ul style="list-style-type: none">• Sixteen 1Gb/s nonblocking Fibre Channel connections• Support for longwave or shortwave optical GBICs• Preconfigured with fabric operating software, web-based management tools, and zoning software• Dynamic path rerouting in the event of a link failure• Support for interconnecting (cascading) up to six switches for larger SANs• General 10km link distance support for switch-to-switch connections• Backward compatibility to currently shipping Fibre Channel SAN switches• Automatic data routing and rerouting, self-healing, and highest scalability

Planning NCS Intracluster Connections



The cluster interconnect is the data path over which cluster nodes (servers) communicate. This type of communication is called intracluster communication.

The cluster interconnect is used to pass information from one cluster node to another. At a minimum, it consists of one network interface controller (NIC) in each cluster node and a cable to connect the NICs. The ProLiant cluster nodes use the node-to-node interconnect data path to:

- Communicate individual resource and overall cluster status.
- Send and receive heartbeat signals between the cluster nodes that indicate that the partner servers are operational and functioning.

Because NCS does not mirror data between servers, the bandwidth consumed by the heartbeat is insignificant. Only the heartbeat and information about the health of the active cluster nodes are sent over the links.


NCS uses a standard Ethernet NIC and adapters at either 10MB/s or 100MB/s for intracluster communications. A LAN connection or dedicated direct link connects each cluster node.



Note

NCS 1.6 does not require a separate (dedicated) network connection for intracluster communications. Place the cluster interconnect on the public LAN. If a dedicated interconnect is used, ensure that every management station has access to the interconnect network.

Planning Cluster Backup and Restore

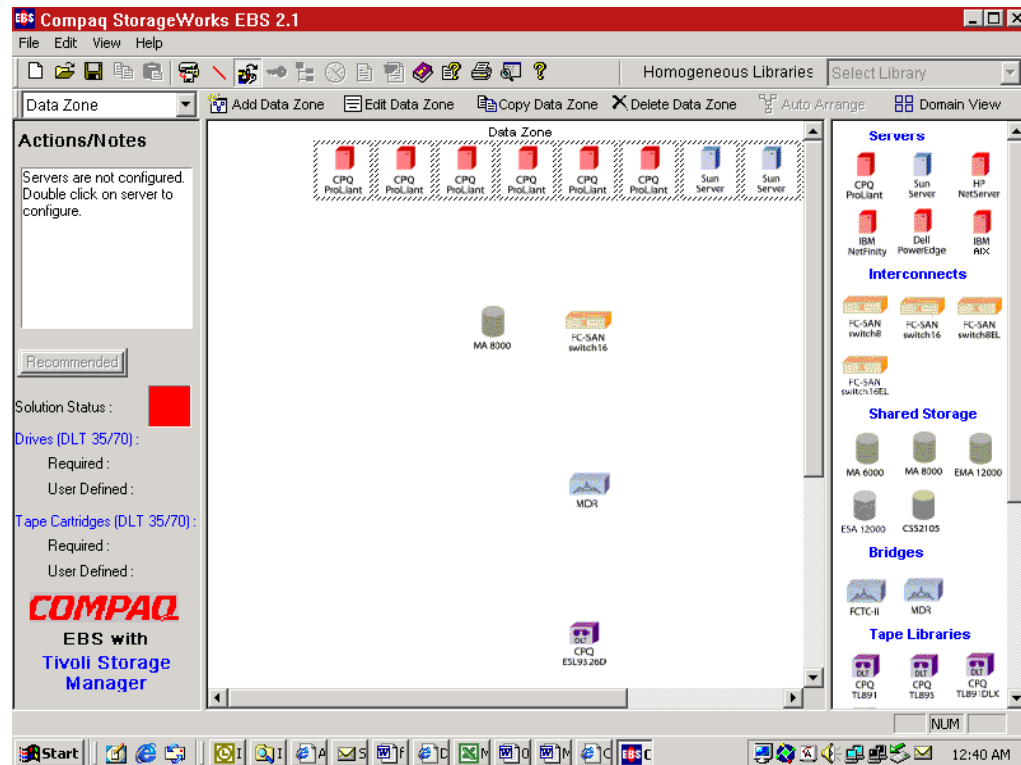


Backup of a Variety of Servers

The need for cluster backup is essential for both business-critical and nonbusiness-critical data. As data becomes an increasingly valued company asset, it must be backed up on a regular basis to ensure its security and availability. Compaq fault-resilient hardware, implemented within a clustered environment, provides a high-degree of application availability but does not prevent the deletion or corruption of files.

A solid and consistent tape backup strategy should be an integral part of high-availability strategy. In a cluster environment, additional cluster-specific configuration issues must be addressed to help select the method that best suits the desired high-availability requirements.

Compaq StorageWorks Backup Sizing Tool



Compaq StorageWorks EBS 2.1

Sizing (or optimizing) a tape backup system is resource-intensive, complex, and highly dependent on the expertise of the person designing the system. Expanding product lines, options, operating systems, and backup methodologies further complicate the sizing, so recommendations might not always be accurate.

To help quickly, easily, and accurately selecting, pricing, and deploying the optimal Enterprise Backup Solution (EBS) for the customer's environment, Compaq has developed a tool to weigh the complex array of variables and trade-offs. The StorageWorks Backup Sizing Tool takes basic parameters about the customer's tape backup environment—such as backup windows, amount of data to be backed up, tape rotation scheme, and number of servers—and produces a bill of materials for the recommended configuration of the EBS.


The StorageWorks Backup Sizing Tool provides a guide through the sizing process using a series of questions about the customer's applications, the system configuration, and the customer's requirements relating to price, performance, and capacity. Many rules, trade-off decisions, and assumptions must be considered to ensure successful results.

INTERNET




The Compaq StorageWorks Backup Sizing Tool is available as a downloadable executable file at:







<http://www.compaq.com/products/storageworks/ebs>



Enterprise Backup Components



In addition to the Fibre Channel components previously discussed, EBS solutions require specialized components. Tape libraries connect to a modular data router, which then connects to the SAN.

Component	Features
	Modular Data Router <ul style="list-style-type: none">• Modular with four slots for Fibre Channel, LVD SCSI, HVD SCSI, Ethernet management port• Internally terminated SCSI modules• Rack-mountable, 1U
	TL891 DLX Multi-Module Configuration <ul style="list-style-type: none">• Creation of multi-module TL891 virtual library• Control of up to six modules• Up to 7.2TB capacity• Storage up to 800GB of data (2:1 compression) in a single module
	TL895 DLT Library <ul style="list-style-type: none">• High capacity and high performance• Two to seven 35/70 DLT drives• 96 cartridges (up to 6.7TB storage capacity)• Data transfer rates up to 35MB/s (5MB/s per drive)

Component	Features
	MSL5026SL Library (with SDLT drives) <ul style="list-style-type: none"> • 5.72TB (2:1 compression) of storage • 5U form factor • Up to seven units in a rack-mount configuration • 26 media slots • Hot-plug drives
	SSL2020 <ul style="list-style-type: none"> • 1TB of storage in a 4U form factor • 10TB of storage in a 20U form factor • Up to 10 units in a rack-mount configuration • 100 media slots • Removable magazines
	ESL9198DLX Tape Library (with 40/80 LVD DLT drives) <ul style="list-style-type: none"> • Up to eight 40/80 drives, LVD, 198 cartridges • Simple upgrade path for future "plug-in" features • Native capacity of 7.92TB • Throughput of up to 172.80GB/hr
	ESL9198SL Tape Library (with SDLT drives) <ul style="list-style-type: none"> • 198 slots and up to 8 Compaq hot-plug SDLT drives • Up to 21.78TB of native data storage • Transfer rates up to 316.8GB/hr • Capability for multiunit scalability • Can be scaled with ESL9326SL or ESL9198SL libraries.
	ESL9326DX Tape Library (with 40/80 HVD DLT drives) <ul style="list-style-type: none"> • Up to sixteen 40/80 drives, HVD, 326 cartridges • Automated backup and restoration • 326 slots and up to 16 Compaq 40/80 DLT drives • Capacity of 13.04TB • Throughput of up to 345.6GB/hr
	ESL9326SL Tape Library (with SDLT drives) <ul style="list-style-type: none"> • High-capacity unattended backup and restore • Consolidation of backup data into a single device • 326 slots and up to 16 Compaq SDLT drives • Up to 35.86TB of native data storage • Throughput up to 633.6GB/hr

Component	Features
	ESL9326D DLT Library <ul style="list-style-type: none">• High capacity and performance• Six to sixteen 35/70GB DLT drives• 326 data cartridges (up to 22.8TB storage capacity)• Data transfer rates up to 240GB/hr (15GB/hr per drive)• Hot-pluggable power supplies, DLT drives, and fans
	SSL2020TL AIT Library <ul style="list-style-type: none">• Support of up to two AIT 50GB drives• Up to 2TB storage capacity• Field-scalable modular design for configurations with up to five modules and 100 slots• 19-cartridge magazine• 4U rack space

Cluster Backup and Restore Design Considerations

Data backups in a cluster environment can be performed either locally, across a LAN, or through a SAN. Each solution has limitations from both the hardware and software perspectives. Most backup software has varying degrees of cluster awareness.

Server-Based (Local) Backup

In a server-based backup strategy, each node in the cluster has its own backup device. Backups can be performed at scheduled times and the backup will be performed on all local disks, including any shared cluster disks that are currently controlled by that server.

The benefit of this strategy is that backup performance is increased because the backup device is directly attached to the same server as the disks and no network bandwidth is used.

However, careful planning is required to ensure that shared cluster disks reside on the server where the backup software expects to find them. If the backup software is not cluster-aware, errors will occur if the backup software attempts to back up a disk that is not currently owned by the local server.

Conversely, the backup software cannot be configured to back up a shared cluster disk that has been failed over from its partner node, resulting in data that has **not** been backed up.

LAN-Based (Remote) Backup

In a LAN-based backup strategy, the files from both nodes in the cluster are backed up over the network. The obvious drawback to this method is that performance could suffer because of the slower throughput capabilities of the network. However, determining which servers own which volumes is solved in a simple manner:

- The server running the backup connects to the cluster volume name or a cluster file share resource instead of connecting to one particular server when backing up cluster-enabled volumes.
- The backup software is configured to back up the assigned logical network drive rather than the physical drive. A connection to the cluster-enabled volume can be located and maintained regardless of which node is controlling the volume, and therefore, no additional configuration is required.

SAN-Based Backup

Using a SAN-based strategy allows the backup device to be placed on a separate network from the public LAN and allows each node to have its own virtual backup device. The Fibre Channel-based SAN allows for a centralized backup/restore solution.

Failure During Backup

Examples of server failure scenarios during a tape backup operation include:

- **The tape backup server is not a member of the cluster** — The tape backup server is connected to the cluster drives through a cluster volume name or share. If one of the nodes in the cluster fails, the tape backup software temporarily halts. The tape software should be pointed to virtual cluster drive resources and not physical drives.

Because of the auto-reconnect feature for cluster file shares, the backup software should be able to reconnect and continue with the backup, regardless of which node is actually controlling the drives.
- **The tape backup server is a member of the cluster** — The tape backup server should be configured to back up the cluster drives as logical network drives accessed through a file share, even if the drives are controlled by the tape backup server itself.

This scenario maintains the backup even if one of the drives is switched over to the other node during the backup. Because the backup software is using the logical network drive that was accessed using the cluster file share, the backup will continue.
- **The tape backup software is running on the cluster node that fails** — The tape software can be configured as a cluster group to fail over to the other node. This presents problems that cannot be overcome by noncluster-aware tape software. Because the backup software is unaware of the cluster, the only behavior that can be configured after a failure is for the backup software to switch to the partner node and restart the backup from the beginning.

Another drawback of noncluster-aware tape software is that if it is halted during a backup, it does not typically keep a log that can be used to restart the process in the middle of the backup, but only from the beginning. Depending on the backup options supported by the tape software, it might be possible to minimize this effect by backing up only files with their archive bit set, which would bypass files that have been backed up recently and not changed. In many cases, this can be achieved with an incremental backup.



Warning

In each of these scenarios, files could be corrupted if there is a fault. Until more cluster-aware backup solutions are available, the safest alternative is to restart the backup session from the beginning if a fault occurs.

Cluster Restoration

Restoring data in a clustered environment is similar to restoring data in a nonclustered environment, with the following exceptions:

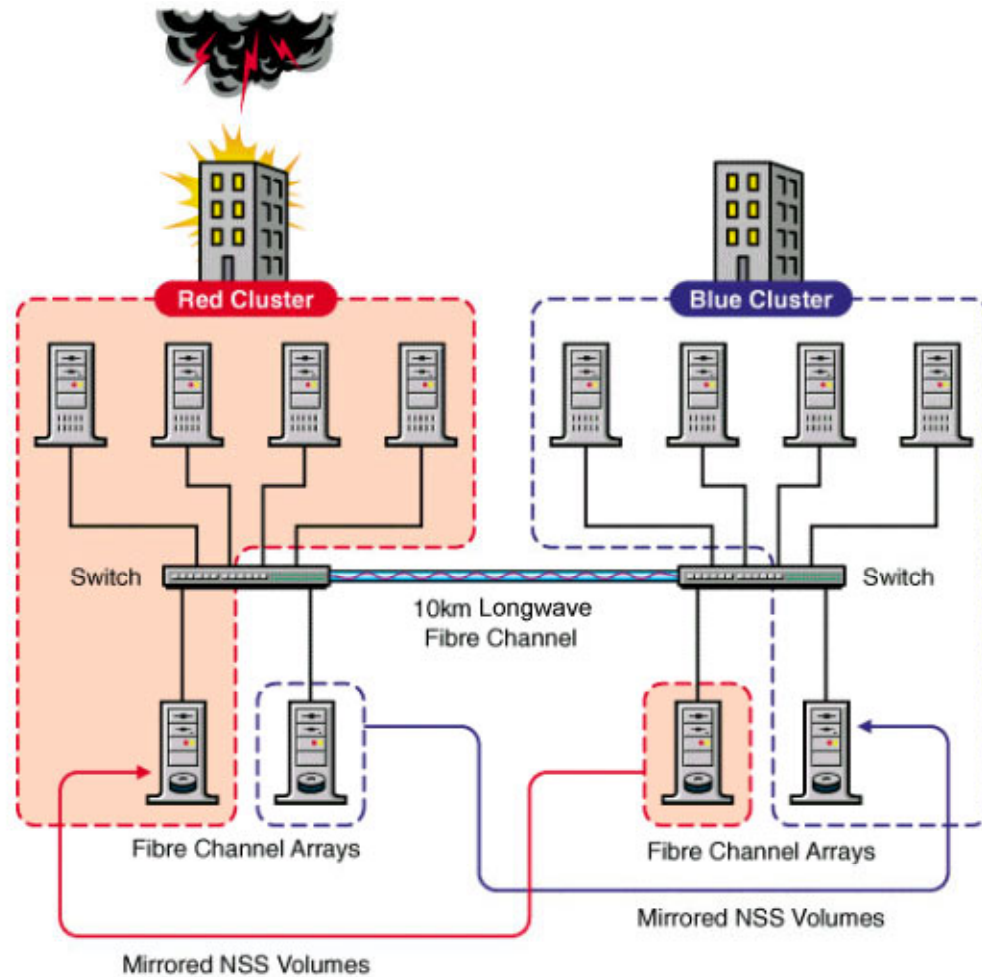
- In the event of a total failure of a single cluster node, use Novell ConsoleOne to delete that node from the cluster.
- The server must be rebuilt and reinstalled into the tree and the appropriate support packs must be reinstalled. Then Novell Licensing Services (NLS) must be reinstalled.
- When the server is brought online, it must be installed into the tree.
- NCS installation software must be run. Choose *Edit Existing Cluster* and the cluster object to edit. After the newly created server is chosen, installation is complete. Restart the new server.

NetWare 6 Backup Simplification

NetWare 6 clusters simplify the backup and restore processes. When a NetWare 5 cluster volume is backed up, the same physical node from which it was backed up must restore the volume. If restored by a different node, there is a possibility of corrupting the trustee ID list.

In contrast, when a NetWare 6 cluster volume is backed up, any node can be used to restore the volume. With any cluster node, the trustee ID list and user space restrictions remain intact.

NSS Mirroring with Novell Cluster Services



Novell Storage Services (NSS) mirroring is intended for installations that require disaster recovery. NSS mirroring is not a replacement for normal hardware RAID drive arrays.

INTERNET

Although Novell supports NSS mirroring with NCS, this configuration has not been qualified for use with Compaq ProLiant Cluster Solutions for NetWare software. For more information on NSS Mirroring with Novell Cluster Services, go to:

<http://www.novell.com/documentation/lg/ncs/configenu/data/a6prnhz.html>

Learning Check

1. List five hardware components that should be considered when planning a cluster.

.....

.....

.....

.....

.....

2. Which three of the following ProLiant servers would be appropriate for a clustering solution?

- a. ProLiant DL320
- b. ProLiant DL380 G2
- c. ProLiant DL580
- d. ProLiant ML350
- e. ProLiant ML510
- f. ProLiant ML570

3. List five features of Compaq Intelligent Fault Resilience.

.....

.....

.....

.....

.....

4. In a cluster, what does shared storage provide?

- a. Larger capacity storage devices
- b. Ease of backup
- c. 99.999% uptime
- d. Continued access to critical information if one of the servers in the cluster fails
- e. Fault-tolerant storage

5. Which RAID level provides the lowest risk of data loss?
 - a. RAID 0
 - b. RAID 1+0
 - c. RAID 4
 - d. RAID 5
 - e. RAID 3/5
6. SCSI priorities are based on SCSI IDs. Which SCSI ID has the highest priority?
 - a. 0
 - b. 1
 - c. 15
 - d. 8
 - e. 7
7. When ranking SCSI priorities, which application should receive the highest priority?
 - a. NetWare Loadable Modules, executable files, and system files
 - b. Log files used to rebuild transactions after a critical failure
 - c. Database files
 - d. Temporary files and TTS files
8. What are the three main components of a Novell Cluster Services cluster?

.....

.....

.....

9. Which of the following factors must be considered when deciding the amount of shared storage to provide in an NCS cluster? Select all that apply.
 - a. Client operating systems
 - b. Printer configurations
 - c. Clustered applications and dependencies
 - d. Type of interconnection
10. When configuring cluster interconnects, they must be configured to provide maximum data availability. How should this be done?

.....

.....

.....

.....
11. How should a tape backup solution running on a server that is not a member of a cluster be configured?
 - a. The tape backup software should be configured to point to the IP address of the shared storage device.
 - b. The tape backup server should be configured to back up the cluster drives as a logical network drive accessed through a file share, pointing to a physical drive.
 - c. The tape backup software should be connected through a cluster volume name or share. The software is pointed to a virtual cluster drive resource and not a physical drive.
 - d. The tape backup server should have a drive configured using the MAP command in a NetWare login script.
12. Which fiber switch is only supported in the RA4100 platform?
 - a. Fibre Channel Storage Hub 7/12
 - b. Compaq FC-AL Switch
 - c. Compaq SAN Switch 8-EL
 - d. Compaq SAN Switch 8

Compaq ProLiant Cluster Solutions for NetWare 6

Module 3

Objectives

After completing this module, you should be able to:

- Describe the Compaq ProLiant Cluster hardware used in solutions for NetWare.
- Describe the Compaq ProLiant DL380 Generation 2 (G2) Packaged Cluster solutions for NetWare.
- Describe the high-availability Fibre Channel solutions for NetWare.
- Identify Compaq ProLiant Cluster Solutions for NetWare software components.

Overview

The Compaq ProLiant Cluster Solutions for NetWare are designed to reduce the risk and costs of downtime resulting from hardware or software failures. These clustering systems provide the highest levels of system and application availability for customers with high system demand.

Compaq supports the Novell Cluster Services (NCS) high-availability clustering solution on the Compaq ProLiant server platform. NetWare Cluster Services (NWCS) supports NetWare 5.1 servers and allows 2 to 12 nodes. NCS 1.6 supports NetWare 6 servers and allows 2 to 12 nodes. Mixed NCS versions or mixed NetWare versions on the same cluster are not supported.



Note

Although NCS 1.6 supports 32 cluster nodes, Compaq ProLiant Cluster Solutions for NetWare currently supports 12 nodes. If support for greater than 12 nodes is desired, it is possible with the help of Compaq Professional Services or Novell Consulting Services.

Compaq product support for the NCS clustering solution takes these forms:

- ProLiant DL380 Generation 2 (G2) Packaged Cluster
- ProLiant Cluster HA/N100
- ProLiant Cluster HA/N200
- ProLiant Cluster HA/N500


The HA/N100, HA/N200, and HA/N500 Fibre Channel-based cluster solutions support mixed and heterogeneous server configurations, making them exceptionally scalable and versatile.

Features

All ProLiant clusters feature:

- **High availability at a minimum cost** — Built from industry-standard components, ProLiant clusters can be implemented for business-critical applications and data at a much lower cost than traditional, proprietary cluster solutions, without compromising availability.
- **Investment protection** — Both existing and new Compaq servers are certified for ProLiant cluster configurations. This means that clusters can be built using existing servers or a mix of old and new servers. A cost-effective migration path is available to upgrade from the HA/N100 to the HA/N200.
- **Integrated solutions** — ProLiant clusters integrate hardware and software to provide a total solution for business-critical environments. Compaq servers, storage systems, interconnect options, system management software, and implementation documentation have been thoroughly tested in cluster configurations.

Components



A typical ProLiant cluster solution consists of the following hardware elements:

- Servers
- Cluster interconnect
- Shared storage subsystems including:
 - Storage arrays with their array controllers
 - Fibre Channel interconnects for Fibre Channel storage subsystems
 - Host bus adapters (HBAs) for Fibre Channel storage subsystems

The following table shows the hardware components certified for Compaq ProLiant Cluster Solutions for NetWare.

Component	Compaq/NCS Certified Products
Compaq servers	ProLiant DL380 G2 ProLiant DL580 ProLiant DL760 ProLiant ML530 ProLiant ML570 ProLiant ML750
Interconnection	Ethernet and Fibre Channel
Shared storage system	RA4100, MA8000, EMA12000, RA8000, ESA12000, or Smart Array Cluster Storage Enclosure
Operating system software	NetWare 6
Compaq configuration software	SmartStart and Support Software CD Options ROMPaq Array Configuration Utility Online Array Configuration Utility Command line interface (CLI) or StorageWorks Command Console (SWCC) System Configuration Utility
Clustering software	Novell Cluster Services 1.6
Cluster management, monitoring, and administrative interface software	Novell ConsoleOne Novell Remote Manager Compaq Insight Manager (all versions) Compaq SANworks Secure Path for NetWare and Compaq SANworks Secure Path for NetWare for the RA4100
Compaq troubleshooting utilities	Fibre Channel Fault Isolation Utility Compaq SANworks Secure Path for NetWare and Compaq SANworks Secure Path for NetWare for the RA4100 Compaq Insight Manager 7

Compaq ProLiant DL380 G2 Packaged Cluster



The Compaq ProLiant DL380 G2 Packaged Cluster Solution consists of two Compaq ProLiant DL380 G2 servers and a Compaq Smart Array Cluster storage system. It is redundant controller-capable, has Wide Ultra 3 speed, and up to 14 hot-pluggable hard drives. It provides over 1TB of shared storage space in a sturdy, recyclable configuration fixture.

The Compaq ProLiant DL380 G2 Packaged Cluster Solution is a rack-mounted solution that is best for remote locations or data center deployment. It complements the IT investment by providing high availability for business-critical applications and by reducing downtime.

This cluster solution supports the following industry-standard operating systems:

- Microsoft Windows NT 4.0 Enterprise Edition SP6a
- Windows 2000 Advanced Server SP2
- Novell NetWare 5.1 SP3 or greater
- Novell NetWare 6 SP1 or greater
- LifeKeeper for Linux 4.0



Note

Continued support is available for the ProLiant CL380 Packaged Cluster Solution. For more information, see Appendix A of this course.

Features

The ProLiant DL380 G2 Packaged Cluster features:

- Two-node shared SCSI cluster running industry-standard operating systems
- Industry-standard DL380 G2 servers
 - Intel Pentium III FC-PGA2 1.26GHz (dual capability) with 512KB Level 2 ECC cache
 - ServerWorks HE-SuperLite chipset
 - Triple Peer PCI architecture with 133MHz frontside bus
 - 256MB PC133 MHz registered ECC SDRAM memory (2 x 1 interleaved memory), expandable to 6GB or 4GB with 2GB online spare memory
 - Three PCI expansion slots — Two HP 64-bit/66MHz and one standard 64-bit/33MHz
 - Two Compaq NC3163 Fast Ethernet Network Interface Cards (NICs) Embedded 10/100 Wake-on-LAN (WOL)
 - Smart Array 5i controller (integrated on system board)
 - Up to six 1-inch Wide Ultra3 SCSI hot-pluggable hard drives or five hot-pluggable hard drives and one AIT or 20/40GB DAT hot-pluggable tape drive support
 - Internal hot-pluggable capacity 436.8GB standard (6 x 72.8GB 1-inch hard disk)
 - 400W hot-pluggable power supply and fans (optional redundancy for both)
- Smart Array Cluster Storage enclosure with SCSI shared storage (based on Smart Array 5300 series controller technology)
 - 4U rack-mount SCSI storage, dual bus, Integrated Ultra3 Smart Array controller (redundant optional)
 - Fourteen hot-pluggable drive bays, up to 1TB storage
 - Battery-backed 128MB cache, upgradeable to 256MB cache
 - Hot-pluggable redundant power supplies and fans
 - Protection by Compaq Services Warranty Plans
 - Firmware cloning

- Specialized configuration fixture:
 - Ease of set up and configuration
 - Protection of the units during shipping and transfer
- Color coding and labeling to simplify installation
- Prefailure warranty: drives, processors, and memory
- Redundant components throughout cluster including power supplies, array controllers, and the dual SCSI bus
- Advanced Data Guarding (ADG)
- RAID level migrations and array expansion
- ROM-Based Setup Utility (RBSU)
- Preboot Executable Environment (PXE)
- Remote Insight Lights-Out Edition

INTERNET

No kit is needed for packaged cluster solutions. For a complete list of option kits and part numbers, visit:

http://www.compaq.com/products/quickspecs/10902_div/10902_div.html

Server Integrated Smart Array 5i Controller

This controller, located on the primary bus, provides two channels of embedded RAID support for the operating system and data drives as well as support for tape backup systems. This controller doubles the data rates of the Ultra2 controller in the previous generation server to 160MB/s.

The controller has two ports. Port 1 connects to the external Very High Density Cable Interconnect (VHDCI), which connects to the Smart Array Cluster Storage controller through a SCSI cable. Port 2 is designed to connect to internal SCSI hard drives and tape.

The Smart Array 5i (SA-5i) controller is responsible for allowing the operating system array drivers to communicate with the Smart Array Cluster Storage controller.

The SA-5i:

- Monitors the status of each Smart Array Cluster Storage controller.
- Initiates hardware failover between Smart Array Cluster Storage controllers.
- Resends commands to a surviving Smart Array Cluster Storage controller when necessary.



Note

The functions of the SA-5i are also supported for the SA532.

SA-5i features include:

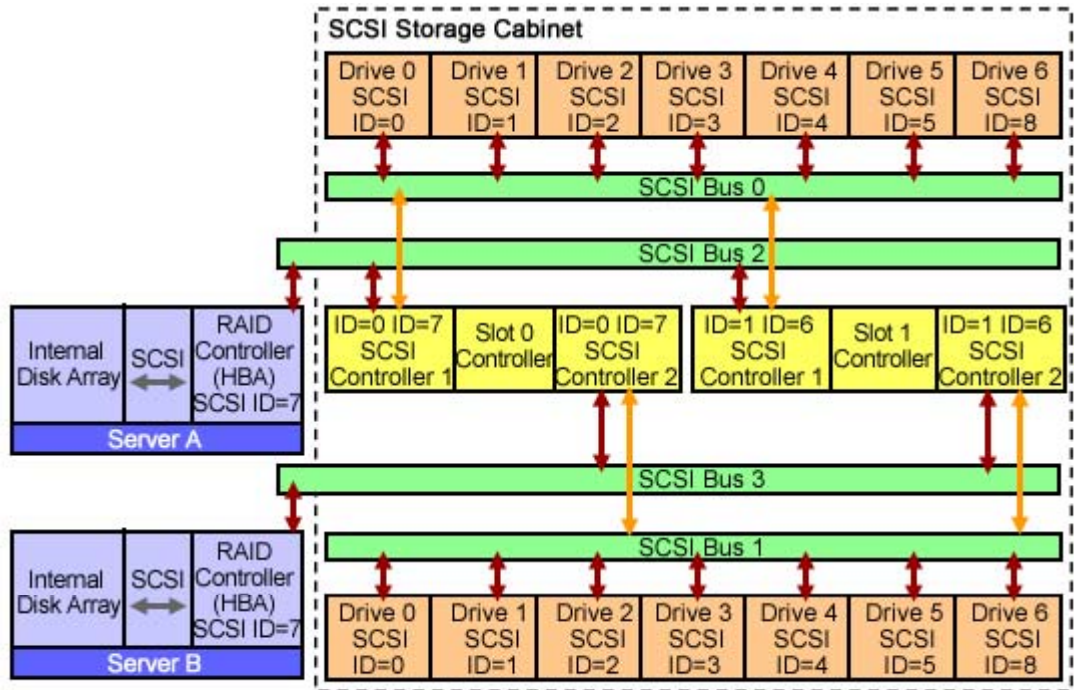
- Support for up to six internal Wide Ultra3 SCSI hot-pluggable hard drives delivering data at 160MB/s per channel
- 32MB total memory for code, transfer buffers, and read cache
- Support for RAID 0, 1, 1+0, and 5
- Offline spare with auto data recovery
- Migration from any RAID level to any other RAID level or from any stripe size to any other stripe size
- Online capacity expansion and spares
- Backward-compatible with Fast SCSI-2, Fast-Wide SCSI-2, and Wide-Ultra SCSI-3 devices
- Support for Compaq Universal Hot-Plug Tape (ATI and DAT)
- Support for internal and external tape and hard drives



Note

Performance monitoring, prefailure notification, and the prefailure warranty is provided through Compaq Insight Manager 7.


Shared SCSI Block



There are two separate SCSI buses for the server-based SCSI HBAs, eliminating the issue of conflicting SCSI IDs of duplicate SCSI IDs on the same bus. Conflicting SCSI IDs was a frequent source of calls to support services for the ProLiant CL380 Packaged Cluster.

High-Availability Fibre Channel Solutions

ProLiant Cluster HA/N100



HA/N100

The ProLiant Cluster HA/N100 offers a robust, integrated cluster solution that provides high availability for business-critical databases, large business applications, and email or file and print services.

The HA/N100 offers Fibre Channel-based clustering using one or more storage subsystems. Each storage system uses a single array controller.

The HA/N100 features:


- A 2- to 12-node solution
- RA4100, MA8000, or EMA12000 Fibre Channel storage with no redundancy
- Fibre Channel interconnects
- Support for mixed configurations

NetWare Cluster Basic Documentation Kit

The HA/N100 kit, PN 219142-B22 contains documentation in hard copy and CD formats for basic RA4100, MA8000, and EMA12000 clusters for NetWare 5.1 and NetWare 6. It does not include SANworks Secure Path.

This kit provides new information about ProLiant NetWare clusters and contains pointers to associated information. The HA/N100 kit is not required if either the HA/N200 or HA/N500 kit is purchased.

ProLiant Cluster HA/N200



HA/N200

The ProLiant Cluster HA/N200 is an RA4100 solution only. It achieves high availability through redundant Fibre Channel components. The HA/N200 provides high availability for applications and data in a business-critical NCS environment. The ProLiant HA/N200 can be configured with no single point of failure between the server and storage.

The HA/N200 delivers a high level of availability with Compaq SANworks Secure Path for NetWare for the RA4100. This program supports multipath I/O management and allows for redundant Fibre Channel loops or switched fabrics using:

- Dual Fibre Channel array controllers.
- Dual HBAs in the servers, connected to one or more RA4100 storage subsystems.

The HA/N200 features:

- A 2- to 12-node solution
- RA4100 Fibre Channel storage only, with redundancy using SANworks Secure Path for NetWare for the RA4100
- Fibre Channel interconnects

NetWare Cluster Redundant Path Kit

The HA/N200 kit, PN 219143-B22, contains the basic documentation from the HA/N100 kit in CD and hard copy formats for the RA4100 only. This kit also contains the software for SANworks Secure Path for NetWare for the RA4100 storage enclosure (two licenses) with a card that has the part number to order more licenses.

RA4100 SAN Solution


Configure the ProLiant HA/N100 and ProLiant HA/N200 to provide the RA4100 SAN Solution. This solution allows multiple servers or clusters access to primary or secondary data storage on common Fibre Channel networks.

The RA4100 delivers the Selective Storage Presentation (SSP) feature that allows ProLiant HA/N100 and ProLiant HA/N200 clusters to support heterogeneous servers and heterogeneous operating systems on a homogeneous platform.

For example, a third-party Windows NT server could share the same RA4100 with a ProLiant NetWare server. However, each server only sees its own storage area.

A Fiber Channel Arbitrated Loop (FC-AL) switch configures the HA/N100 and HA/N200 for the RA4100 SAN solution. To expand storage capacity for multinode ProLiant HA/N100 or HA/N200 clusters to nine storage boxes, use the optional StorageWorks FC-AL Switch 3-port Expansion Module in place of the 12-port hub.

ProLiant Cluster HA/N500



HA/N500

The ProLiant Cluster HA/N500 provides the highest level of data availability in the ProLiant cluster family. This Fibre Channel-based clustering system forms a complete enterprise solution that enables redundant connection between the server and the RA4100, MA8000, or EMA12000.

This platform can be configured to provide an environment in which there are no single points of failure in the storage subsystem.

SANworks Secure Path software must be used to enable support for redundant 64-bit/66MHz PCI-to-Fibre Channel HBAs in each server.

The HA/N500 features:

- A 2- to 12-node solution
- RA4100, MA8000, or EMA12000 Fibre Channel storage in redundant configurations
- Fibre Channel interconnects
- No current DRM support for stretch clusters

NetWare Cluster Redundant Path Kit

The HA/N500 kit, PN 219144-B22, contains the basic documentation from the HA/N100 kit in CD and hard copy formats for the RA4100 and the MA8000/EMA12000 storage families. This kit also contains the SANworks Secure Path software for NetWare (two licenses) with a card that has the part number to order additional licenses. The mixture of RA4100, MA8000, and EMA12000 storage systems is supported.

Features Comparison

The following table compares the ProLiant HA/N100, the ProLiant HA/N200, and the ProLiant HA/N500.

Items	HA/N100	HA/N200	HA/N500
Redundant Fibre Channel loop	No	Yes	Yes
FC-AL switch/switched fabric for the RA4100 SAN solution	Yes	Yes	Yes (with RA4100 only)
ProLiant servers	2 to 12	2 to 12	2 to 12
RA4100	Yes	Yes	Yes
MA8000, EMA12000	Yes	No	Yes
64-bit/66MHz HBA	One per server	Two per server	Two per server
RA4000 controller	One per storage array	Two per storage array	Two per storage array
HSG80 controller	One	No	Two
SANworks Secure Path for NetWare family	No	Yes	Yes

Software Components

To configure and manage a cluster, several types of software are needed:

From Novell

- Novell NetWare 5.1 or 6 operating system software, which includes:
 - Management drivers
 - Interconnect NIC drivers
 - Storage subsystem drivers
 - Support packs
- Storage management software Novell Storage Services (NSS) 3.0 — Support for manual setting of Shareable for Clustering flag
- Cluster software — NCS 1.6
- NWDEPLOY utility
- Novell ConsoleOne
- Novell Remote Manager (formerly NetWare Portal Management)
- Cluster troubleshooting software tools
 - iMonitor
 - iManage
 - NSS Verify
 - NSS Rebuild

Novell NetWare 6 Operating System Software

NetWare 6 includes the latest version of Novell eDirectory 8.6, which provides secure access to all authorized network services with a single login and makes it possible to manage the network from a single location.

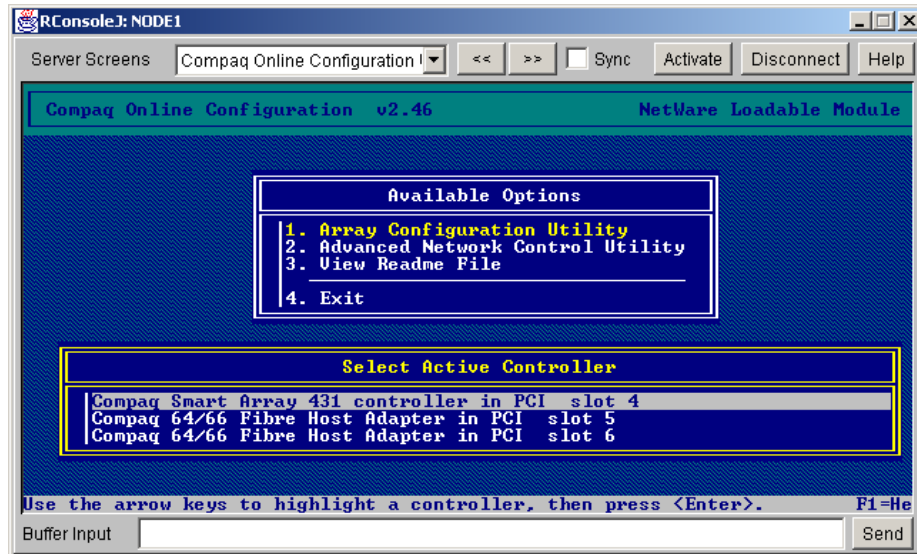
NetWare 6 also includes:

- NCS with support for two nodes
- Novell Remote Manager for browser-based management
- Centralized IP management based on open standard DNS and DHCP utilities
- Improved NSS 3, which is the native file system

From Compaq

- Cluster management software — Compaq Insight Manager 7
- Device drivers — CPQFC.HAM 2.52 or later
- StorageWorks Command Console (SWCC)
- Command line interface (CLI) — For use with the RA8000, MA8000, ESA12000, and EMA12000
- Array controller Software (ACS) 8.5 or 8.6 — For use with the RA8000, MA8000, ESA12000, and EMA12000
- SANworks Secure Path
- Offline Array Configuration Utility (ACU) — For use with RA4100 (from SmartStart 5.3)
- CPQONLIN — For use with RA4100 (from SmartStart 5.3)
- SmartStart and Support Software CD

Array Configuration Utility



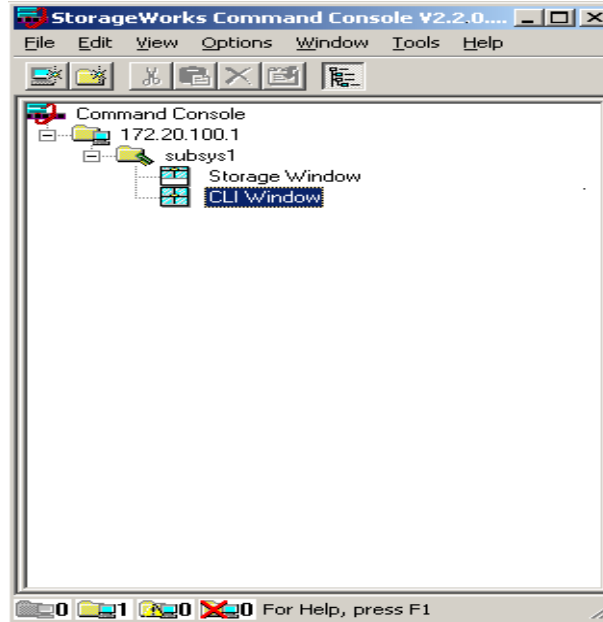
Use the offline Array Configuration Utility (ACU) or CPQONLIN utility provided with SmartStart to create logical arrays (storage sets) to be used for cluster resources and shared volumes. When creating the logical arrays, select the RAID level. Use RAID 1 or RAID 5 for fault tolerance.



Note

The Array Configuration Utility XE is not supported for NetWare.

CLI and SWCC



Command line interface (CLI) is the most direct interface to the HSG80 controller. The CLI commands manage the subsystem by viewing and modifying the configuration of the controller and its attached storage devices.

Use the CLI to start controller diagnostic and utility programs. Although the CLI provides the most detailed level of subsystem control, a graphical user interface (GUI) is also available.

The MA8000 and EMA12000 storage subsystems can also be configured using the SWCC. The SWCC GUI simplifies the configuration by eliminating the need to memorize the various line commands needed with the CLI.

The Secure Path Agent notifies SWCC clients immediately when a fault is detected.

Use SWCC to check for RAID array system faults. SWCC has a CLI window through which CLI commands can be entered.

ACS

Compaq Array Controller Software (ACS) is firmware used on the HSG80 controller. ACS resides on the controller through a PCMCIA card and requires an operating system-specific solution kit for each operating system.

ACS is a separately priced software item and is not included in the price of the controller, nor is it covered by any hardware maintenance agreement on the controller and other hardware. Choose the ACS version appropriate for the controller type and storage configuration. If software support and periodic upgrades are required, a separate software support agreement must exist for ACS.

ACS 8.6-1 includes completion of modular storage support, increased drive and SAN connectivity, and support for new drives and proactive services.

Solution kits to provide access to the functional improvements available in ACS 8.6-1 have been developed for the 10 most popular operating systems. Each kit contains the latest version of SWCC.

INTERNET For ACS solution kit documentation, refer to the Compaq website:
<http://www.compaq.com/products/storageworks/techdoc/raidstorage/>

Comparing ACS 8.5F to ACS 8.6-1

ACS 8.6-1 supports up to 12 nodes in a fully redundant configuration. ACS 8.5F is limited to eight nodes in a fully redundant configuration.

- **ACS 8.5F** — Used with the HSG80 controller, this version supports 512MB of mirror-capable, write-back cache with read-ahead capability. It supports transparent and multibus failover modes, and also supports dynamic RAIDset expansion.
- **ACS 8.6-1** — Used with the HSG80 controller, this version supports:
 - 72GB drives in a 1-inch form factor.
 - An increased number of maximum supported host connections (from 64 to 96).
 - 14 drives on a single SCSI bus.
 - 84 drives behind one controller pair with 6 fully loaded shelves of 14 drives each.
 - RAID storageset size of up to 1.024TB.



Note

Compaq ProLiant Cluster Solutions for NetWare supports ACS 8.5F, ACS 8.6-1F, and ACS8.6-1G.

The following table summarizes the abilities of each ACS version.

Version of Array Controller Software	Support
ACS 8.5F	FC-AL and FC-SW with HSG80
ACS 8.6-1G	FC-AL configurations
ACS 8.6-1F	FC-SW environments

Learning Check

1. Novell Cluster Services (NCS) 1.6 supports up to _____ nodes. The Compaq ProLiant Cluster Solutions for NetWare currently supports up to _____ nodes.
 - a. 32 and 12
 - b. 12 and 32
 - c. 24 and 8
 - d. 8 and 24
2. Which three of the following are features of the Compaq ProLiant Cluster Solutions for NetWare?
 - a. High availability at a minimum cost
 - b. Operating system independence
 - c. Support for the entire line of ProLiant Servers
 - d. Integrated hardware and software solutions
 - e. Protection of investment
 - f. Five-year on-site 7/24 warranty with 4-hour response time guaranteed.
3. The Compaq ProLiant DL380 G2 Packaged Cluster Solutions can house up to _____ Ultra 3 hot-plug hard drives.
 - a. 32
 - b. 12
 - c. 16
 - d. 14
 - e. 8

4. The Compaq ProLiant DL380 G2 Packaged Cluster Solution can support _____ shared-storage space.
 - a. 750MB
 - b. 750GB
 - c. 900GB
 - d. 1TB
 - e. Over 1TB
5. One key difference between the ProLiant Cluster HA/N100 Fibre Channel solution and the ProLiant Cluster HA/N200 solution is:
 - a. The HA/N200 supports up to the full 32 nodes.
 - b. The HA/N100 offers full redundancy.
 - c. The HA/N200 has dual Fibre Channel array controllers and dual HBAs.
 - d. The HA/N100 only supports NetWare 6.
6. The highest level of data availability in the ProLiant cluster family is provided by the:
 - a. ProLiant Cluster HA/N200
 - b. ProLiant Cluster HA/N500
 - c. ProLiant Cluster HA/F200
 - d. ProSignia Cluster HA/F500
7. The storage management software NSS 3.0 will support which of the following options?
 - a. Management of traditional NetWare volumes, including the ability to resize volumes and make the volume smaller
 - b. Remote management of file system security settings using Active Directory.
 - c. Manual setting of Shareable for Clustering flag
 - d. Backward compatibility to NetWare 3.2

8. Which software does **not** come from Compaq?
 - a. Compaq Insight Manager 7
 - b. SWCC
 - c. NWAdmin
 - d. ACS
 - e. CLI
9. What is the most direct interface to the HSG80 controller?
 - a. ACU
 - b. SWCC
 - c. CPQONLIN
 - d. RAID
 - e. CLI

Shared Storage Systems for Novell Cluster Services

Module 4

Objectives

After completing this module, you should be able to:

- Provide an overview of how shared storage is used with Novell Cluster Services (NCS).
- Describe how to configure Fibre Channel storage components.
- Plan configurations for high availability.

Overview

Compaq ProLiant Cluster Solutions for NetWare configured with the Compaq StorageWorks storage subsystems deliver high availability with automatic failover of all network resources and responsibilities in the event of a failure. Should one server in the cluster lose access to the shared storage and network, the other server assumes the responsibilities of the downed server.

ProLiant Cluster Solutions for NetWare:

- Provide support for up to 12 nodes.
- Provide support for full Fibre Channel Switched Fabric (FC-SW) and Fibre Channel Arbitrated Loop (FC-AL) storage area networks (SANs), allowing customers to use the cluster for medium- and large-scale SAN environments.
- Are based on an architecture in which clustered servers share access to a common set of hard drives.
- Require all shared data to be stored in an external storage system for the data to be readily accessible from each cluster node.

Configuring Shared Storage

Choose levels of RAID protection and create logical volumes with the Compaq Array Configuration Utility (ACU).

When configuring the storage array, allocate space for the NCS. On the partition, each server writes a status signature and updates its epoch.

In the event that the cluster-interconnect fails, the other cluster servers will check the status signature periodically to determine whether that server is operational.

**Note**

Consistently perform disk maintenance from the same server.

To configure or reconfigure the shared storage array after the NCS installation process, use CPQONLIN on the SmartStart CD to:

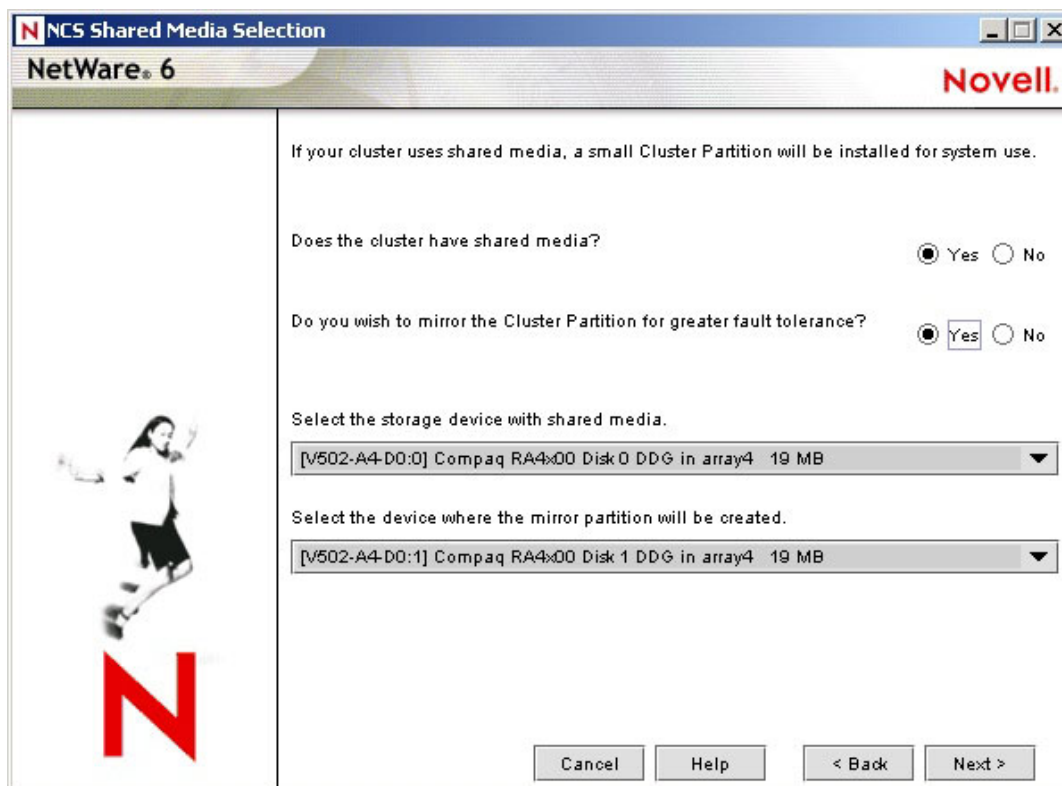
- Configure logical drives to be used for cluster resources and shared volumes.
- Allocate at least 15 to 25MB of free space to be used for the Cluster Services partition.
- Assign online spare drives to the shared storage arrays.

While it is possible to install NCS without shared storage, Compaq does not support this configuration, nor does Novell recommend it.

**Important**

The Cluster Services partition will be created during the NCS installation.

Creating the Cluster Services Partition



The NCS installation procedure automatically allocates one cylinder of one drive of the shared storage for the Cluster Services partition. If the drive where the Cluster Services partition is to be created is larger than 10GB, more than 10MB of free space is needed. The NWDEPLOY screen shows a typical Cluster Services partition allocation.

For increased fault tolerance:

- Allocate an additional 15 to 25MB of free space for a redundant Cluster Services partition (this additional space is used during the NCS installation, if required for redundancy).



Note

The specifications for NCS require a minimum of 10MB for the partition but because most hard disk drives (HDDs) today are greater than 10GB, a minimum of 15MB is recommended in this course.

- Select the appropriate RAID levels and configure online spares.



Important

Restart all cluster servers one at a time after reconfiguring the storage arrays.

Configuring the Storage System

To configure the storage system:

1. Power off the servers.
2. Power on the storage system.



Important

Power on the Smart Array Cluster Storage before powering on the servers.

3. Power on one of the servers and log in to the operating system.
4. Insert the Smart Array Cluster Storage Support Software CD.
5. Install the drivers, CPQONLIN, and other components on the server.
6. Repeat steps 3 through 5 for the second server.
7. Load cpqonlin.nlm on one of the servers and select the Smart Array Cluster controller to configure the shared storage hard drives. Refer to the *Compaq Smart Array Cluster Storage User Guide* for more information.

Applying Power

Before applying power to the shared storage system, all components of the storage system must be installed and connected to the storage interconnect, if applicable. Hard drives must be installed in the shared storage system so they can be identified and configured.

For shared SCSI, apply power in the following order:


1. Shared storage array
2. Servers

The Fibre Channel storage system must be powered up in the following order:

1. Storage hubs or switches (power is applied when the AC power cord is plugged in)
2. Fibre Channel arrays
3. Servers

Configuring for High Availability

Identifying and Eliminating Single Points of Failure



Shared Storage Points of Failure

A complete failure in the shared storage subsystem can be catastrophic because data is no longer accessible to clients. Failures can occur in either of two areas:

- The cluster node interconnect or network connections
- The path between the cluster nodes and the storage subsystems

A failure of any one of the following components in a nonredundant system will disrupt the flow of data between the server and external storage:

- Array controllers
- Gigabit Interface Converters (GBICs)
- Fibre Channel host bus adapters (HBAs)
- Fibre Channel cables
- Fibre Channel hubs or switches
- External array or server power supplies

The method for reducing the possibility of a disruption of storage services involves adding redundant backup components to the I/O path between the servers and storage subsystem. When used in conjunction with Compaq Secure Path, these redundant components provide alternate I/O paths so that if an active component fails, disk operations fail over to an alternate I/O path.



Note

The use of Secure Path is not required to manage redundant I/O paths of the ProLiant DL380 G2 Packaged Cluster Solution.

There are several ways to reduce the possibility of a disruption of storage services from the storage array. Each method involves adding redundant backup components to the data path so that if a component fails, a backup is in place to take over and allow disk operations to continue. A recommended configuration implements a redundant controller option.

In this configuration, the following components are added:

- A redundant interconnect device (hub or switch)
- A redundant controller in each storage array
- A redundant HBA in each server
- Cables and GBICs for the server-to-interconnect connections and the interconnect-to-storage connections
- A redundant power supply in each array and server


Redundant components are supported for:

- Network interface controllers (NICs)
- Fibre Channel HBAs
- Array controllers
- Fibre Channel interconnects
 - Fibre Channel hubs
 - FC-AL switches
 - Fibre Channel switches

INTERNET

For information on drivers and firmware required, see the certification matrix for the Compaq ProLiant Cluster Solutions for NetWare at:
<http://www.compaq.com/highavailability>

Configuring Storage Arrays for High Availability



Controllers Configured for Multibus Failover

ProLiant Cluster Solutions for NetWare can be configured to take advantage of several mechanisms designed to eliminate potential points of failure, which significantly increases cluster availability. Potential failure points are the hardware or software cluster components whose failure prevents the cluster from operating.

Certain failures might result in a failover in which an operational cluster node takes over the operations of the failed node. Other failures do not cause a failover and manual intervention is required to resolve the problem and restore the cluster to full operation.

In either case, there is a disruption in cluster operation when clients are unable to access their applications or data. Design a cluster to eliminate as many potential failure points as possible.

ProLiant DL380 G2 Packaged Cluster


In the ProLiant DL380 G2 Packaged Cluster, the high-availability improvement is provided through clustering server failover. Additionally, all major components have redundant options, including processors, memory, power supplies, fans, and shared storage controllers. Online spare memory supports a prefailure warranty while providing memory redundancy.

Planning Highly Available Cluster Configurations for Fibre Channel

In a cluster that uses FC-AL or FC-SW topologies, increase the availability of the I/O paths to the cluster's shared storage in the following ways:

- **By adding redundant Fibre Channel HBAs and array controllers to provide redundant I/O paths between the cluster servers and shared storage** — This requires the use of SANworks Secure Path to automatically manage the redundant paths.
- **By adding a separate, redundant arbitrated loop or switched fabric** — This protects against the failure of one of the hubs or switches. Configuring a redundant FC-AL or FC-SW requires additional Fibre Channel HBAs, array controllers, hubs or switches, and fiber optic cables. Redundant hubs or switches provide the highest level of availability in a cluster that uses the FC-AL or FC-SW topology.

Redundant Controller Option



MA8000 Redundant Controller Option

Configure storage arrays so that each has dual power supplies and connections to the SAN. To provide the highest level of fault tolerance, configure disks that are in the storage array as RAID 1 or RAID 5 logical drives.

The preceding image shows the active and inactive, and standby data paths from each server to its volumes in the storage arrays. If a component in the primary data path for either server fails, the corresponding standby component takes over and disk operations continue over the standby path.


The HBAs, storage hubs, array controllers, and cables are no longer possible single points of failure.

In the RA4100, two array controllers cannot be in the same storage system in an active/active configuration.

However, with the MA8000, two HSG80 array controllers in the same storage system can have an active/active configuration. This is called multibus-failover mode for the HSG80 controllers.

When the two HSG80 controllers are in an active/standby configuration, this is called transparent failover mode.

HBA Configurations Supported with the RA4100



An active/active HBA configuration compared to an active/standby state can be determined by how the hardware is physically attached. It is possible to have an active/active HBA configuration by ensuring that each HBA is connected to an active RA4100 array controller.

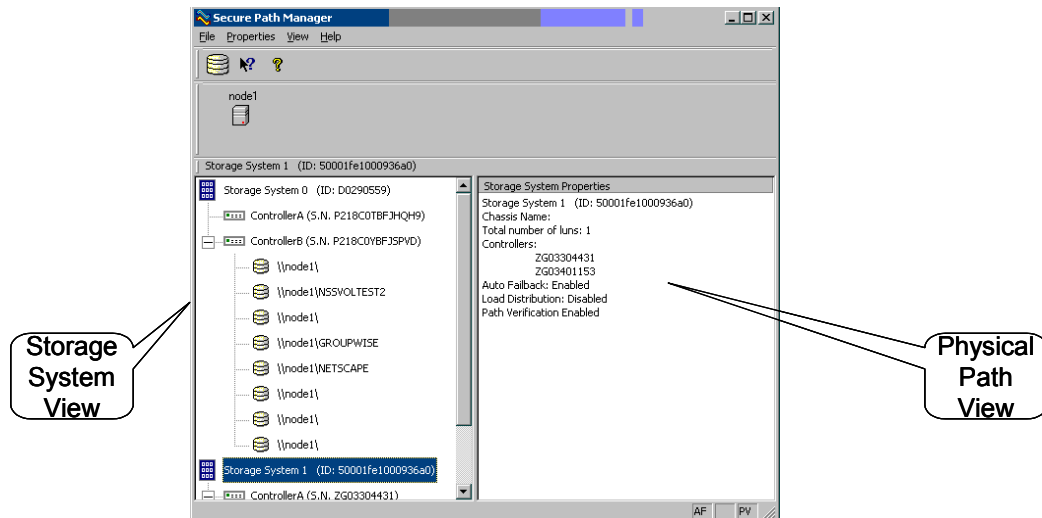
An active/active HBA configuration is not required, but it can lead to more efficient utilization of the HBAs within a system. This configuration can be achieved only if there are at least two RA4100 shared storage systems on the SAN to which the cluster is attached.



Important

During power up, the slot 0 controller (top slot in the RA4100 rack-mount model, right-hand slot in the tower model) will be the default primary controller path. Slot 1 (bottom slot in the rack-mount model, left-hand slot in the tower model) will assume the inactive standby path.

SANworks Secure Path for NetWare



SANworks Secure Path is a high-availability data path management software product that provides continuous data access for storage subsystems.

Redundant hardware, advanced RAID technology, and automated failover capability are used to enhance fault tolerance and availability.

Secure Path, in conjunction with an RA4100 or MA8000 storage subsystem, effectively eliminates the following components as single points of failure:

- Storage system
- Array controllers
- Interconnect hardware (cables, hubs, or switches)
- HBAs

Secure Path enables a dual-controller Compaq RA4100, RA8000, ESA12000, MA8000, or EMA12000 to be cabled on two independent buses, using two separate HBAs in a single server. If a failure occurs on one path's HBA, cable, or controller, the failure is detected and I/O is automatically rerouted to the alternate path. This is done at the hardware level and is transparent to the operating system and cluster services software.

Secure Path consists of three main components:

- Enhanced cpqfc.ham driver
- Server-based agent
- Secure Path Manager (SPM)

SPM is the cluster management component of Secure Path. SPM provides a GUI utility that:

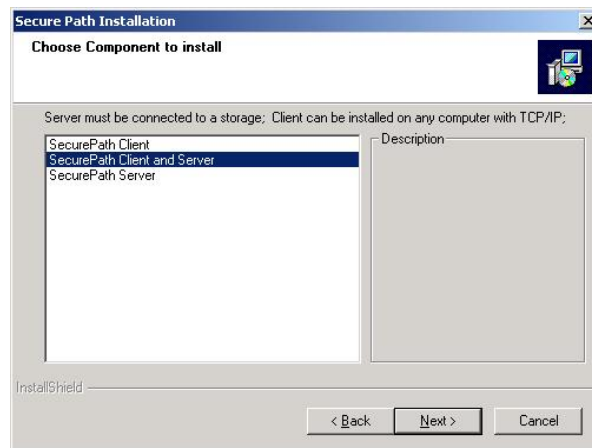
- Reports the status of the I/O paths in each redundant FC-AL or FC-SW.
- Reports disk status (path assignment, failover, and failback activity).

Secure Path allows the administrator to manually load balance logical volumes between the dual storage paths in the redundant pair. This task can be performed online, with no interruption of services, using the GUI interface.

Use SPM from a Windows environment to monitor and manage a Secure Path environment. SPM automatically detects and indicates path failures, and provides the capability to move RAID Array storage sets across controller pairs. SPM displays specific information about the state of RAID storage systems and I/O paths configured for high-availability storage access.

Use SPM to set various properties and modes associated with a managed storage profile and to set failback policy.

Installing Secure Path for NetWare



Secure Path Installation Screen

For redundant components, after configuring the storage subsystem, install and configure Secure Path from the Compaq SANworks Secure Path CD. The details for installing Secure Path are given in the *Compaq SANworks Secure Path for NetWare Installation and Reference Guide*.

INTERNET

To download a copy of the guide, choose the document title from the following website: <http://www.compaq.com/products/storageworks/Storage-Management-Software/sptechdoc.html>

Installing Secure Path involves the following basic tasks:

1. Install the Secure Path driver and Secure Path agent on the cluster servers.
2. Install Secure Path Manager on one or more workstation clients.

The following sections provide operational information regarding SPM:

- Responding to a lost host connection
- Verifying a path
- Repairing a path
- Detecting and identifying path failures
- Identifying controller failovers
- Responding to failover events

Responding to a Lost Host Connection

SPM monitors the connection status for each active host that is a member of the current storage profile.

A server icon displays for each host in the window frame located immediately underneath the tool bar. The host's name is listed above the icon and a cluster name is listed below if it is a member of a cluster.

SPM monitors its connection with each member of a storage profile and indicates a loss of connection to a particular host with a red X. The red X can also indicate that the Secure Path agent is not running on the host.

When investigating possible problems with lost host connections, consider the following:

1. **Connection loss** — A loss of connection does not necessarily mean that Secure Path has stopped protecting storage on that host. If the host is still running, the problem is most likely due to a network connectivity problem. Another possibility is that the Secure Path agent is not running and only the Secure Path remote management functions are lost. The Secure Path cpqfc.ham multiple path driver is still protecting availability to the storage.
2. **Cluster member** — SPM continues to report storage information based on data received from the surviving host or hosts. Check the cluster management utilities to determine whether storage resources have failed over to a surviving host.
3. **Host communications** — SPM automatically re-establishes communication to a host when the connection becomes available.
4. **Path Verification** — With Path Verification (PV) enabled, Secure Path periodically runs diagnostics on all preferred and alternate paths to determine their current state. If a path is diagnosed with a problem that would prevent reliable I/O operations from completing, it is marked as failed and no further I/O operations are permitted on that path.

If PV is enabled, the failure will show in the bottom right corner of the SPM window in the status bar.



Note

Path Verification must **not** be disabled for NetWare.

Verifying a Path

Choose *Verify a Path* for SPM to determine the current state of a path. To verify a path:

1. Click the path.
2. Right-click and select *Verify Path*.

SPM generates a pop-up message when the verification completes to indicate the result of the operation. No state change occurs as a result of this operation.

Repairing a Path

Choose *Repair a Path* for SPM to restore access to a failed path after the problem has been corrected. To repair a path:

1. Click a path in the failed state.
2. Right-click and select *Repair Path*.

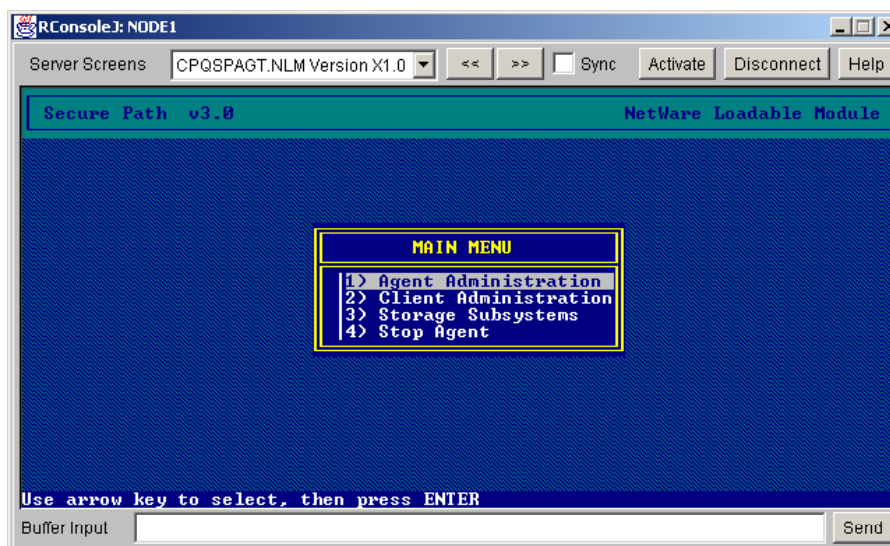
If the repair action is completed successfully, the state of the path changes to *Available* if its mode is *Alternate*, or to *Active* if its mode is *Preferred*.



Note

Repairing a path is also referred to as manual failback.

Detecting and Identifying Path Failures



SPM periodically monitors the status of all systems in the storage profile at a rate determined by the polling interval. To indicate failures, icons are used in the Storage System view and path states are set to *failed* in the Physical Path view.







The Secure Path Agent notifies StorageWorks Command Console (SWCC) clients immediately when a fault is detected. Monitor SPM status routinely to check for occurrences of failover events that could compromise either the performance or availability of storage resources.

Availability is compromised when a configuration includes only two configured paths to a storageset and one path is lost because of component failure. Secure Path cannot fail over to a redundant path if a subsequent fault occurs in this situation.

The SPM client is not required to be running for Secure Path to protect path availability. The cpqfc.ham device driver running on the host controls the Secure Path automated path protection capability.

Path Failure Icons

Several types of icons display in the SPM Storage System view to indicate the presence of a path failure. Recognizing these icons will help determine the specific storageset and path associated with the failure.

Icon	Indication	Explanation
	Storage system path failure detected	Secure Path has detected a failure of at least one, but not all, paths to the RAID Array storage system. Browse the storage system to determine the affected controller and storagesets.
	Storage controller path failure detected	Secure Path has detected a failure of at least one, but not all, paths to the storage controller. Browse the storage controller to determine the affected storagesets. Unless the Path Verification property is enabled, Secure Path only detects failures for paths with active I/O. One or more paths could be failed over to other storagesets owned by the same controller, but not yet detected by Secure Path. However, Secure Path will perform path or controller failover of these drives and indicate the failure if subsequent I/O occurs to any storagesets. If PV is enabled, Secure Path will automatically detect the failure of paths to all the affected storagesets on the controller and will immediately perform whatever path or controller failover activity is necessary to maintain availability.
	Storageset path failure detected	Secure Path has detected a failure of at least one, but not all, paths to the storageset. Click the storageset to highlight it and examine the Physical Path view information in the right pane to determine the specific nature of the path failure.
	Storage system failure detected	All paths to the affected storage object have failed.
	Storage controller failure detected	All paths to the affected storage object have failed.
	Storageset failure detected	All paths to the affected storage object have failed.

Identifying Path Failovers

To identify the source of path failover activity, first check the Storage System view for path failed icons, and then examine the Physical Path view of the affected storageset. Check for paths that indicate failed status.

Whether there is one or more paths to a particular storageset in the failed state depends on the following conditions:

- **Was I/O active on the affected storageset?** — Secure Path determines path failures by detecting the failure of I/O operations to complete. This means that if I/O was not active on a broken preferred path, the fault is not detected and the state of the path is not marked as failed until I/O operations occur.
- **Is PV enabled?** — PV periodically tests the viability of all paths and automatically detects faults on all preferred and alternate paths. This means that a controller failover on installations with multiple paths to a storageset results in failed states for both the preferred and alternate paths to the failed controller.

Identifying Controller Failovers



Failed Storage Controller Icon

A RAID Array controller failure will cause Secure Path to change the ownership of a given storageset to the surviving controller. Failover will occur only for those storagesets with active I/O operations.

If a controller failover has occurred, use the PV feature to check the viability of all configured paths. Although it can be enabled at anytime, PV requires approximately two minutes per storageset to verify the integrity of all paths in the storage profile.

The PV diagnostics will identify the failing controller in the Storage System view. Check for the failed storage controller icon. SPM will show all storagesets located on the failing controller have been failed over to the surviving controller. Because all alternate paths to the faulty controller will also transition to the failed state, storageset path failure icons will display for each storageset on the surviving controller.

Responding to Failover Events

When investigating possible problems with failovers, questions to ask include:

- Are there additional available paths remaining to the storageset or has this failure totally eliminated the ability to survive any subsequent failures?
- What caused the failure?

Most storage channel problems are caused by failures in the interconnect hardware. Use SWCC to check for RAID Array system faults, and visually inspect the switches or hubs for LED or LCD hardware fault indications.

Learning Check

1. Describe how shared storage provides a high-availability solution.

2. List four features of the Compaq ProLiant Cluster Solution for NetWare?

3. The _____ utility found on the SmartStart CD is used to configure the shared storage array after the installation process.
4. What is the minimum amount of free space that should be allocated for the Cluster Services partition?
 - a. 500MB to 1GB
 - b. 500GB to 1TB
 - c. 10MB to 25MB
 - d. 100MB to 500MB
5. Put the following steps in the order to start a Fibre Channel Storage System.
 - a. Start servers.
 - b. Start storage hubs or switches.
 - c. Start Fibre Channel arrays.

6. List the steps for configuring the ProLiant DL380 Packaged Cluster to be sent to another location.
.....
.....
.....
.....
.....
.....
.....
.....
7. What utility is used to configure the ProLiant DL380 Packaged Cluster?
 - a. ORCA
 - b. cpqonlin
 - c. Console One
 - d. Install.NLM
 - e. ACU
8. A complete failure in the shared storage subsystem can be catastrophic. List the two areas where failures can occur.
.....
.....
9. Which of the following components have redundant options in the ProLiant DL 380 G2 Packaged Cluster?
 - a. Processor
 - b. Memory
 - c. Power Supply
 - d. Fan
 - e. Shared Storage Controller

10. List two ways to increase the availability of the I/O path to the clusters shared storage in a high-availability Fibre Channel cluster configuration.
-
-
11. What are the three main components to SANWorks SecurePath?
- a. ACU
 - b. Client-based agent
 - c. Server-based agent
 - d. Enhanced CPQFC.HAM driver
 - e. SPM
 - f. CLI
12. If a system in the Secure Path storage profile does not respond, what is the path state set to in the Physical Path view?
- a. Down
 - b. Inactive
 - c. Unresponsive
 - d. Failed
 - e. Off line

Objectives

After completing this module, you should be able to:

- Identify supported network interface adapters.
- Describe the methods used to maintain high availability in network subsystems.
- List supported transport protocols.
- Explain the effect of failover on network capacity.

Network Adapters Supported

Standard Compaq Ethernet adapters are the network interface controllers (NICs) of choice for Compaq ProLiant clusters. Either 10Mb/s, 100Mb/s, or 1000Mb/s Ethernet can be used.

The following table shows the NICs supported and their part numbers.

Server	NIC	NIC Option Part Number
ProLiant DL360	NC3163 Fast Ethernet NIC Embedded 10/100 (standard)	
ProLiant DL380 G2	NC3123 Fast Ethernet NIC PCI 10/100 (optional)	174830-B21
	NC3134 Fast Ethernet NIC 64 PCI Dual Port 10/100 (optional)	138603-B21
	NC6136 Gigabit Server Adapter, 64-bit/66MHz, PCI, 1000 SX (optional)	203539-B21
	NC7131 Gigabit Server Adapter, 64-bit/66MHz, PCI, 10/100/1000-T (optional)	158575-B21
ProLiant DL580	NC3134 Fast Ethernet NIC 64 PCI Dual Port 10/100 (standard)	138603-B21
ProLiant DL760	NC3123 Fast Ethernet NIC PCI 10/100 (optional)	174830-B21
ProLiant ML750	NC6136 Gigabit Server Adapter, 64-bit/66MHz, PCI, 1000 SX (optional)	203539-B21
	NC7131 Gigabit Server Adapter, 64-bit/66MHz, PCI, 10/100/1000-T (optional)	158575-B21
ProLiant 8000	NC3131 Fast Ethernet 64 PCI Dual Port 10/100 controller (standard)	338456-B21
ProLiant 8500	NC6136 Gigabit Server Adapter, 64-bit/66MHz, PCI, 1000 SX (optional)	203539-B21
	NC7131 Gigabit Server Adapter, 64-bit/66MHz, PCI, 10/100/1000-T (optional)	158575-B21
	Compaq NC3123 Fast Ethernet NIC PCI 10/100 (optional)	174830-B21
ProLiant ML330	NC3163 Fast Ethernet NIC Embedded 10/100 (standard)	
ProLiant ML350	NC3123 Fast Ethernet NIC PCI 10/100 (optional)	174830-B21
ProLiant ML370	NC3134 Fast Ethernet NIC 64 PCI Dual Port 10/100 (optional)	138603-B21
	NC6136 Gigabit Server Adapter, 64-bit/66MHz, PCI, 1000 SX (optional)	203539-B21
	NC7131 Gigabit Server Adapter, 64-bit/66MHz, PCI, 10/100/1000-T (optional)	158575-B21
ProLiant ML530	NC3123 Fast Ethernet NIC PCI 10/100 (standard)	174830-B21
ProLiant ML570	NC6136 Gigabit Server Adapter, 64-bit/66MHz, PCI, 1000 SX (optional)	203539-B21
	NC7131 Gigabit Server Adapter, 64-bit/66MHz, PCI, 10/100/1000-T (optional)	158575-B21
	NC3134 Fast Ethernet NIC 64 PCI Dual Port 10/100 (optional)	138603-B21



Note

This is not a comprehensive listing of supported network controllers for NetWare 6. Any NIC certified for NetWare 6 can be used with Novell Cluster Services 1.6.

Configuring Communication Links for High Availability


Both the cluster interconnect and the public network link affect the availability of a cluster.

If the cluster interconnect fails and a server becomes unable to communicate with the other servers and perform split-brain arbitration, the failed server is removed from the cluster and its shared resources are migrated to another server.

Network service is disrupted while the cluster detects the failure, migrates the service to another server, and restarts the service on that server. The disruption might last a short time, or it might not. Compaq addresses the disruptions by providing features that can keep downtime to a minimum. The features are part of the Compaq Advanced Network Services and involve configuring either:

- Two Compaq Ethernet adapters.
or
- Two ports on a single adapter, so that one is a hot standby for the other.

Compaq Advanced Network Services



Compaq Advanced Network Services is a feature of the CPQANS.LAN driver when used in conjunction with N100.LAN and N1000.LAN drivers.

It provides:


- Network Fault Tolerance (NFT) that:
 - Is available for the Compaq NICs that use an Intel chipset.
 - Creates a team of two to four NICs to provide automatic redundancy for the network controller.
 - Uses primary and secondary controllers.
- Adaptive Load Balancing (ALB) that:
 - Creates a team of two to four NICs to increase transmission throughput.
 - Is only effective in an Ethernet switched environment.
 - Becomes the NFT option only if used with a hub or repeater.

Some Compaq NICs support ALB, which provides the benefits of NFT and the added advantage that the NICs are grouped to increase NIC throughput.

These features reduce the possibility of downtime because of network failure by allowing the administrator to configure a redundant path for each interconnect. In this configuration, each interconnect is configured with a primary path and a standby path. If a component in the primary path fails, communication continues over the standby path.

Redundant interconnects greatly enhance the availability of ProLiant clusters.

Network Adapter Teaming



Network adapter teaming combines two or more physical NICs into a single logical NIC with multiple IP addresses, known as NIC teaming. Two embedded NICs ensure that the server can always keep an active link. If the primary network connection fails, the job of processing network requests automatically fails over, retaining the network connection. This feature must be configured.

NIC teaming is also tightly integrated with Compaq Insight Manager, providing proactive notification when the primary NIC fails. This configuration, when coupled with a dedicated interconnect for cluster communications, provides redundant paths for both client and cluster communications.

**Note**

Teaming redundant pairs should not be used for dedicated intracenter heartbeat connections, unless NetWare 6 SP1 has been implemented.

Advantages of Teaming

The primary advantage of this option is fault tolerance. If a NIC fails in a teamed scenario, none of the IP addresses will become inaccessible. This is because the IP addresses are not bound to a single physical NIC, but rather to the whole logical NIC. The software that facilitates NIC teaming automatically realizes that one NIC is having a problem and stops using it. Generally this process occurs with little or no impact to the user.

Another advantage of teaming is speed. When NICs are teamed, they basically pool all their available bandwidth. So if three 100Base-T NICs are teamed, then there is an available bandwidth of 300Mb/s. This is advantageous when one IP address is using only a small portion of the bandwidth that it would normally use on one card and another IP is maximized.

Dual-Port NICs

Most NICs have a single port to which a single network cable connects. Ordinarily, if two distinct network communication paths are needed from a single server, two NICs are placed in the server, and two expansion bus slots are used. A dual-port NIC, however, has two ports, each of which supports its own network connection. Only one expansion bus slot is used.

The two ports of a dual-port NIC can be configured to be redundant. In this configuration, one of the NIC ports is configured as a hot backup for the other. The primary port operates normally, sending and receiving data. The second port remains in a standby state until the primary port encounters a failure. When the primary port encounters a failure, the standby port takes over. Data flow is not interrupted during the transition.

Types of Teaming

All Compaq ProLiant Ethernet network adapters support the following three types of teaming:

- **Network Fault Tolerance (NFT)** — Provides simple redundancy with two to eight network adapters in a fault-tolerant team. Each server can support up to eight teams where one adapter per team is defined as the primary adapter. All other adapters are secondary.

NFT teaming functions at any speed, on any media. It is switch-independent and can be split across Layer 2 switches, but must be in the same Layer 2 domain.

- **Transmit Load Balancing (TLB)** — Was formerly known as Adaptive Load Balancing (ALB). TLB supports adapter fault tolerance and load balancing of IP traffic being transmitted from the server. It incorporates all the features of NFT and load balancing.

TLB is also switch-independent and can be split across Layer 2 switches; however, all TLB team members must be in the same Layer 2 domain.

TLB provides both failover and load balancing of transmit IP traffic across all adapters for increased performance. With TLB, traffic received by the server is not load balanced. The primary adapter is responsible for receiving all traffic destined for the server.

- **Switch-assisted Load Balancing (SLB)** — Was formerly known as Fast EtherChannel (FEC) Gigabit EtherChannel (GEC) teaming. For even higher performance, it incorporates all the features of NFT and TLB to provide:
 - Adapter fault tolerance.
 - Load balancing of all traffic being transmitted from the server.
 - Load balancing of traffic received by the server.


With SLB teaming, there are from two to eight network adapters in a team and up to eight teams in a server. All adapters transmit and receive using the same speed. This approach must be used in conjunction with an intelligent switch that supports this type of teaming, and all ports must be connected to the same switch.

Adapter Failover

The most common reasons that network adapters fail over are:


- **Physical link failure** — A physical link failure is anything that causes the link LED indicator on the back of the network adapter to go out, such as a pulled cable, switch power loss, or other such condition.
- **Heartbeat failure** — A heartbeat (not the cluster heartbeat) failure causes Compaq ProLiant network adapters to fail over when detected, such as when heartbeat packets are not successfully transmitted and received from one network adapter to another through the network infrastructure. They are only transmitted to increment the receive-counter on a network adapter that has been idle to verify that the network adapter still has network connectivity.

Redundant Hubs and Switches



Private Interconnect with Ethernet Hub and Cables

An Ethernet hub or switch connects the clusters in a public LAN. Up to 12 heterogeneous ProLiant servers can be used in an NCS cluster configuration.



Redundant NICs and Interconnects with Ethernet Hub and Cables

If a single hub or switch is used to connect both adapters in a fault-tolerant team, the hub or switch could fail.

For greater reliability, connect the redundant NICs to separate hubs or switches.

Supported Transport Protocols

TCP/IP

LAN communications are an integral part of NCS clustering. All clustering data must travel the LAN using the TCP/IP protocol for proper intracenter communications.

TCP/IP is the only communication protocol suite supported in an NCS cluster environment. Client machines on the public LAN connect to services running on the cluster using the TCP/IP service. All servers in the cluster must be configured with TCP/IP and must be on the same IP subnet.

Volume resource objects are configured with their own IP address.

Most services and applications run from shared storage and clients map to virtual servers. If the service fails, it fails over to another server in the cluster. Because the client mapping is to a virtual server, remapping is not necessary. Clients reconnect transparently and typically do not notice a server failure, depending on their activity and the type of client operating system that is running.

IP Planning

Before installing a NetWare cluster, it is essential to plan the IP addressing scheme. NCS is closely integrated with the NetWare 6 operating system. Unlike past versions of NetWare, NetWare 6 and NCS make extensive use of the TCP/IP protocol stack.

A NetWare cluster configuration requires static IP addresses for:

- Cluster nodes for LAN access.
- Cluster Master IP Address.
- Cluster resources and volumes (applications and services).

IP address planning makes managing a cluster easier and ensures availability of IP addresses for the cluster nodes and cluster resources. Clients can use either static or dynamic IP addresses. A recommended technique for assigning IP addresses is to reserve blocks of IP addresses for related items in the cluster.

Examples of potential reserved blocks of IP addresses to use for this purpose are included in the following table.

Object	IP Address Range	Subnet Mask
Cluster nodes	172.20.xxx.1 to 172.20.xxx.12	255.255.0.0
Cluster Master IP Address	172.20.xxx.15	255.255.0.0
Cluster resources and volumes	172.20.xxx.xxx to 172.20.xxx.xxx	255.255.0.0

Network Capacity

The cluster nodes must have enough network capacity to handle requests from the client machines and to handle failover and failback events gracefully.

Ensure that both servers can handle the maximum number of clients that can attach to the cluster. If node 1 encounters a failure and its applications and services fail over to node 2, node 2 needs to handle access from its own network clients as well as those that normally connect to node 1.

It is important to understand the effect of failover on network I/O bandwidth. When the cluster encounters a server failover event, only one server is responding to network I/O requests in a two-node cluster. If it is already operating near maximum capacity, taking on the load of the failed server might cause server degradation in response to client requests.

In multinode clusters, the resources of a failed server can be distributed to multiple surviving nodes. This is known as Fan-Out Failover™ and prevents any one server from becoming overloaded. Even if more than one node fails at a time, any of the operational servers can be assigned to restart failed applications and mount associated volumes automatically. Distributing the load across many nodes allows users to continue accessing resources without any noticeable slowdown in response time.

Ensure that the network speed and protocol of the surviving node will sufficiently handle the maximum number of network I/Os necessary to support critical services.

Learning Check

1. Which of the following best describes the supported network cards for Novell Cluster Services 1.6?
 - a. Any NIC certified for NetWare 6
 - b. All token ring cards
 - c. All Compaq network cards
 - d. Any Ethernet card
 - e. Any NIC that has a Linux driver
2. List the two features of Compaq Advanced Network Services that can be configured to minimize downtime caused by a NIC failure.

.....

.....
3. Choose two network card drivers that work with the Compaq Advanced Network Services driver (CPQANS.LAN).
 - a. NE2000.LAN
 - b. N1000.LAN
 - c. TOKEN.LAN
 - d. N10.LAN
 - e. N100.LAN
4. How many NICs can be used to provide automatic redundancy with the Compaq Advanced Network Services?
 - a. 1
 - b. 2
 - c. 2 to 4
 - d. 2 to 8

5. What is ALB?
 - a. Automatic Link Bridging
 - b. Automatic Load Balancing
 - c. Automatic Load Bridging
 - d. Automatic Link Balancing
6. What is Network Adapter Teaming?

.....

.....

.....
7. Which of the following are types of NIC teaming supported by Compaq ProLiant Ethernet network adapters? Select all that apply.
 - a. SFT
 - b. NFT
 - c. ALB
 - d. TLB
 - e. SLB
8. What are the two most common reasons for network adapter failover?
 - a. Dust
 - b. Physical link failure
 - c. Heartbeat failure
 - d. Static electricity
 - e. Power surge
9. Which communication protocol is used to transport clustering data?
 - a. TCP/IP
 - b. IPX
 - c. SPX
 - d. RIP
 - e. NLSP

Objectives

After completing this module, you should be able to:

- Describe Novell Cluster Services (NCS).
- List and define the functions of Cluster Resources.
- List and define the functions of the cluster system architecture.
- Identify the major steps and the hardware and software requirements for installing or upgrading to NCS.
- Identify NetWare 6 client configuration parameters for NCS.

Novell Cluster Services



Novell Cluster Services

Novell offers NCS 1.6 to complement NetWare 6. NCS is the clustering software system implemented on ProLiant servers to provide an active/active, general-purpose solution that ensures high availability and manageability of critical network resources including data, applications, server licenses, and services.

NCS is a multi-node, eDirectory-enabled clustering product for NetWare 6 that supports failover, failback, and migration (load balancing) of individually managed cluster resources.



Note

NetWare 6 ships with a two-node license for NCS. Support for all 32 nodes is present but must be activated with additional licenses.


NCS provides important features to help ensure and manage the availability of network resources, such as:

- **Scalability** — Increases the size of a cluster by adding additional servers (up to 12) while the cluster is online.
- **Online cluster software updates** — Updates the cluster software while the cluster is online.
- **Novell Storage Services (NSS) 3.0** — Updates nodes with NSS information and adds new NSS information to the cluster.
- **Improved volume mount protection** — Prohibits attempts to manually double-mount a cluster volume.
- **ConsoleOne client** — Permits multiple instances of the ConsoleOne client to be opened on different management clients.

Functionality in NCS includes:

- Management and protection of data storage pools containing multiple volumes.
- NetWare Remote Manager (NRM), a Java-enabled web browser, providing a fully functional web-based alternative to ConsoleOne for configuring and managing NCS 1.6 clusters.
- Simple Network Management Protocol (SNMP) support for improved cluster management and monitoring of events.
- Simple Mail Transport Protocol (SMTP) support for notification of cluster events and status through email for up to eight destinations.
- Diagnostic tools, including an event log that can be viewed from ConsoleOne or NRM.
- Several resource templates to simplify creating and managing resources.

Typical Cluster Configuration



Although there are a variety of ways to configure a cluster, every configuration contains common elements. The following list describes a typical cluster configuration:

- Two to 12 Compaq ProLiant servers
- Shared storage
- Client
- Supported Fibre Channel interconnect devices including:
 - Fibre Channel storage hub 7 or 12
 - Fibre Channel Arbitrated Loop (FC-AL) Switch (RA4100 only)
 - Fibre Channel Switched Fabric (FC-SW) family

- Latest versions of NetWare client software (only needed if access is required to specific applications using NetWare Core Protocol {NCP} protocol is required)
- Compaq StorageWorks devices including:
 - RA4100
 - RA8000
 - MA8000
 - ESA12000
 - EMA12000
 - Smart Array Cluster Storage

Master Node

The first node to load in a cluster is designated as the master node. The master node assumes the responsibility for maintaining cluster membership and status information.

During installation, NCS 1.6 generates a special cluster resource called the Master IP Address. This IP address runs on the master node and remains with the master node regardless of which server is currently serving as the master node.

Cluster Heartbeat

NCS 1.6 includes a heartbeat statistics tool. The “heartbeat” of a cluster refers to a group protocol that cluster nodes use to transmit packets at regular intervals over the LAN that connects the nodes. These heartbeat transmissions enable nodes to keep track of one another and to detect potential node failures.

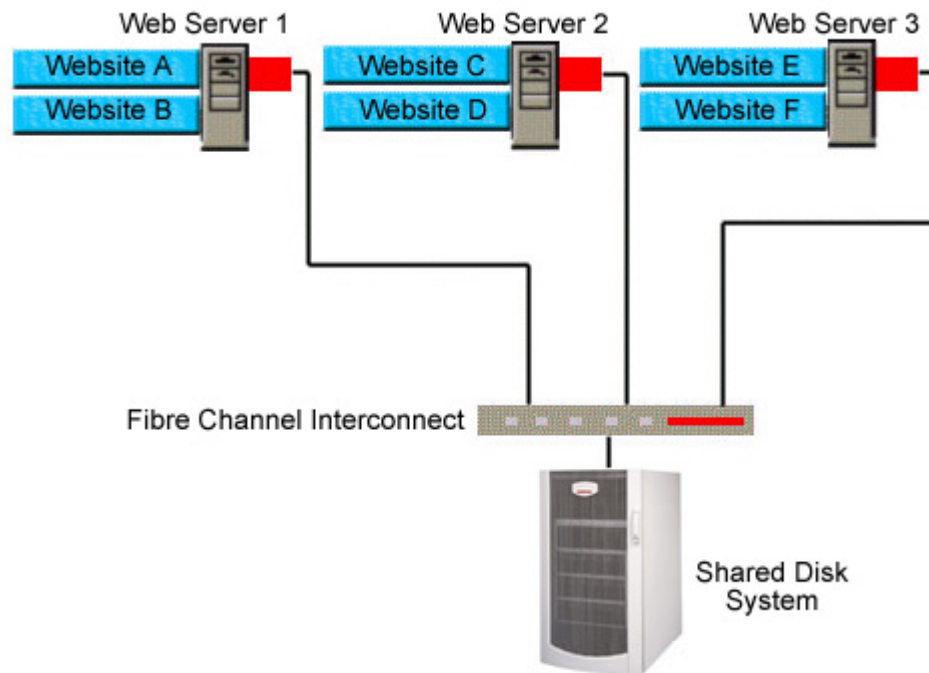
Cluster nodes also use the heartbeat protocol to trigger writing to a special partition on the shared storage array. The data the nodes post to this partition plays a role in avoiding a split brain (discussed later in this module).

The heartbeat statistics provide viewing and tuning of heartbeat settings on both the LAN and SAN.

Example

The eight-second heartbeat threshold on the LAN can be changed. Eight seconds marks the default amount of time that other nodes will wait to receive a heartbeat response packet over the LAN. If they have not received a heartbeat packet within the eight-second period, they will take action.

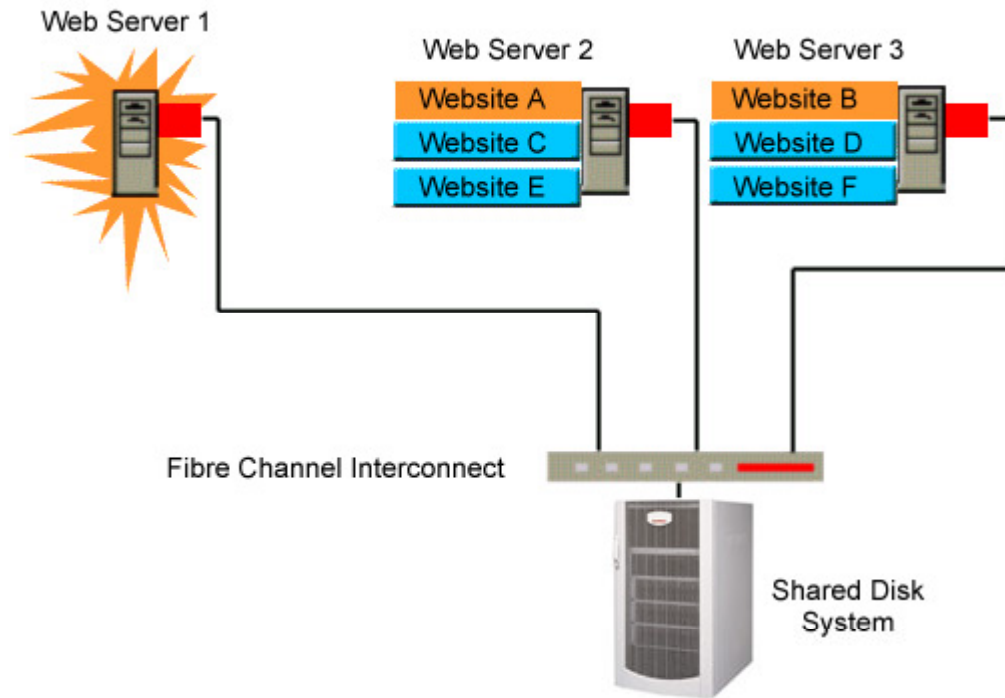
Failover and Failback



All Cluster Nodes Operational

In this example, resources for the intranet are stored on two volumes owned by a single node. Each volume is failed over to a different node when the original node fails.

Fan-Out Failover™



One Cluster Node Failed Using Fan-Out Failover™

When a server failure occurs, cluster resources and services are switched to another server. Failure occurs in one of two ways:


- All resources move from a failed node to a surviving node.
- Resources are distributed among several surviving nodes.

During resource configuration, the cluster is set up to specify if NCS should move all the resources from a failed node to a single surviving node or if NCS should distribute the resources among several surviving nodes. The ability to fail over application resources to multiple servers is known as application fan-out. Novell has trademarked this feature as Fan-Out Failover™.

When the failed server becomes available again, the resources or services that were previously running on the server can be switched back in a process called failback.

Failover and failback can occur automatically or manually. The event often goes unnoticed by users. The failover and failback processes are designed to occur quickly so that users regain access to applications within seconds and, in most cases, without having to log in again. Users typically only notice an hourglass on their monitors for a brief period of time.

Cluster Resources



Cluster Container

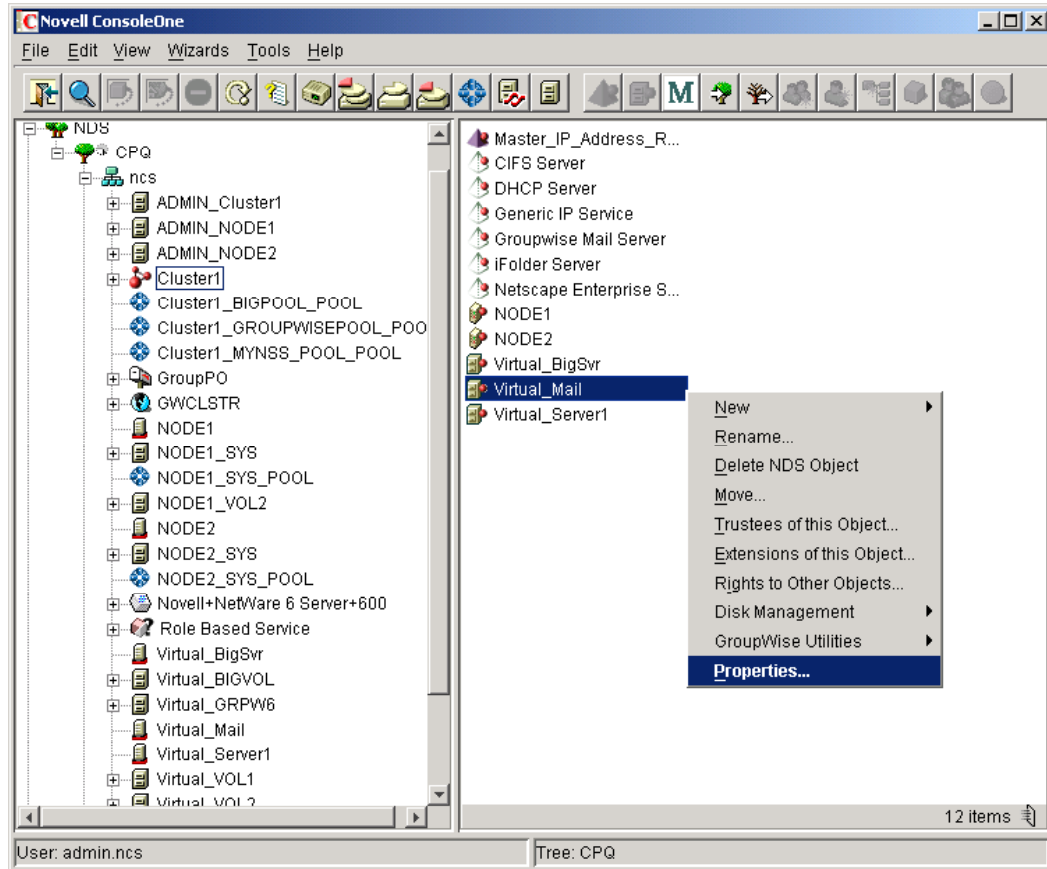
The eDirectory cluster container object contains leaf objects that represent the resources of the cluster. The following objects belong to the cluster container:

- Cluster resources
- Resource templates
- Volume resources
- Cluster nodes
- Master IP Address

The properties of the cluster apply to the cluster as a whole, rather than to each individual node. In effect, the cluster is viewed as a single object and is called a single system image.

The cluster-wide settings, such as membership and communications properties, are configured from a single logical point on the network.

Resource Objects



Cluster resources are defined as either applications or services. A cluster resource must be created for every resource or application that runs on NCS clusters. Every cluster resource is represented in the eDirectory tree by a cluster resource leaf object.

Cluster resources are created for:

- Web servers
- Email servers
- Databases
- Other server-based applications or available services

The cluster resource object holds the configuration information for the applications and services, including the preferred node where each resource should run and the designated standby nodes where resources are transferred in the event of a failure.



Note

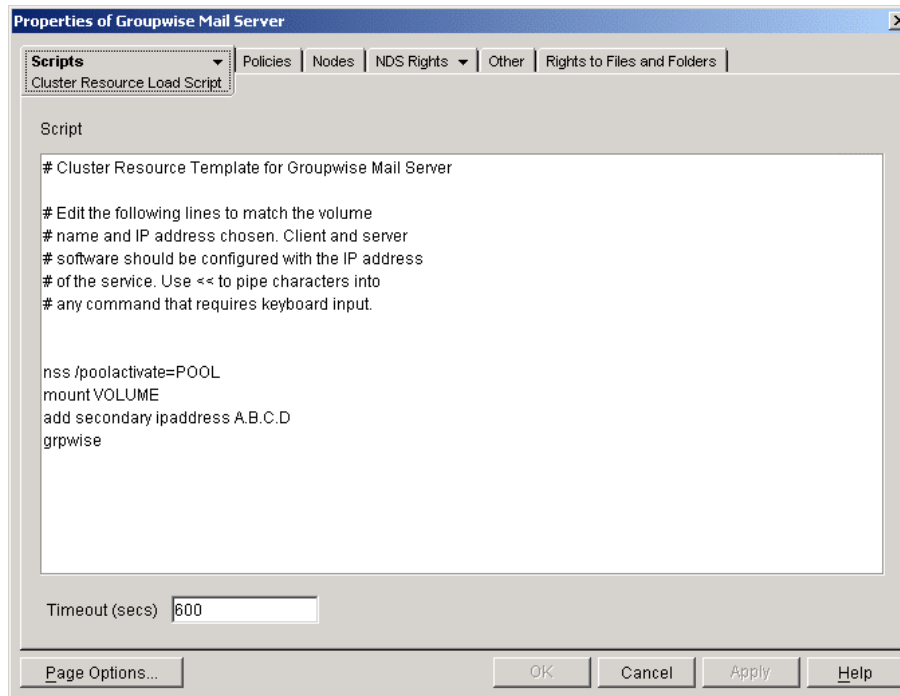
Novell distinguishes volume resources from other cluster resources. A different type of eDirectory object represents clustered volumes.

If a template is being used for the cluster resource, the template performs additional resource configuration automatically. If the resource object is not created using a template, manually configure several cluster resource properties.

The properties of a cluster resource include:

- Load Scripts
- Unload Scripts
- Policies
- Nodes
- NDS Rights
- Rights to Files and Folders

Load Scripts



Sample Load Script

A load script is required for each resource, service, or volume in the cluster. The load script specifies the commands to either start the resource or service on a server or to mount the volume on a server. Any commands used in the .ncf file that runs from the server console can be used in the load script.

The following events must be initiated using a script:

- **Volume mount and dismount** — Commands must be included to mount the volume on the new node and dismount the volume on the recovery node if the object is a cluster-enabled volume:
 - **Mount volume** — NSS/Activate=volume
 - **Dismount volume** — NSS/Deactivate=volume
- **Secondary IP address migration between nodes** — Commands must be included to activate the appropriate virtual IP address on the new node and deactivate the IP address on the recovery node if the object uses an IP address resource.



Note

Load and unload scripts can have no more than 600 characters each. Exceeding this number can cause Cluster Resource Manager (CRM) warning messages when a resource joins or leaves a cluster.

Events in scripts include:

- Database or file system consistency checks
- Addition of newly mounted volumes to the default directory search path of the server
- User notification messaging
- Administrative notification messaging

An unload script can be added to specify how the application or resource should terminate, depending on the requirements of the cluster application or resource. An unload script might not be required by all resources or applications, but it can ensure that a resource unloads before it loads on another node during a failback or manual migration.

If there is no unload script, the resource's secondary IP address will not be released. This prevents the new node from assuming the secondary IP address. Manually forcing the resource to close can cause a comatose state, prohibiting an online state until the original problem is remedied.

Unload Scripts

Properties of Generic IP Service

Scripts | Policies | Nodes | NDS Rights | Other | Rights to Files and Folders

Cluster Resource Unload Script

Script

```
# Cluster Resource Template for Generic IP Service

# Add commands to stop the service here:

# Edit the following lines to match the volume
# name and IP address chosen. Client and server
# software should be configured with the IP address
# of the service. Use << to pipe characters into
# any command that requires keyboard input.

# del secondary ipaddress A.B.C.D
# nss /pooldeactivate=POOL /override=question
```

Timeout (secs)

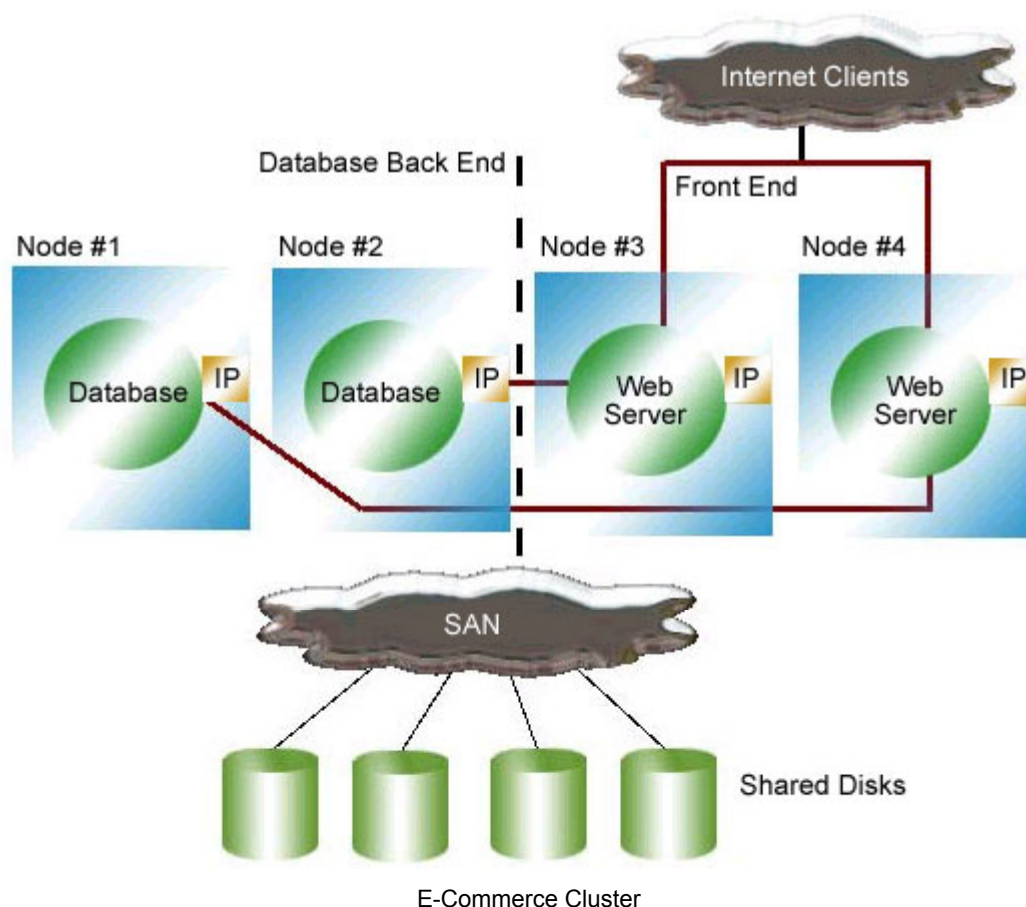
Page Options... OK Cancel Apply Help

Unload Tab Page

Depending on the cluster application or resource, an unload script can be added to specify how the application or resource should terminate. An unload script might not be required by all resources or applications, but it ensures that a resource unloads cleanly before it loads on another node during a failback or manual migration. Consult the application vendor or documentation to determine if unload commands should be added.

The timeout setting for scripts is set on the load script and unload script Properties pages. The default timeout is set to 600 seconds, or 10 minutes. If the script does not complete execution within the time specified, the cluster will terminate the script. If a load script is terminated, the remaining commands in the script will not be executed, and when viewed in ConsoleOne, the resource will show an error condition. If an unload script is terminated, the cluster will proceed to invoke the load script on the failover node.

Sample Use of Load and Unload Scripts



Load and unload scripts are among the most powerful features of NCS and can customize the way NCS handles failover and failback of resources.

Example

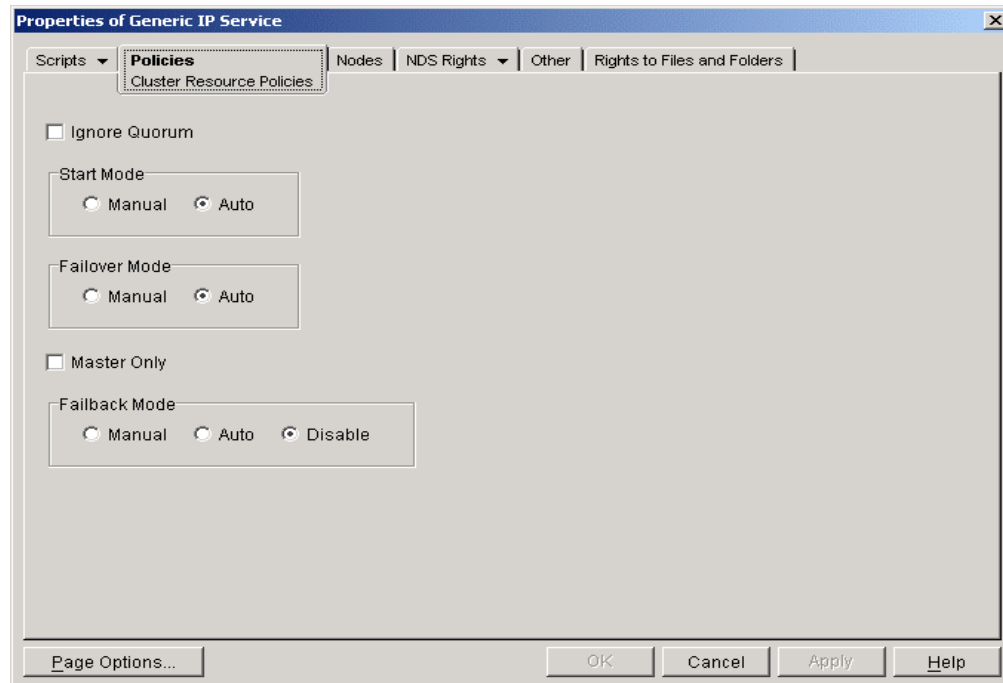
A four-node cluster supports a web-based database application containing the company's shipping and receiving data. The application provides decision support services (DSS), which allow managers to run complex queries.

Two nodes run the web server and the other nodes run the back-end database. The nodes hosting the database server can support both the database and the web server in the event of a failover. The nodes hosting the web server, however, cannot support both services running at full speed.

To prevent the loss of critical services when the nodes hosting the database fail, the system administrator configures the database service load script to detect the presence of the web server running on the same node.

If both services are running on the same node, the load script passes a message to the web application server instructing it to disable the DSS user interface, ensuring that the more critical shipping and receiving functions remain available.

Policies



When a node in the cluster fails, all cluster resources set to fail over automatically migrate to the specified surviving nodes in the cluster without administrator intervention. Configuring a cluster resource to fail back automatically ensures that it will move back to its preferred node after that node rejoins the cluster.

The manual setting forces the cluster to wait for administrative intervention before performing failover or failback, allowing a network administrator to manually control how the resources are migrated. This method might be necessary if the resource cannot be properly started or stopped in the load and unload scripts.

Setting Failover and Failback Modes

Failover and failback of cluster resources are configured to occur manually or automatically from the Policies page.

- **Automatic failover and failback** — To move applications or resources automatically to specified nodes in the event of hardware and software failures, set the failover mode to *Auto*.

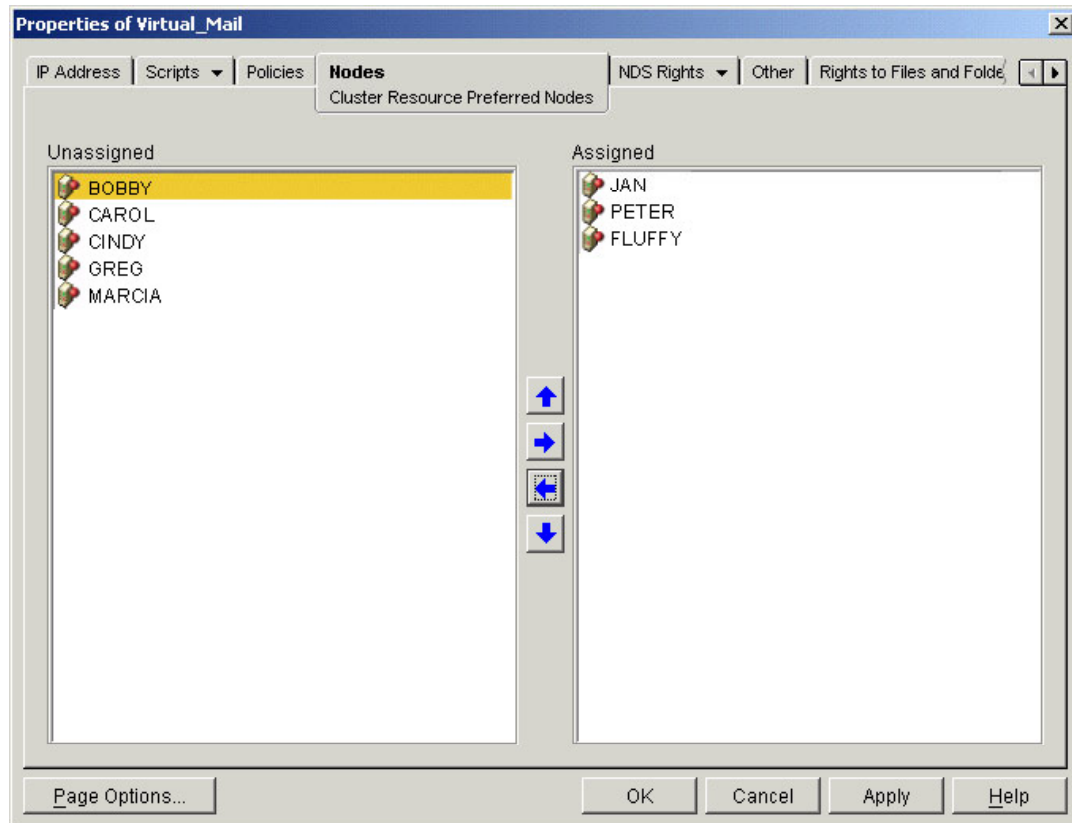
Setting automatic failback ensures that applications or resources automatically move back to their preferred node when hardware or software problems are resolved and the failed node is back online. The preferred node is the first server in the list on the Nodes property page. Using this setting, resources will only automatically fail back to the preferred node.

- **Manual failover and failback** — The Failover mode can be set to *Manual* if intervention is required after a failure occurs and before the resource is moved to another node. Setting the failover mode to manual allows time to restore failed nodes or migrate resources on other nodes before allowing the resource to move.

Manual failback works in much the same way as manual failover. Manual failback prevents a resource from moving back to its preferred node after that node is back online.

Select the *Master Only* node to ensure that a highly available resource runs only on the Master Node.

Nodes



The Nodes page is used to configure the preferred nodes assigned to the cluster resource. The preferred nodes are the nodes on which the resource can be loaded. When resources are created or volumes are added to the cluster, the nodes in the cluster are assigned automatically to the resource or volume.

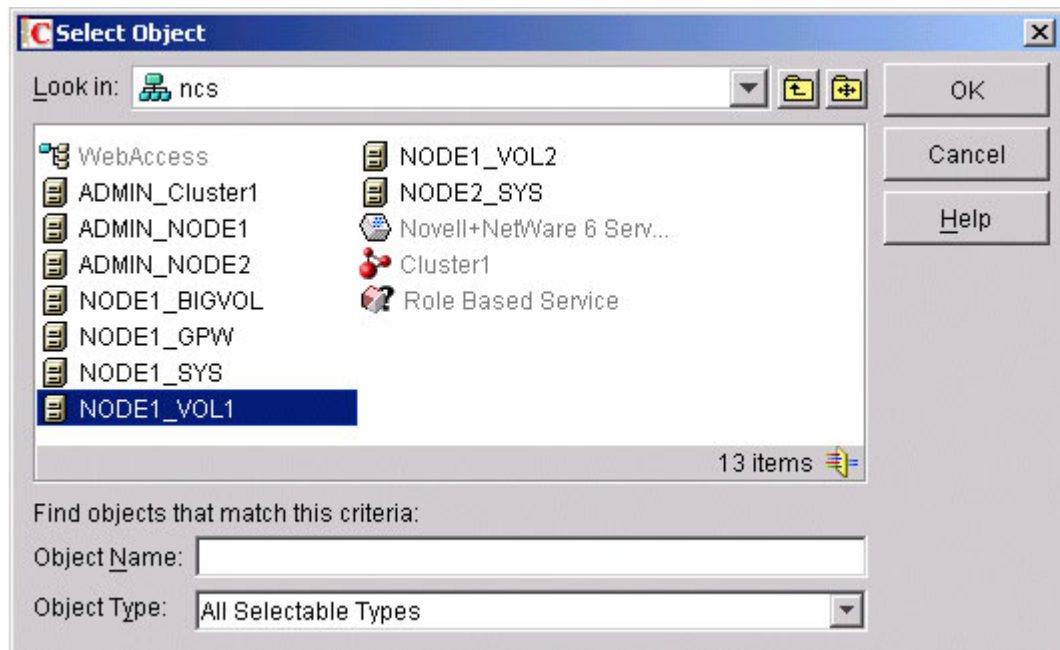
The first node in the Assigned list is the node on which the resource loads when the cluster starts. The order of the remaining assigned nodes determines the failover sequence. A failure of the first node causes the resource to fail over to the next node in the Assigned list, and so forth.

The resource will fail back automatically only to the first node listed on the Nodes property page.

Example

If a cluster resource is assigned to nodes 1, 2, and 3, in that order, and nodes 1 and 2 are both offline, the resource will not fail back from node 3 to node 2 when node 2 comes back online. When node 1 comes back online as well, the resource will fail back to node 1 if the resource has been configured for automatic failback.

Resource Templates



To eliminate the need to repeatedly configure identical properties for similar cluster resources, NCS allows the use of resource templates.

Resource templates can be created for any server application or resource to be added to a cluster. They can be used when configuring server applications to run on a cluster.

The process for creating a cluster resource template is:

1. Set the preferred node assignments.
2. Set the automatic or manual failover and failback modes.
3. Configure the load and unload scripts.

Templates provided with NCS:

- GroupWise Mail Server
- DHCP Server
- Netscape Enterprise Server
- Generic IP Service
- Cluster Resource Load and Unload scripts

Application Resources

Cluster-aware applications can create their own resource templates. These templates allow the administrator to configure application resource parameters such as IP addresses, failover and failback modes, and load and unload scripts. This allows the cluster-aware application to be seen by network clients as a single system image.

Each application creates a resource template for its critical services, for example:

- A database server might create a virtual transaction coordinator resource.
- A web server might create a virtual HTTP server service resource.
- A mail server might create virtual Post Office Protocol 3 (POP3) and Simple Mail Transfer Protocol (SMTP) service resources.

Each resource template contains the following information about the application:

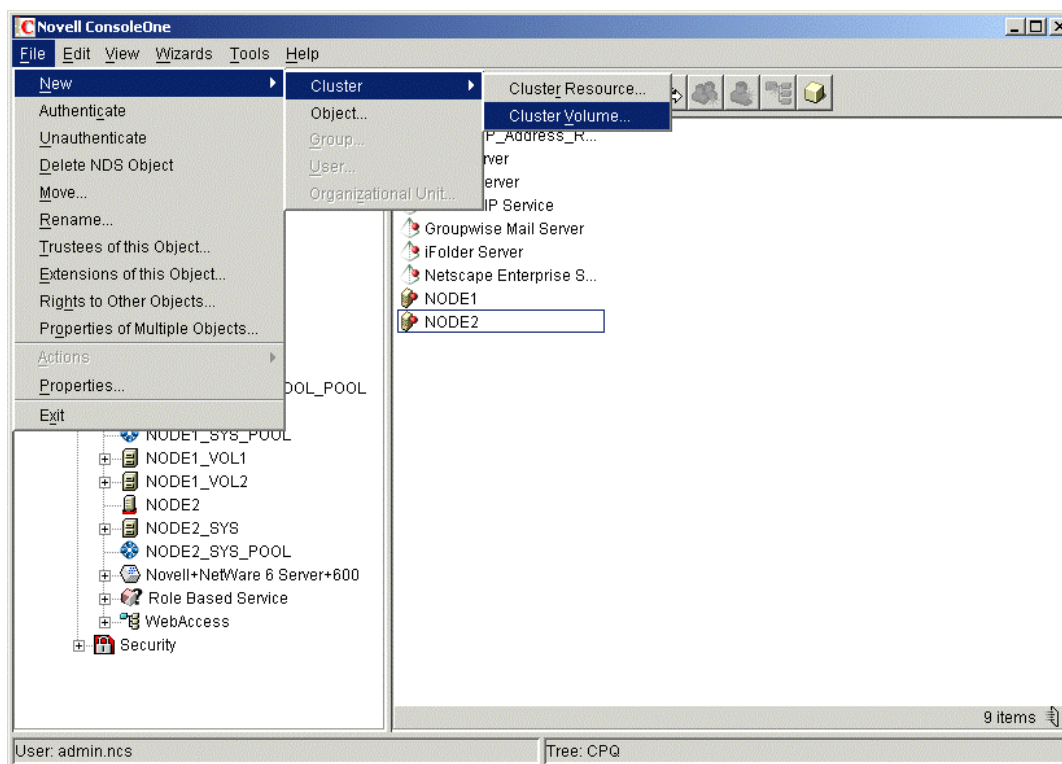
- Virtual IP address
- Load and unload scripts
- Virtual NetWare Core Protocol (NCP) server

Volume Resources

If a shared disk system is part of the cluster and volumes on the shared disk system are to be highly available to NetWare clients, those volumes must be added to the cluster as resources. Adding volumes on the shared disk system to the cluster enables them to be moved or mounted on different cluster servers during failures or when migration is necessary. Server applications that do not require NetWare client access to volumes, do not have to be added to the cluster.

In an NCS shared disk environment, NSS volumes are used. The volumes can be shareable cluster-enabled volumes or regular volumes. Both volume types can fail over to another server node when a node failure occurs, but the method for doing so differs between the two volume types.

Cluster-Enabled Volumes



Cluster-enabled volumes are bound to a virtual server known as a virtual NCP server. Clients attach to the virtual NCP server that acts as a proxy for the actual server hosting the volume. Applications and users do not see a physical server. They see a virtual NCP server and a cluster-enabled volume.

Each cluster-enabled volume has a unique IP address assigned to its virtual NCP server. Rather than connecting to the IP address of a physical server, clients connect to the virtual NCP server through its own IP address.

This configuration allows users to maintain the link to the virtual NCP server and its cluster-enabled volume when it moves from server to server in response to failures.

Virtual server names for cluster-enabled volumes can now be customized. This eliminates DNS confusion.



Note

eDirectory needs to synchronize for an NCP server to display in the directory tree.

With cluster-enabled volumes, users can map drives using the name of the cluster-enabled volume and its virtual NCP server. This feature allows drive mappings to remain valid when volumes fail over to different nodes in the cluster.

This capability:

- Is critical for server applications that require clients to maintain a constant connection to a specified network volume and server.
- Provides location transparency for client connections to the cluster and its shared volumes.

Cluster-enabled volumes are required to achieve transparent client reconnect to a mapped drive or to an application that uses the server name as part of the directory path to a volume. They are not necessary if a transparent client reconnect is not required or if the server applications do not require direct client access to volumes.

Example

Applications that refer to data locations using the server name as part of the directory path must use cluster-enabled volumes. A connectionless database application that uses IP to connect to network clients would not require a cluster-enabled volume.

Because users and applications only see the relationship between the cluster-enabled volume and the virtual NCP server, and not a physical server, they are shielded from the effects of node failures. So when node failures occur, cluster-enabled-volumes can provide uninterrupted service to users and applications.




Note

Auto-reconnect only works for volumes that are cluster-enabled or volumes that are mounted using the load scripts of a cluster resource.

Cluster-enabled volumes have properties similar to other cluster resources. Network administrators can specify the nodes that will be eligible to mount a cluster-enabled volume as well as its preferred node. The cluster-enabled volume must also be configured with its own load and unload scripts. Additionally, network administrators must set the failover and failback modes of the volume resources.

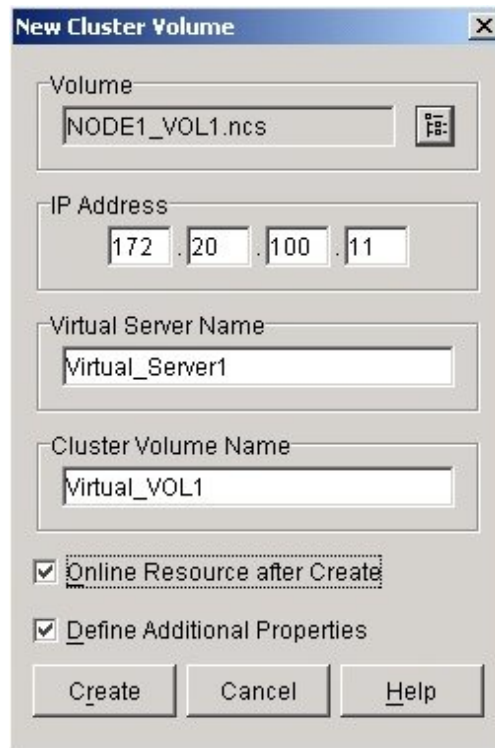
Cluster Volume Failover



Cluster Volume Failover

When a node fails, the cluster-enabled volume remounts automatically on a surviving node and the virtual NCP server then functions as a proxy for the new hosting server. In most cases, clients will not know that a failover has occurred because NSS volumes mount quickly.

New Cluster Volume

The image shows a Windows-style dialog box titled "New Cluster Volume". It contains several input fields and checkboxes. The "Volume" field is a text box containing "NODE1_VOL1.ncs" with a browse button to its right. The "IP Address" field consists of four separate boxes containing "172", "20", "100", and "11". The "Virtual Server Name" field is a text box containing "Virtual_Server1". The "Cluster Volume Name" field is a text box containing "Virtual_VOL1". Below these fields are two checked checkboxes: "Online Resource after Create" and "Define Additional Properties". At the bottom are three buttons: "Create", "Cancel", and "Help".

New Cluster Volume

Volume
NODE1_VOL1.ncs

IP Address
172 . 20 . 100 . 11

Virtual Server Name
Virtual_Server1

Cluster Volume Name
Virtual_VOL1

☒ Online Resource after Create

☒ Define Additional Properties

Create Cancel Help

Three eDirectory objects are created for a cluster-enabled volume:

- **Cluster Volume** — Represents the cluster-enabled volume.
- **Cluster Volume Resource** — Contains policies and load and unload scripts for the cluster-enabled volume.
- **Cluster Virtual Server** — Represents the virtual NCP server for the volume.

The Cluster Volume Resource parameter is similar to a generic Cluster Resource. It specifies preferred nodes, load and unload scripts, and failback modes.

Regular Volumes

Transparent client reconnection to a mapped drive requires cluster-enabled volumes. Because regular volumes do not create eDirectory objects or virtual NCP servers, clients attach to the volume through their owner node. If the node fails, the volume object is inaccessible.

Transparent reconnection to client/server applications that do not depend on drive mappings to a specific server can be achieved by tying the volume to the application resource using scripts.

Regular volumes can be migrated along with the application resource by entering the appropriate volume mount and dismount commands in the resource load and unload scripts of the application. In this situation, migration of regular volumes is controlled directly by the applications to which they are tied.

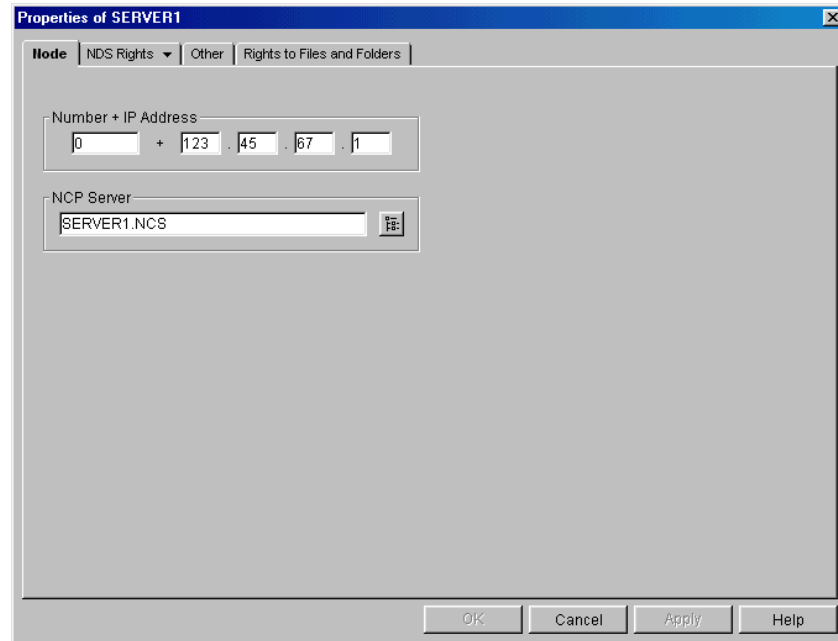
Any number of regular volumes can be tied to an application resource using scripts.

Many applications use regular volumes to store their data and executable files.

Example

Netscape web servers use regular volumes to store local data.

Cluster Nodes



Use the Cluster Node Properties page to view or edit the cluster node number or IP address of the selected node or to view the context for the NetWare Server object.

To view or edit cluster node properties in ConsoleOne, select the *cluster object*, right-click the desired cluster node on the right side of the ConsoleOne display screen, and select *Properties*. When the Cluster Node Property page displays, select the *Node* tab.

Number + IP Address

Number + IP Address field specifies the cluster node number and IP address for the selected node. If the cluster node number or IP address changes for the selected node, the new information is not automatically updated in eDirectory. Edit the information and click *Apply* to update the information in eDirectory.

NCP Server

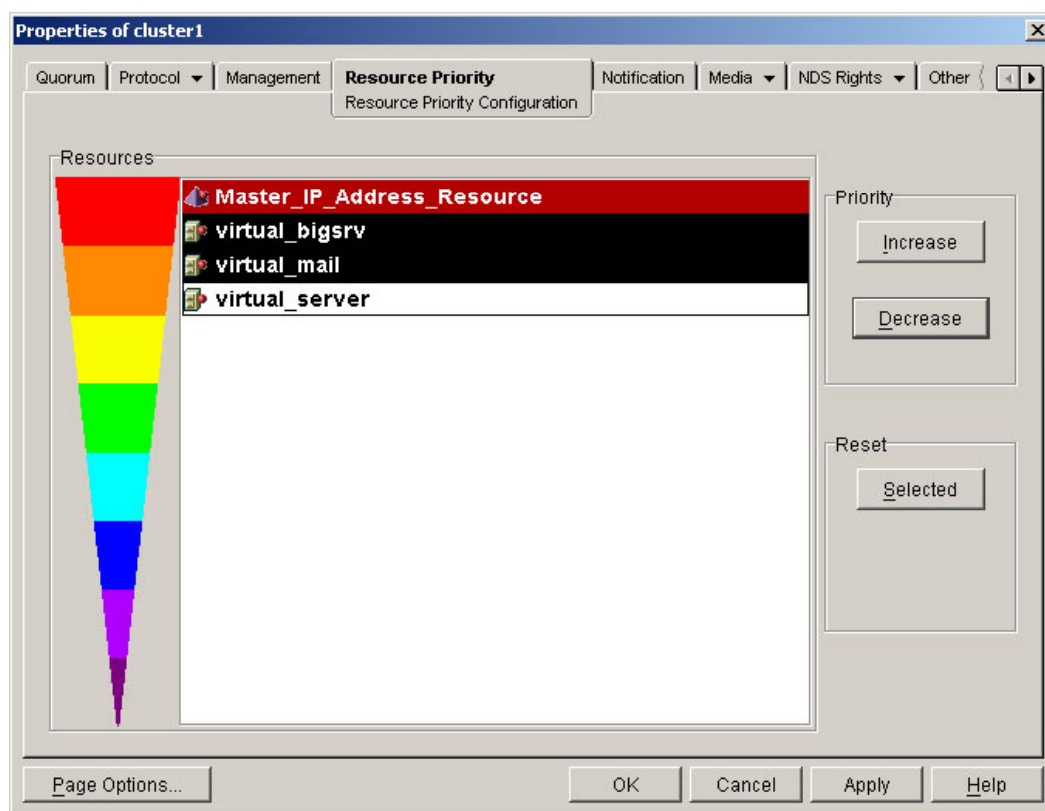
The NCP Server field displays the context for the NetWare Server object. This field cannot be edited.



Note

All servers in a cluster must be in the same eDirectory tree, must be configured with TCP/IP, and must be on the same IP subnet.

Resource Priority



Resource Priority controls the order in which multiple resources start on a given node when the cluster is brought up or during a failover or failback. For example, if a node fails and two resources fail over to another node, the resource priority determines which resource loads first.

This feature ensures that the most critical resources load first and are available to users before less critical resources.

To view or change resource priorities using ConsoleOne:

1. Right-click the cluster object.
2. Select *Properties*.
3. On the Cluster Object property page, select the *Resource Priority* tab.
4. To change the priority for a resource, select the resource in the list and then click the *Increase* or *Decrease* button to move the resource up or down in the list. This feature changes the load order of the resource relative to other cluster resources on the same node. You can also select a resource and then click *Selected* to reset the resource back to its default load order.
5. Click *Apply* to save changes made to resource priorities.

eDirectory Cluster Objects

NCS leverages the power and extensibility of eDirectory to configure, manage, and monitor a cluster. All cluster configuration information including failover and failback policy and cluster resource properties reside in the eDirectory database. During the installation of NCS, the eDirectory schema is extended to include the cluster container object.

When network clients connect to the cluster through the directory services protocol, the integration of directory services enables:

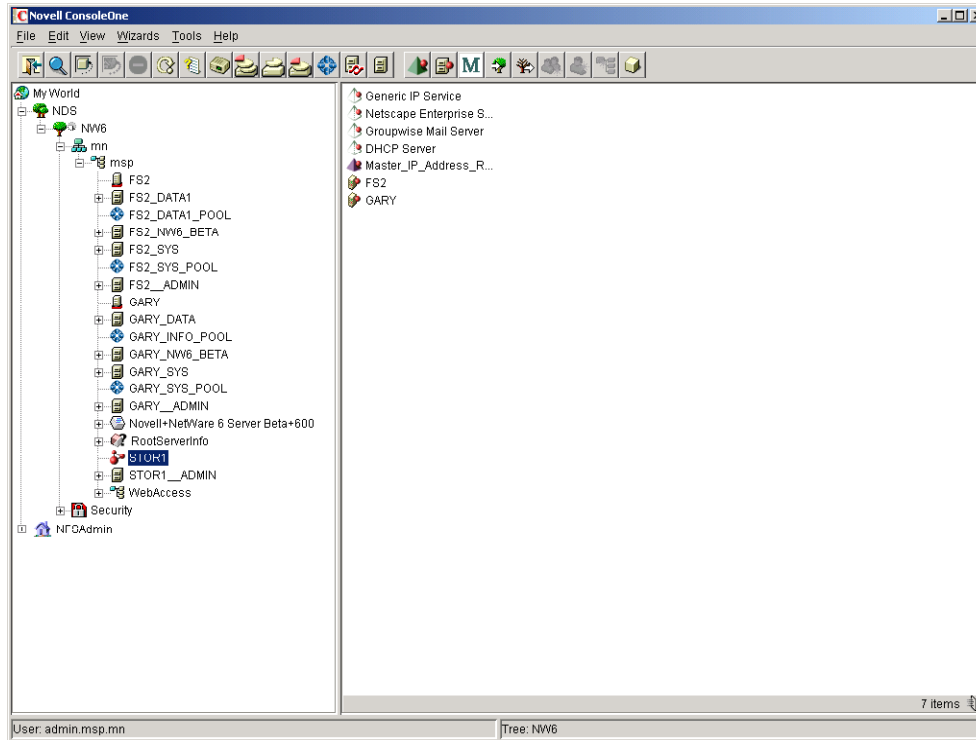
- Transparent client reconnect following failover.
- Consolidated network management using standard eDirectory management tools.

The cluster-related eDirectory objects are administered through either ConsoleOne or NRM. ConsoleOne is a Java-based GUI application running on a client workstation on the network, while NRM can be accessed from any web browser and does not require NetWare Client32 software.

Some server console commands include:

- View
- DHCP {context}
- Stats {display, clear}
- Cvsbind {add, del} {resource} {IP address}

ConsoleOne Properties



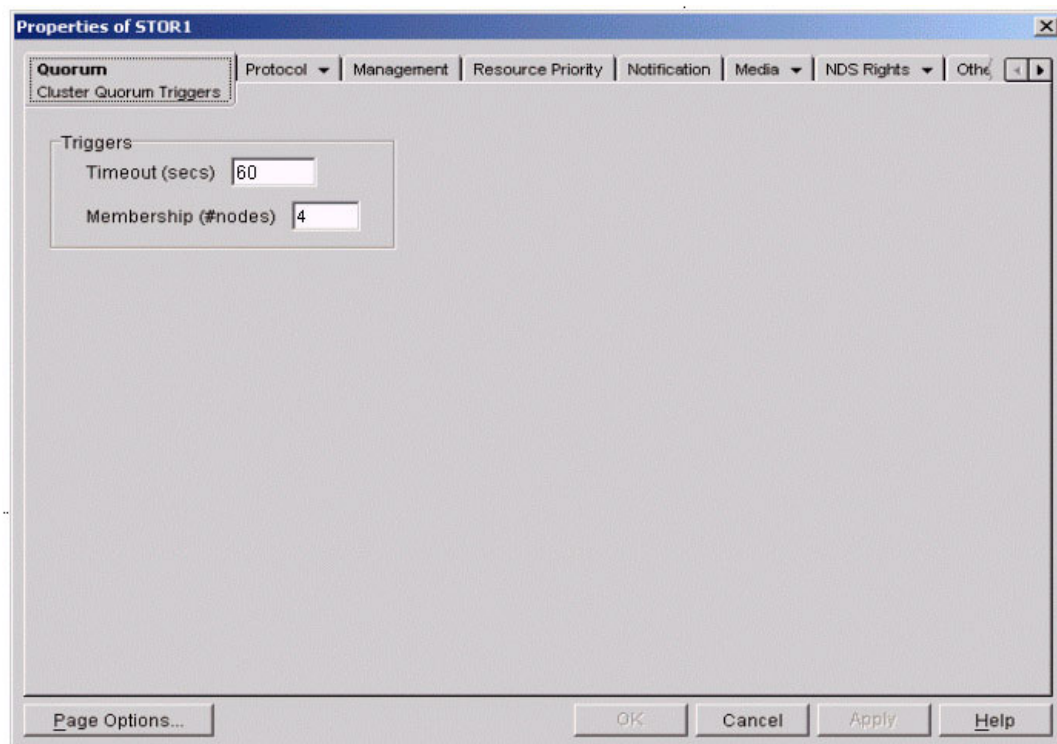
ConsoleOne Properties

The ConsoleOne properties dialog for the cluster container divides the properties into the following pages:

- Quorum
- Protocol
- Management
- Resource Priority
- Notification
- Media
- NDS Rights
- Other
- Rights to Files and Folders

The NDS Rights, Other, and Rights to Files and Folders tabs will not be discussed in this course.

Quorum Membership



The Quorum membership is the number of nodes that must be running in the cluster before resources start to load. When first bringing up servers in the cluster, NCS reads the number specified in the Membership (# nodes) field and waits until that number of servers is up and running in the cluster before it starts loading resources. Set the membership value to a number greater than one so that all resources do not automatically load on the first server that is brought up in the cluster. By default, this number will equal the total number of nodes currently defined in the cluster.

To edit the Quorum page Membership (# nodes) and Timeout (secs) properties in ConsoleOne, right-click the cluster object and select *Properties*. When the Cluster Object Property page displays, select the *Quorum* tab.

Example

In the preceding graphic, four servers must be running in the cluster before any resource will load and start.

Timeout

Timeout specifies the amount of time to wait for the number of servers defined in the membership field to be up and running. If the timeout period elapses before the quorum membership reaches its specified number, resources will start loading automatically on the servers that are currently running in the cluster.

Example

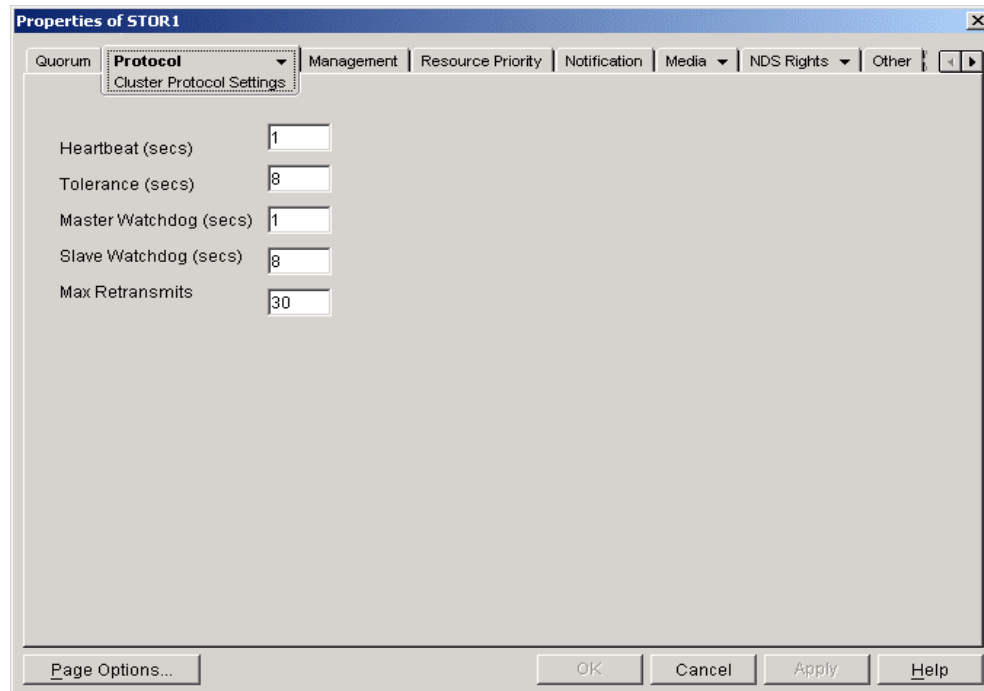
If a membership value of 4 is specified with a timeout value equal to 30 seconds, and after 30 seconds only two servers are up and running in the cluster, resources will begin to load on the two servers that are running in the cluster.

To illustrate what would happen if these parameters were set too low, consider a situation where a five-node cluster is restarted. The Quorum page Membership (# nodes) parameter is set to one and the Quorum Timeout (secs) parameter is set to one minute. One of the servers starts and forms the cluster before the other two have completed initialization and the time specified in the Quorum Trigger timeout expires.

NCS assumes that the other four servers have failed and begins to load the resources owned by the other servers. If the first server does not have enough resources to support the applications owned by the entire cluster, it might fail. In addition, when the other servers come online and failback is enabled, the resources must be failed back to their owners, causing a further interruption of services.

The quorum timeout trigger is set to 60 seconds by default. This parameter can be increased to 600 seconds, or 10 minutes, which would ensure that all the server nodes are up before resources begin initializing.

Protocol Settings



Use the Cluster Protocol Properties pages to view or edit the transmit frequency and tolerance settings for all nodes in the cluster, including the master node. The master node is the first node brought online in the cluster, but if that node fails, any of the other nodes in the cluster can become master.

To view or edit Cluster Protocol properties, in ConsoleOne, right-click the cluster object and select *Properties*. When the Cluster Object Property page displays, select the *Protocol* tab. This tab has two pages, Internals and Settings. The Internals page allows viewing of the script used to configure the cluster protocol settings. Use the Cluster Protocol Settings page to make changes to cluster protocol properties.

Setting Cluster Protocol Properties

Use the following properties to view or edit the transmit frequency and tolerance settings for all nodes in the cluster, including the master node.

The default values are:

- **Heartbeat (seconds)** — The amount of time between transmits for all nodes in the cluster except the master.

For example, setting this value to one means that nonmaster nodes in the cluster will send a signal that they are “alive” to the master node every second.

The default is one second.

- **Tolerance (seconds)** — The amount of time the master node waits for all other nodes in the cluster to signal that they are alive.

For example, setting this value to four means that if the master node does not receive an “I’m alive” signal from a node in the cluster within four seconds, that node will be removed from the cluster.

The default is eight seconds.

- **Master Watchdog (seconds)** — The amount of time between transmits for the master node in the cluster.

For example, if you set this value to one, the master node in the cluster will transmit an “I’m alive” signal to all the other nodes in the cluster every second.

The default is one second.

- **Slave Watchdog (seconds)** — The amount of time the master node has to signal that it is alive.

For example, setting this value to five means that if the nonmaster nodes in the cluster do not receive an “I’m alive” signal from the master node within five seconds, the master node will be removed from the cluster and one of the other nodes will become the master node.

The default is eight seconds.



Note

Max Retransmits is not currently used with NCS, but will be used in future versions.

The default settings typically do not have to be changed, but a number of noncritical factors can cause missed heartbeats, resulting in false failover:

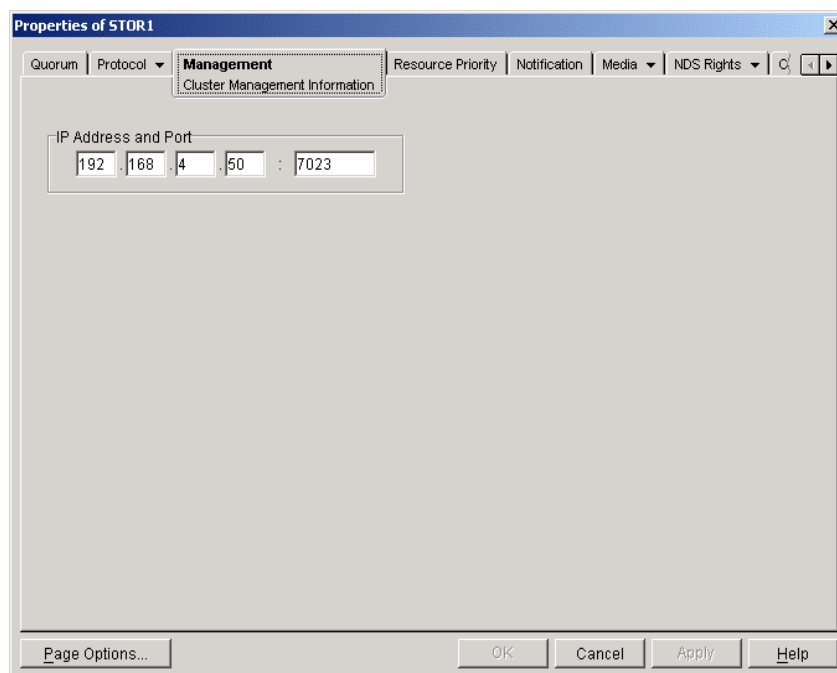
- If the cluster does not have a dedicated interconnect, or if the dedicated interconnect is severed, the heartbeat will be sent over the client LAN. Excessive traffic on the client LAN can cause dropped packets.
- A bad link or switch on the interconnect network can cause intermittent network errors.
- Excessive I/O or processor utilization can cause a node to fail to send out the heartbeat on time.

The cluster nodes should be monitored closely when:

- The cluster is first brought online.
- Services are added to the cluster.
- Usage of cluster services rises significantly.

If false failovers occur, verify the cluster Protocol Settings parameters before beginning other diagnostic procedures.

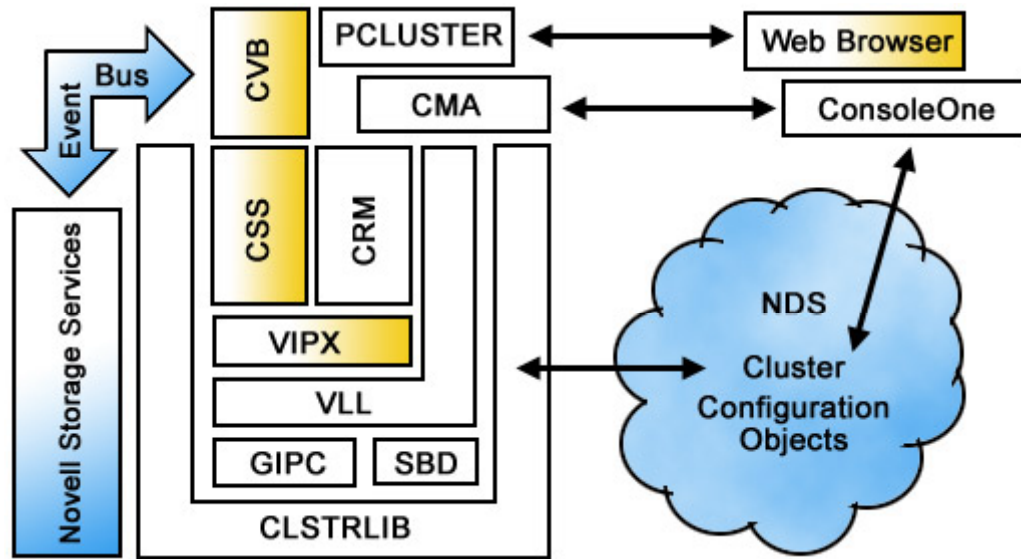
Cluster Management Port



The cluster management port specifies the TCP/IP port number that the cluster management tools use to connect to the cluster. The default cluster port number is 7023 and is automatically assigned when a cluster is created. The cluster port number does not need to be changed unless a conflict is created by another resource using the same port number. If there is a port number conflict, change the port number to any other value that does not cause a conflict.

To view or edit the cluster port property, in ConsoleOne, right-click the cluster object and select *Properties*. When the Cluster Object Property page displays, select the *Management* tab. If the default IP port setting is in potential conflict with an existing port, it can be altered in the IP address and port field in this window.

System Architecture



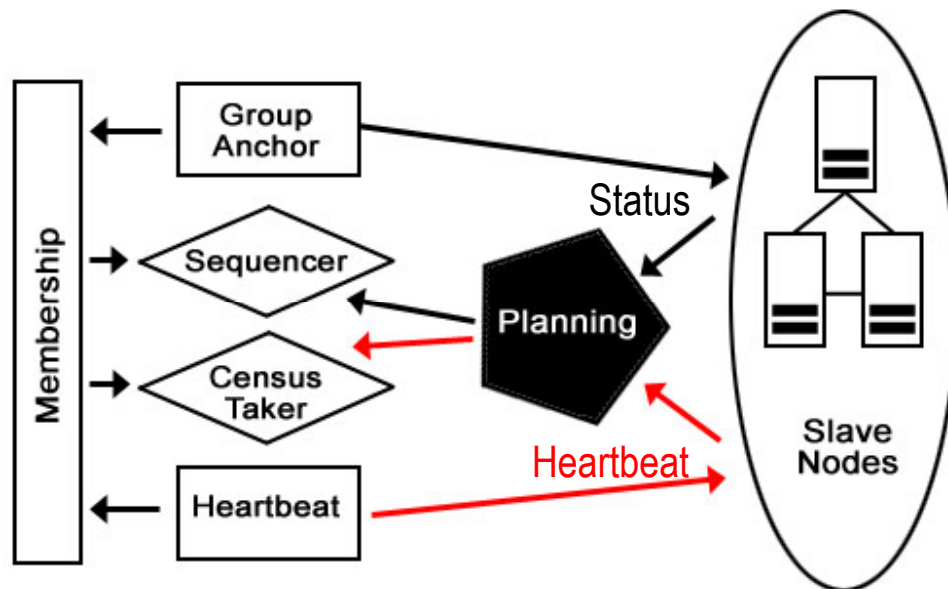
Cluster Configuration Library

The Cluster Configuration Library (CLSTRLIB.NLM) stores NDS cluster data. CLSTRLIB.NLM is the interface between eDirectory and NCS. The cluster objects in the eDirectory database contain configuration information, properties, and cluster settings. The master node in a cluster accesses eDirectory and stores the information from the cluster objects in local memory. The master node then sends this information to the other nodes in the cluster so they do not need to query eDirectory. The other cluster library NLMs on each node then access the data, now stored in local memory to perform their tasks.

When changes are made to cluster objects in eDirectory, the Cluster Configuration Library updates the local version of the configuration data on the master node. The master node sends the configuration updates to the slave nodes so the slave nodes can update their local copy of the data.

Each node maintains a local record of the history of the cluster. This record is updated whenever a node joins or leaves the cluster. Each time a change is recorded, the new configuration is saved and assigned an epoch number. The epoch number resolves situations where the cluster nodes cannot agree on the current configuration.

Group Interprocess Protocol




The Group Interprocess Protocol (GIPC) module (GIPC.NLM) runs the group membership protocols of the cluster. It controls the heartbeat protocol and tracks changes that affect group membership.

The master node uses the GIPC to track cluster membership because the GIPC guarantees an accurate view of the membership of nodes in the cluster.

The group membership protocols are:

- **Planning** — Filters messages to the current epoch.
- **Heartbeat** — Generates point-to-point heartbeat messages between the nodes.
- **Sequencer** — Enforces sequencing of multicasts and membership changes between nodes.
- **Membership** — Maintains stable membership information.
- **Census Taker** — Generates unstable membership reports.
- **Group Anchor** — Communicates group membership information to nodes.

Cluster Resource Manager



The Client Resource Manager NLM (CRM.NLM) keeps track of where cluster resources are running. It is a layer above the group membership protocols ensuring that as membership changes occur, cluster resources run on the correct nodes in the cluster.

The CRM of the master node tracks the real-time status of all resources running on the cluster. It receives notification from GIPC when a node leaves or joins the cluster and initiates failover or failback according to the policies specified for that resource in the eDirectory database.

The CRM running on the master node records the current state of all resources running on the cluster. This information is distributed to all nodes in the cluster. If the master node fails, a CRM on a surviving node becomes the new master and manage the cluster resources.

The various states that the CRM tracks for cluster resources are shown in the following table.

Resource State	Description
Unassigned	The cluster resource cannot be started because the preferred nodes are not currently members of the cluster.
Offline	The resource is dormant. The properties of the resources can be edited.
Loading	The load script is running to activate the resource.
Unloading	The unload script is running to deactivate the resource.
Comatose	The resource failed to complete the script before the timeout occurred.
Running	The resource has been located and is running.
Ask-Load	The resource is waiting for administrative confirmation of manual failover.
Ask-Unload	The resource is waiting for administrative confirmation of manual failback.
Alert	The resource requires manual intervention.
Sync	The resource is waiting to synchronize with its eDirectory properties.
Quorum Wait	The resource is waiting to activate until enough nodes join the cluster to meet quorum trigger requirements.

Cluster Management Agent

The Cluster Management Agent (CMA.NLM) acts as a proxy for the ConsoleOne management application. It works with ConsoleOne to display and control the state of the cluster, and to migrate and configure cluster objects.

Specifically, the CMA allows ConsoleOne to:

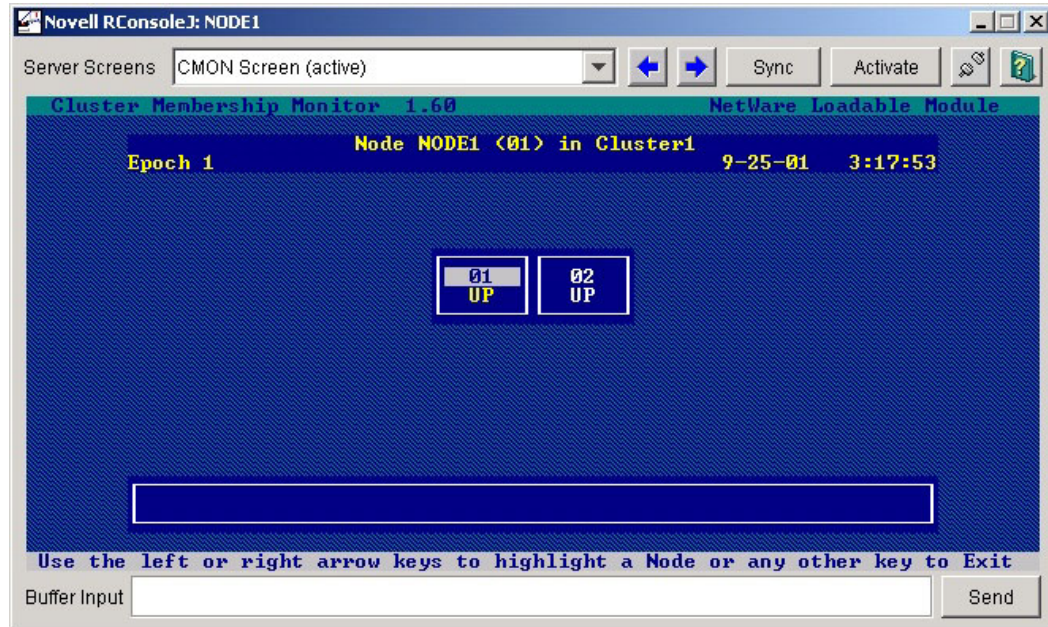
- View the current location of cluster resources.
- Add, move, or remove cluster resources.
- Create or modify load scripts, unload scripts, and script timeout parameters.
- Take resources offline or bring them online.
- Specify failover and failback policies for cluster resources.
- Invoke manual failover and failback.

CMA is installed when NCS is first installed.

Portal Cluster


The Portal Cluster (PCLUSTER) is the NRM snap-in module that interfaces with NRM to manage a cluster from any computer with a browser and Internet access. Any task done in ConsoleOne can be done using NRM.

Cluster Membership Monitor



Cluster Membership Monitor (CMON.NLM) is a console utility that runs on each node in the cluster, allowing network administrators to view the status of the cluster's nodes.

Split Brain Detector



The Split Brain Detector (SBD) module (SPD.NLM and SBDLIB.NLM) detects and protects against split brain faults that can occur when a node loses its Ethernet connection. This node is then unable to transmit or listen for heartbeat packets over the LAN.

The other cluster nodes assume that the silent node is “dead” and attempt to take over the resources of the presumably “dead” node.

However, because the node is not actually dead (but has only lost its Ethernet connection), the node thinks it is the only node “alive” and thus attempts to restart the cluster resources single-handedly.

The SBD module detects such situations and notifies the cluster, which acts immediately to “kill off” one side of the split brain. The cluster “kills off” either the smaller half of the split brain or the half that is not running the master node.

In two-node scenarios, however, neither node is smaller, and each thinks it is the master node. To address this potential problem, NCS 1.6 can detect LAN failure enabling the cluster to protect against split brains in two-node clusters by “killing off” the node that lost its Ethernet connection.

Restoring an SBD Partition

To perform an SBD partition restoration:

1. Use ConsoleOne or NRM to delete the old SBD if necessary.
2. Unload all NCS NLMs using `uldnscs.ncf`.
3. Load `clstrlib` and `vll` manually.
4. Enter the server console command `SBD install`. Proceed through the instructions to create a new SBD partition. A prompt for a device list displays asking whether to create a mirrored SBD partition.

Globally Unique Identifiers

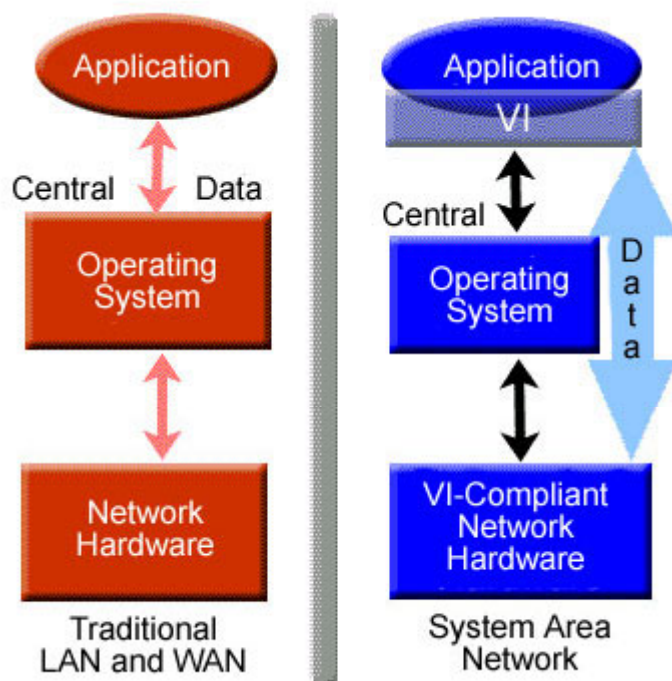
NetWare 6 uses Globally Unique Identifiers (GUIDs) that remain the same across all servers on which the user is a trustee. User space restrictions are enforced in NetWare 6 after a failover. Because there is no longer a need to translate trustee IDs, volume failovers occur much quicker.

This also makes backups simpler because it eliminates the issue of restoring data to the same node from which it was backed up to avoid corrupting trustee IDs. With NetWare 6, data backed up from one node can be restored to another node. This feature improves NDS cluster integration.

**Note**

TRUSTMIG is no longer needed because there are no trustee assignments to migrate. NetWare 5.1 and previous versions used trustee IDs that were a 32-bit identifier unique for each server. As a result, items such as drive space restrictions did not transfer after a failover because it took too long to calculate the ID.

VIPL Extensions



The Virtual Interface Provider Library (VIPL) Extensions (VIPX) module is a Novell extension of the provider library for the Virtual Interface Library (VI) Architecture specification. The VI Architecture specification defines an industry-standard architecture for communication within clusters of servers and workstations. Novell VIPX makes it easier to use the standard VI Provider Library to write programs for NCS 1.6.

Its application program interface (API) is the VIPX. NCS uses the extended APIs for intracluster communication. The Virtual Link Layer (VLL) module passes messages between the other NCS NLMs. By supporting the VIPX APIs, NCS enables future integration with third-party cluster software.

INTERNET More information on the VI Architecture specification can be found at:
www.viarch.org

Cluster System Services

The Cluster System Services (CSS) is for shared memory and distributed locks. The CSS module provides a generic API that any distributed, cluster-aware application can use to enable distributed-shared memory and distributed locking. Distributed-shared memory allows cluster-aware applications running across multiple servers to share access to the same data as though the data were on the same physically shared RAM chips.

Distributed locking protects cluster resources by ensuring that if one thread on one node gets a lock, another thread on another node cannot get the same lock. For example, when one node activates a shared pool, the pool releases its lock to that node. When another node attempts to activate the same pool, that node cannot do so because the pool has already released its lock.

Cluster Volume Broker

The Cluster Volume Broker (CVB) is the single-system image storage manager. The CVB module is a cluster-aware application that enables changes to an NSS volume on one node only, after which CVB distributes that change across all nodes. In other words, CVB ensures that each node in the cluster has the same image of the storage situation at all times.

CVB also checks for the SBD partition. If the SBD partition is not present, the CVB will not allow the Shareable for Clustering flag to be set because it thinks that the cluster is no longer present.



Note

If the CVB is corrupt and there are problems with the cluster flag, it will be necessary to unload the CVB to flag the pool as shareable. Next, update the eDirectory from the properties page, then reload the CVB.

Installing or Upgrading NCS



The NCS installation program must be run when:

- Creating a new cluster.
- Adding new nodes to an existing cluster.
- Upgrading NCS software in an existing cluster.



Note

Because of changes within the NSS 3.0 file structure, cluster servers must be prepared before upgrading them to NetWare 6 and NCS 1.6. To prepare cluster servers for upgrade, run the NetWare Deployment Manager (nwdeploy.exe).

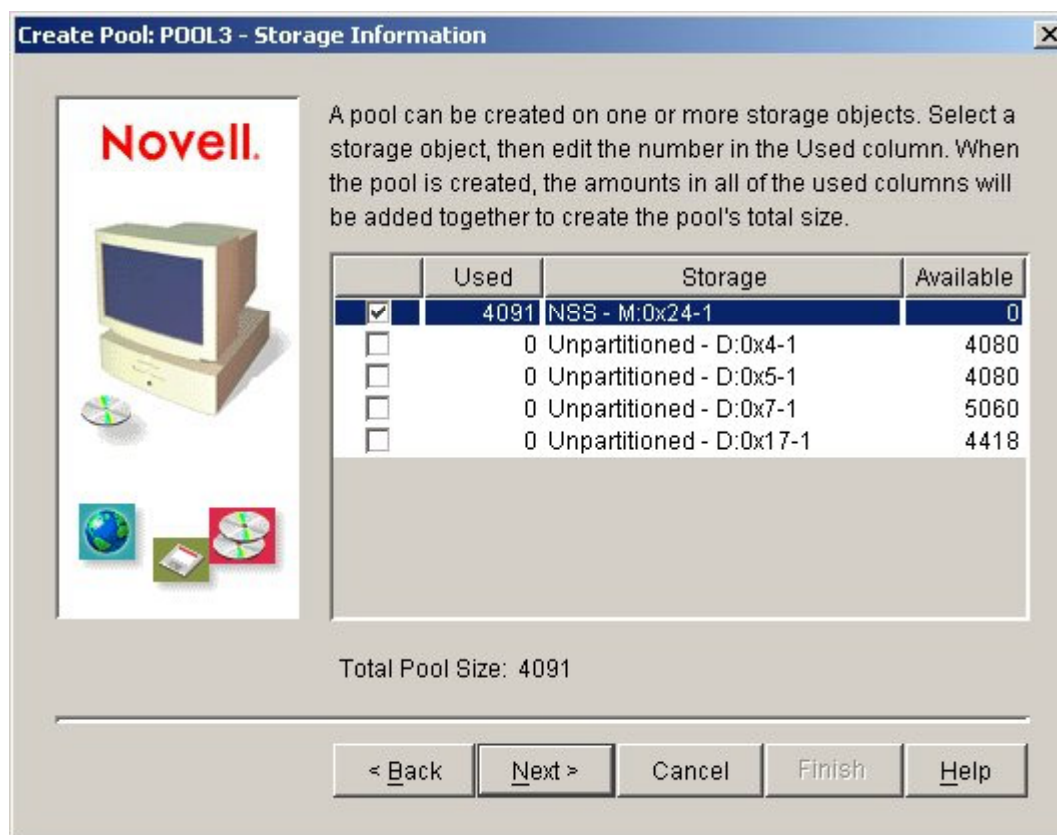
The NCS software includes:

- NCS
- Novell Cluster Services licenses (required only for clusters having more than two nodes)
- Management client running ConsoleOne or NRM

The installation steps are:

1. Ensure that you use the latest Compaq Support Paq (cpqdpoy.nlm).
2. Configure the shared storage with ConsoleOne or NRM.
3. Install the NCS software.
 - a. Execute nwdeploy.
 - b. Assign the cluster Master IP Address.
4. Add NSS-based pools and volumes to the cluster.
5. Create or select Cluster Resource Templates.
6. Create or modify Cluster Resources.
 - a. Configure the load scripts.
 - b. Configure the unload scripts.
 - c. Set the failover and failback modes.
 - d. Assign the preferred nodes.
 - e. Set the resource priority.


Creating Storage Pools and Volumes



An NSS pool is a point of management. NSS is the broker between the local device and the shared storage device. Storage pools are containers for logical volumes.

The storage pool feature in NSS provides more flexibility in planning and configuring shared storage to work with NCS. Multiple cluster-enabled volumes can now be part of a single cluster resource, and volumes can dynamically grow as needed to take advantage of free disk space.

Use ConsoleOne to create storage partitions and pools, to define logical volumes within the pools, and to cluster-enable the volumes.



Storage Area Network Devices (JBOD Disks or RAIDsets)

When NSS detects a Sharable for Clustering flag, NSS will not activate the storage pools running on that device unless NCS is also running. (By default, NSS activates the pools on only local devices.)

Storage pools must be either purely shared or purely local; that is, the devices or portions of devices used to create a storage pool must be either all shared devices or all local devices.

Logical volumes are beneficial because they expand and shrink like balloons, using only as much space as the files and directories within them consume at any given moment. NCS provides protection for data in shared storage systems at the storage pool level. For example, when NSS requests permission to activate a storage pool on a cluster node, NCS grants or denies permission to do so. The rule on which NCS bases this decision is simple: the cluster allows a shared storage pool to be active on only one cluster node at a time.

Hardware and Software Requirements

The following requirements must be met for a successful NCS installation:

- Two or more NetWare 6 servers in the cluster, up to 12 servers per cluster
- ProLiant servers qualified for NetWare 6 (Pentium III at 700MHz and higher)
- 256MB minimum RAM
- 512MB minimum RAM recommended for failing over multiple applications
- IP protocol configuration and location on the same IP subnet
- All servers configured with the IP protocol on the same IP subnet
- All server nodes located within same eDirectory 8.6 tree
- NSS 3 installed
- Minimum of one local disk device (not shared) for a SYS volume residing on each server
- Novell Client 4.8 for Windows NT and Windows 2000 or Novell Client 3.3 for Windows 95 and Windows 98 installed on the workstations managing the cluster
- ConsoleOne 1.3.2 installed on the workstation used to manage the cluster or access provided to a browser to launch NRM
- Minimum of two eDirectory replicas and maximum of six for the partition containing the cluster



Note

More than six eDirectory replicas in the partition will reduce the performance of eDirectory.

- Minimum 300MHz processor and 90MB of memory on the client system used to manage the cluster
- Novell Licensing Services (NLS) installed before beginning the NCS installation

A shared disk system is required for each cluster. Ensure the following:

- Minimum of 15 to 25MB of free disk space is available for the special cluster partition. The 15 to 25MB of free space must **not** be a part of an NSS partition.
- Shared disk system is properly configured and functioning.
- Shared disk volumes are configured to use NSS.
- Shared storage subsystem disks are configured to use a RAID configuration.

**Warning**

Disks not configured with RAID fault tolerance can cause a system failure. NCS will not protect against this type of fault.

**Note**

The complete step-by-step procedure for installation using the RA4100 can be found in the *Compaq ProLiant Cluster for NetWare 6 with the RA4100 Storage Subsystem User Guide*.

For installation using the MA8000 or EMA12000, see the *Compaq ProLiant Cluster for NetWare 6 with the MA8000/EMA12000 Storage Subsystem User Guide*.

Upgrading from NetWare 5.1 to NetWare 6

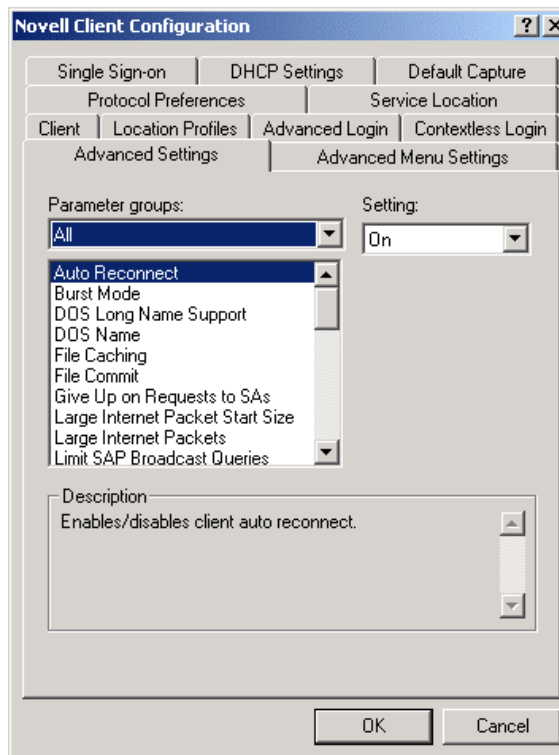
To upgrade the software:

1. Prepare the network and the computer.
2. Specify the hardware and software settings.
3. Create additional disk volumes, if required.
4. Select and install the networking protocols.
5. Set up Novell eDirectory.
6. Install other networking products.

The upgrade process automates the following tasks:

1. Device drivers and LAN drivers for the NetWare 6 operating system are loaded. Outdated drivers are matched with and replaced by new drivers included with NetWare 6.
2. eDirectory is upgraded.
3. NetWare 6 information is added to the autoexec.ncf and startup.ncf files.
4. The NetWare 6 files are copied to the server.

Client Configuration Parameters



When a node failure occurs, any NetWare clients logged in to the failed NetWare server automatically and transparently reconnect to the surviving node in the cluster responsible for taking over the data of the failed node.

Transparent client reconnect preserves users' drive mappings when cluster-enabled volumes are remounted on a surviving server. Transparent client reconnect also supports the failover of open files and file locks on Windows 95 and Windows 98 clients.

Additionally, server-based applications with transaction tracking can be configured so that when a node failure occurs, transactions will proceed uninterrupted by the failure and unnoticed by the user.

Network clients must use TCP/IP as a transport protocol. IPX application failovers cannot be seamless because of limitations of the uniform naming convention (UNC) path names.

Network clients must run NetWare Client32 version 3.1 or later to support transparent failover. The following Client32 parameters must be configured:

- For Windows 95 and Windows 98 clients:
 - Auto Reconnect Level = 3
 - Handle Net Errors = On
 - Name Cache Level = 0
 - Net Status Timeout = 60
 - NetWare Protocol = NDS
- For Windows NT and Windows 2000 clients:
 - Auto Reconnect = On
 - Preferred Network Protocol = IP
 - Protocol Component Settings = NDS

Learning Check

1. List three important features of NCS that help ensure and manage the availability of network resources.
.....
.....
.....
2. What are the objects that belong to the cluster container?
.....
.....
.....
.....
.....
3. Cluster resources are specifically defined as either applications or _____.
4. You are using either Load or Unload scripts and notice that Cluster Resource Manager is generating warning messages when a resource joins or leaves a cluster. What could be a possible cause for these warning messages?
.....
.....
.....
.....
5. A cluster resource is assigned to nodes 1, 2, and 3 in that order. Nodes 1 and 2 are offline. Node2 comes back online. Node 3 is not automatically failing back to Node2. Why not?
.....
.....
.....
6. What is the process for creating a cluster resource template?
.....
.....
.....

7. What are the three eDirectory objects created for a cluster-enabled volume?
.....
.....
.....
8. You have set the quorum membership field to 3. This is the number of nodes currently defined in the cluster. When bringing up the servers in the cluster, NCS reads this field, waits until the servers are up, and attempts to load resources. It cannot do so. Explain why.
.....
.....
.....
.....
9. What is the default time between transmits for all nodes in the cluster except the master?
 - a. 8
 - b. 6
 - c. 1
10. What property specifies the amount of time the master node waits for all other nodes in the cluster to signal that they are alive?
 - a. Master watchdog
 - b. Slave watchdog
 - c. Heartbeat
 - d. Tolerance
11. When changes are made to cluster objects in eDirectory, the Cluster Configuration Library updates the local version of the configuration on:
 - a. The master node, which sends the updates to slave nodes
 - b. The master and slave nodes at the same time
 - c. Only the master node, which does not forward the updates to the slave nodes
 - d. Its own local drive. It does not send updates anywhere.

12. The Cluster Management Agent (CMA.NLM) acts as a proxy for the ConsoleOne management application, allowing ConsoleOne to perform what functions?
-
-
-
-
-
-
13. In restoring an SBD partition, ConsoleOne has been used to delete the old SBD. All NCS NLMs have been unloaded using `uldncs.ncf`. What is the last step to be done before issuing the SBD install command?
-
-
14. List the six main steps to install NCS.
-
-
-
-
-
-
15. When upgrading to NetWare 6, the drivers are loaded, eDirectory is upgraded, and one other step is automated before the files are copied to the server. What is that step?
-

Cluster Monitoring and Management

Module 7

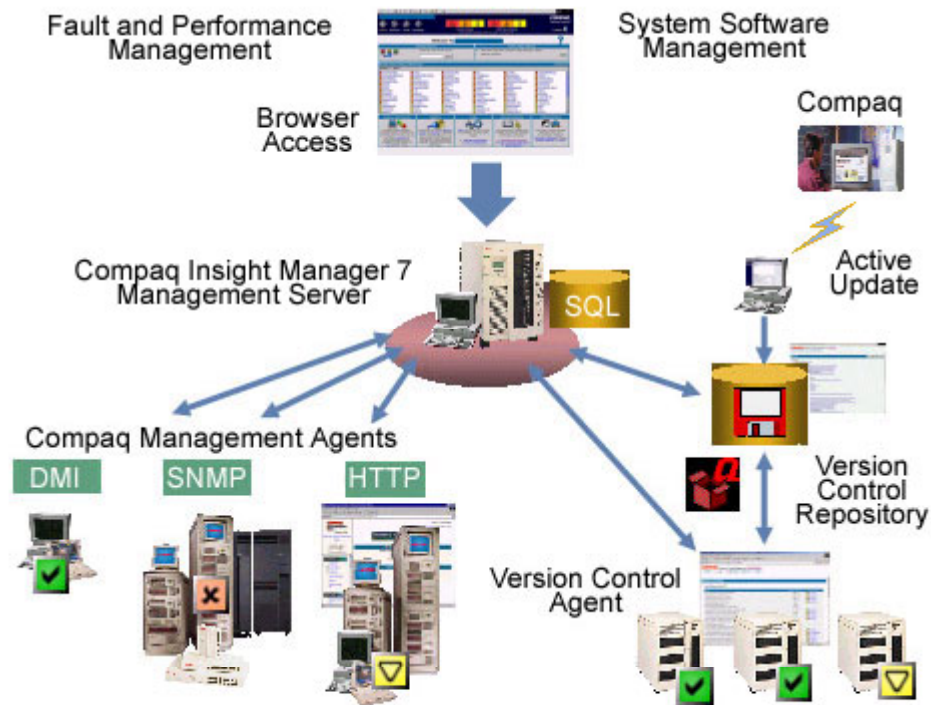
Objectives

After completing this module, you should be able to:

- Describe how to monitor Novell Cluster Services (NCS) using utilities from Compaq Computer Corporation and Novell.
- Describe how to manage NCS using utilities from Novell.
- Describe tuning techniques for improved cluster performance.

Monitoring Novell Cluster Solutions

Compaq Insight Manager 7



Compaq Insight Manager 7 Architecture

Compaq Insight Manager 7 is the core to Compaq Adaptive Infrastructure. It is an easy-to-use web-based enterprise management application that integrates current enterprise technology with the latest advances in web technology. It provides a proactive, automated, and cost-effective solution for managing and checking the health status of distributed systems.

Compaq Insight Manager 7 transforms the management of standards-based, distributed computing environments. By providing browser access to its components, it enables the management of devices and groups of devices anywhere, at any time.

Compaq Insight Manager 7 delivers:

- Device management.
- Event management.
- Intelligent monitoring.
- Alerting (both failure and prefailure).


- Remote maintenance.
- Visual control.
- Cluster monitoring.
- Software deployment on a 1:1 or 1:n basis.

It provides comprehensive system management for Compaq servers, workstations, and clients.

Compaq Insight Manager 7 monitors and manages system resources and provides fault, configuration, and performance management. The Compaq Insight Manager application runs on a Compaq server, acts as the management console, and interfaces with Compaq management agents.

This valuable program is used as a server administration tool to provide notification of hardware failures such as disk errors or network link errors. It can even be configured to provide notification through the network, a pager, or other means.

Cluster Monitor



Cluster Monitor provides enhanced management for clustered servers running on Compaq ProLiant and Compaq AlphaServers. The Cluster Monitor navigation pane displays all discovered clusters, detailed information regarding processor and disk usage, as well as environmental status on individual cluster nodes.

Settings Menu

The Cluster Monitor option in the Settings menu group provides the following options to set up clusters:

- **User Settings** — Selects the clusters that the Cluster Monitor displays to a particular user. This set of clusters (assigned) is called the *cluster management scope* or *scope*. By default, a user does not have any clusters in their scope until an administrator adds them.



Note

Only Compaq Insight Manager 7 administrators can configure a user scope.

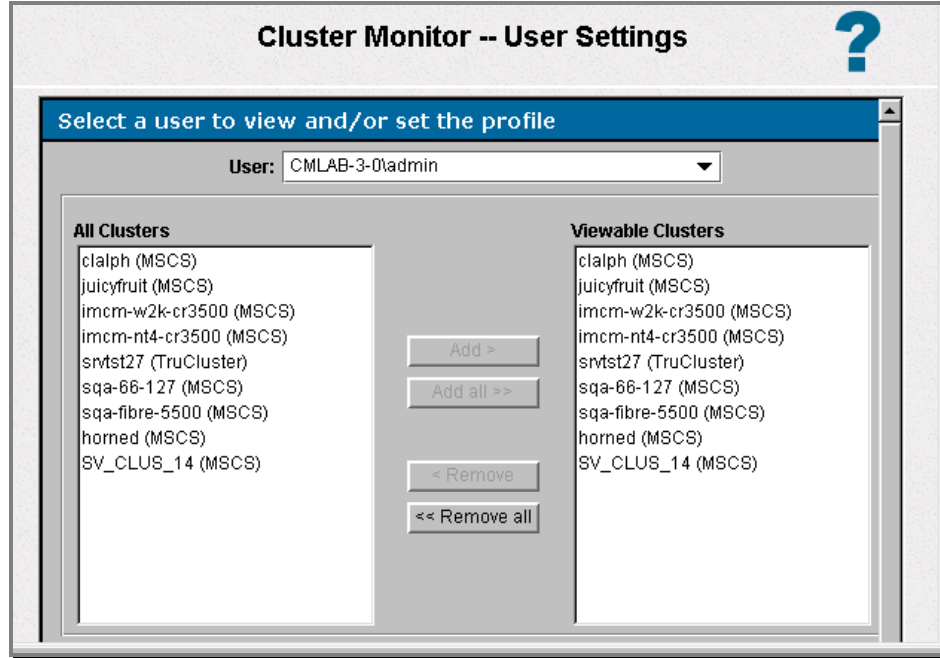
- **Cluster Settings** — Enters problem reporting and administrative information for the managed clusters.
- **Node Settings** — Enters problem reporting and administrative information for nodes.
- **Cluster CMX Settings** — Modifies operational parameters for cluster-level extensions called Cluster Monitor Extensions (CMXs).
- **Node CMX Settings** — Modifies the operational parameters for node-level CMXs.



Note

At the writing of this course, Cluster Monitor will support NCS 1.6 with the deployment of Support Pack 2 (SP2).

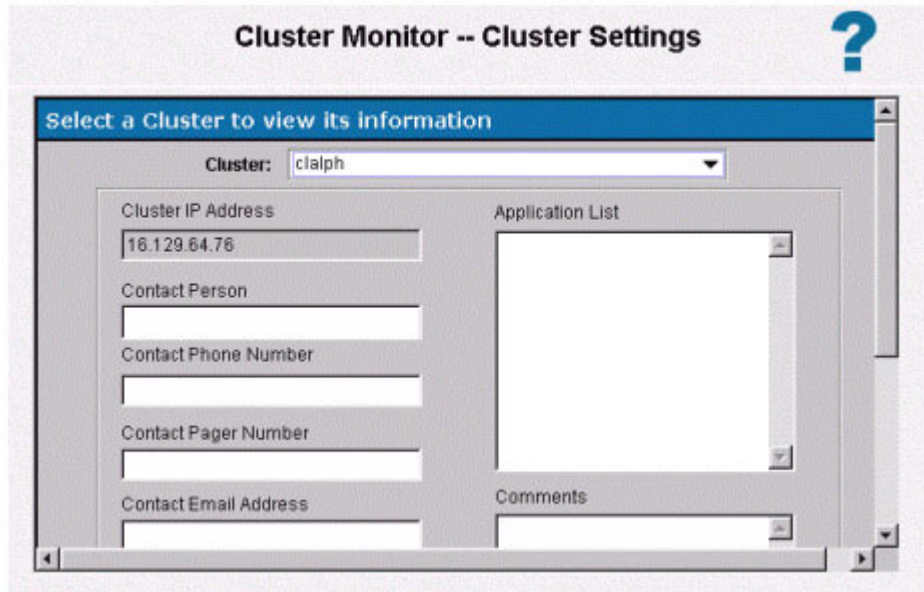
User Settings



The All Clusters and Viewable Clusters dialog boxes under Cluster Monitor — User Settings, determines the results under All Viewable Clusters by clicking *Devices* → *Queries* → *Cluster Monitor*. The display only brings up the clusters that have been specified here for that user's scope.

Cluster Monitor queries act as a filtering mechanism. The most clusters the user can see after running a Cluster Monitor query is the total number of clusters defined in their scope.

Cluster Settings



This page is used to enter problem reporting and administrative information for the clusters managed. This information displays:

- In the cluster monitor data area when selecting the cluster from the hierarchy tree.
- In alerts generated about the cluster.

Compaq Insight Manager 7 is designed to show the service name in parentheses beside the cluster name.

Compaq Intelligent Services Link

The Compaq Intelligent Services Link (CISL) is a service tool that integrates with Compaq Insight Manager 7 to provide advanced service-event filtering and reliable secure-event reporting to Compaq Customer Support Centers (CSC) over an Internet or remote access service (RAS) connection.

It also forwards event status reports back to Compaq Insight Manager 7 for online viewing by customer staff.

Management Using ConsoleOne



NCS can be managed by ConsoleOne, Novell's Java-based graphical user interface (GUI) application running on a client workstation on the network. ConsoleOne offers integrated and centralized management of NCS clusters, and supports only NCS clusters. It enables remote management of the cluster from any workstation with access to NDS.



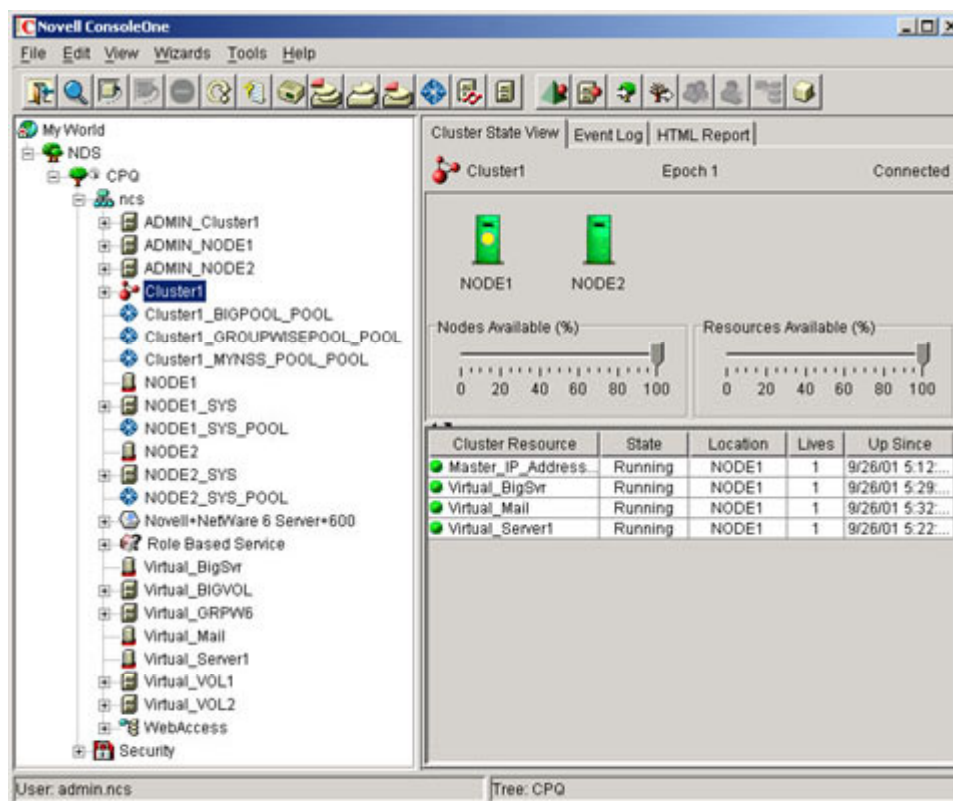
Important

The ConsoleOne workstation must reside on the same network as the cluster interconnect.



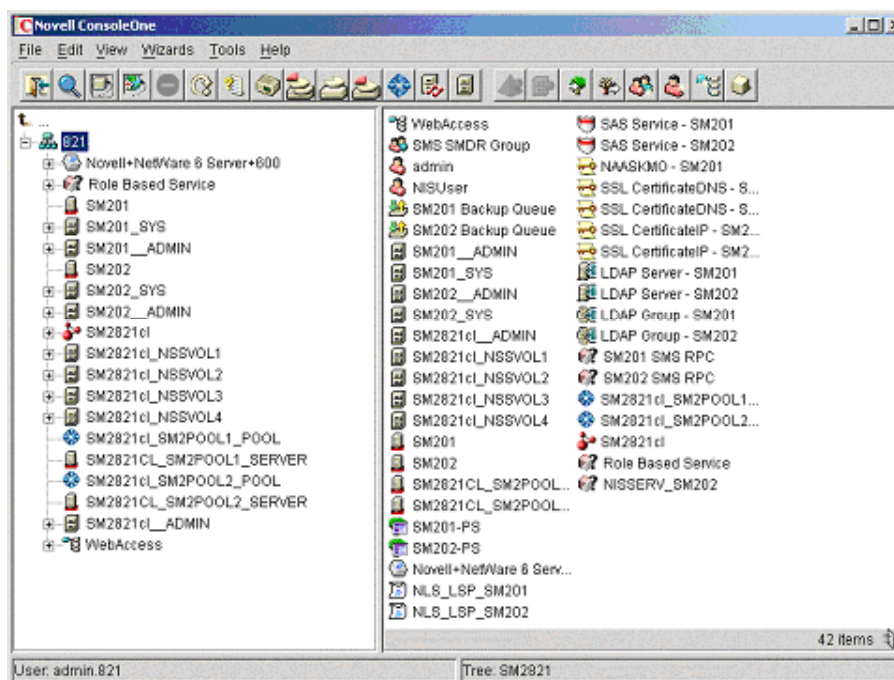
Note

Any task that can be done with ConsoleOne can also be done with NRM.



Any workstation running the ConsoleOne snap-in modules for NCS can manage the cluster. ConsoleOne enables management of the groups, resources, and operating state of the cluster. It is used to:

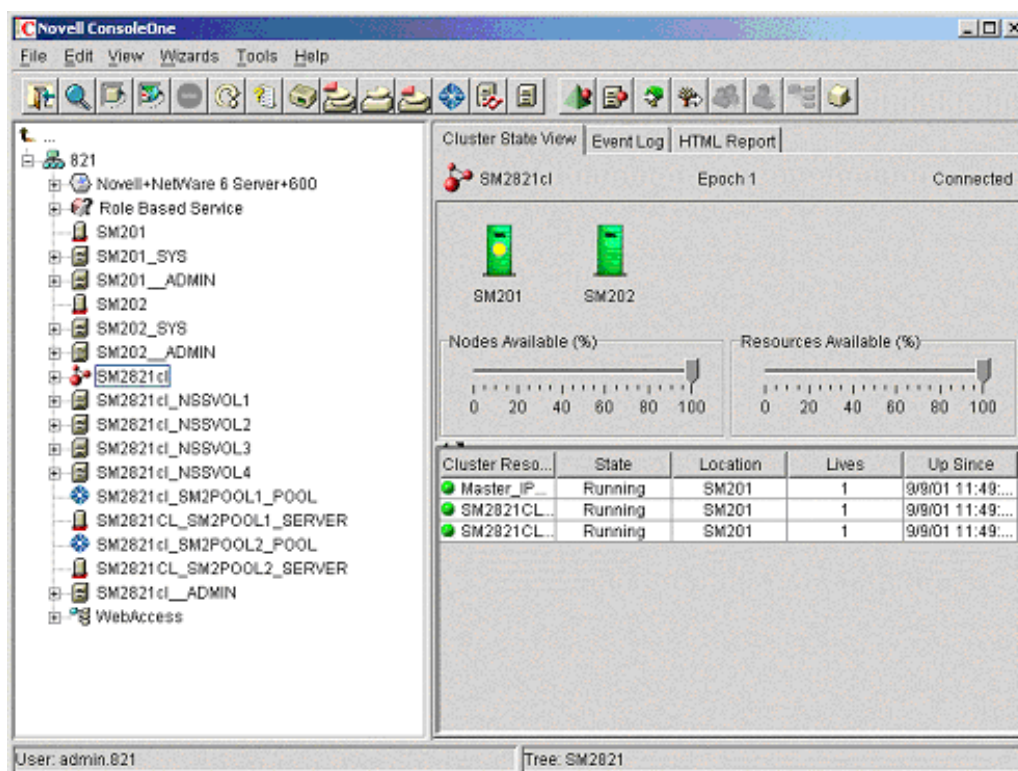
- View the current location of resources.
- Create or modify load scripts, unload scripts, and timeouts for applications.
- Add, move, or remove resources.
- Take resources online or offline.
- Specify failover and failback policies.
- Respond to or invoke manual failovers or failbacks.



Select the cluster object in the directory tree in the left-hand pane. Then, from the View menu, choose:

- Cluster State view
- Console view
- Partition and Replica view

Cluster State View



Cluster State View

The Cluster State view displays information about the status of servers and resources. It shows what volumes and services are running on each server in the cluster as well as the state (offline, running, and so on) of the volume or service.

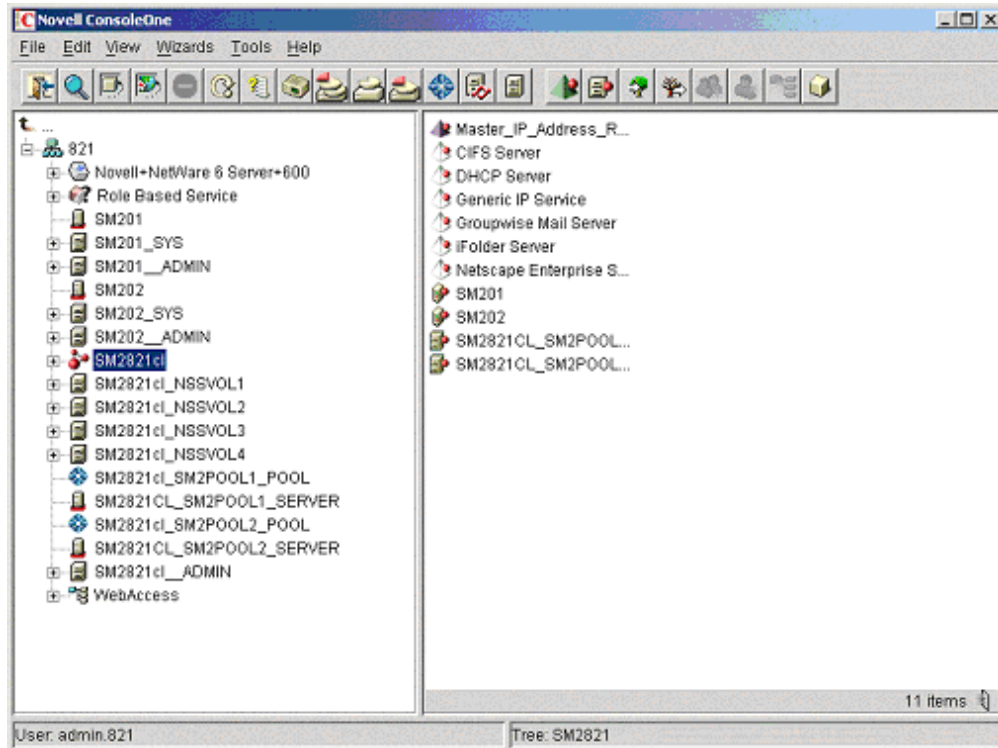
Identifying Cluster States

The Cluster State view in ConsoleOne provides information about the status of servers and resources in a cluster. Cluster servers and resources display in different colors, depending on their operating state.

- A green icon shows that the servers and resources are operating normally.
- A red icon signifies that a cluster has failed and needs administrator intervention.
- A gray icon indicates that the resource is:
 - Unassigned
 - Offline
 - Idle
 - Loading or unloading
- The yellow ball icon designates the master server in the cluster. The master server is the first server in the cluster, but another server can become the master if the first server fails.
- No icon indicates that a server is not part of the cluster or its state is unknown.

The Epoch number in the Cluster State view indicates the number of times the cluster state has changed. A change is recorded every time a server joins or leaves the cluster.

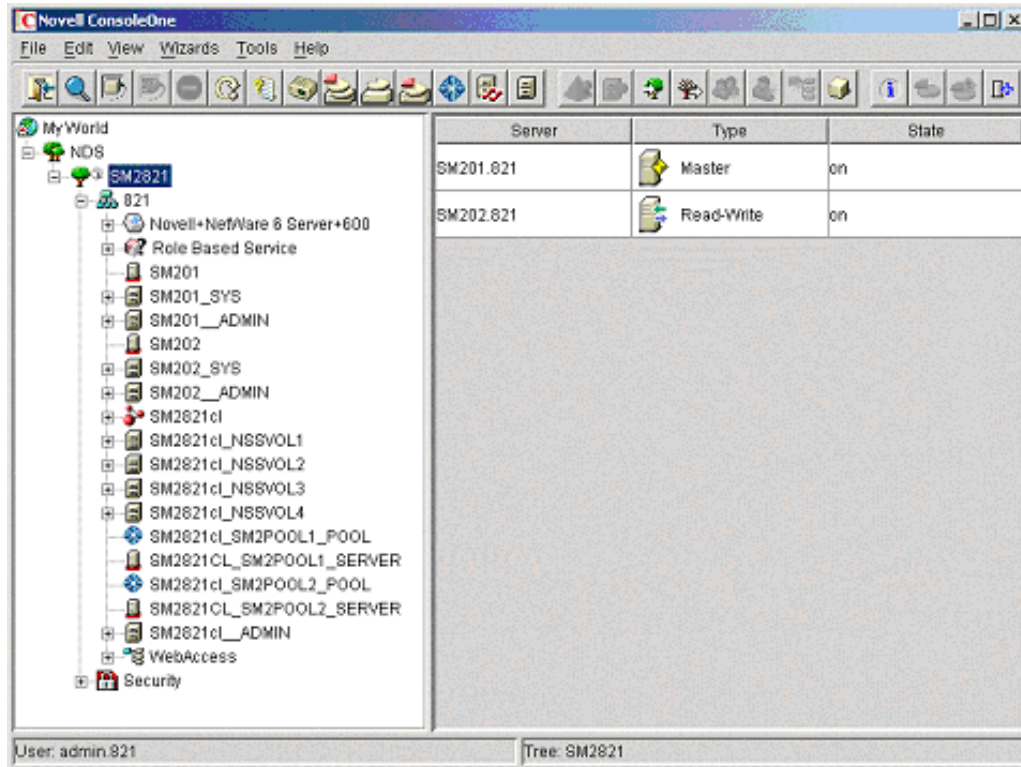
Console View



Console View

The Console view shows the cluster servers, resources, and templates. Change cluster and resource properties and create clustered volumes from this screen.

Partition and Replica View



Partition and Replica View

The Partition and Replica view shows the state of the replicas.

Migrating Resources

Resources can be moved to different servers in the cluster without waiting for a failure to occur.

Example

Move resources to:

- Decrease the load on a specific server.
- Free the use of a server so it can be brought down for scheduled maintenance.
- Increase the overall performance by moving a resource or application to a faster machine.

Migrating resources can balance the load and evenly distribute applications among the servers in the cluster. In ConsoleOne, browse and select the cluster object that contains the resource to be migrated.



Note

A resource will be unloaded from the server if its state is changed to *Offline*. It will not load on any other servers in the cluster and will remain unloaded until it is loaded again. This option is useful for editing resources because resources cannot be edited while loaded or running on a server.

Use ConsoleOne management capabilities to:

- Create shared disk partitions.
- Create NSS storage pools.
- Create NSS volumes.
- Cluster-enable NSS volumes.

NCS Command Prompt Commands

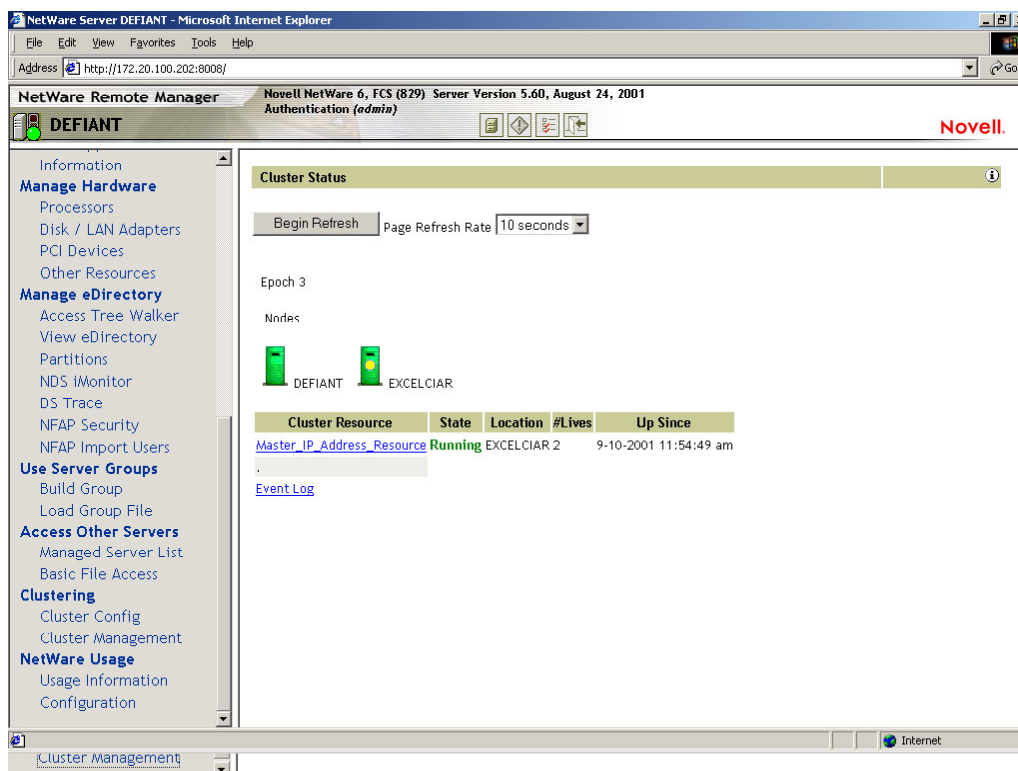
NCS creates two .ncf script files that can be run from the server console and are useful for updating cluster software or troubleshooting cluster problems.

- ULDNCS.NCF unloads NCS software from the server.
- LDNCS.NCF reloads NCS software on the server.

NCS provides several server console commands to help perform certain cluster-related tasks. Enter `Cluster Help` at the console prompt to view information on the following commands and their functions:

- `cluster join`
- `cluster leave`
- `cluster down`
- `cluster view`
- `cluster migrate`

NetWare Remote Manager



NetWare Remote Manager (NRM) is a Java-enabled web browser that provides a fully functional web-based alternative to ConsoleOne for configuring and managing NetWare 6 clusters.

Novell designed the cluster management interface in NRM similar to the ConsoleOne interface. For example, the Cluster Status screen in NRM is similar to the Cluster View screen in ConsoleOne.

NRM can be used to manage the cluster regardless of what nodes are running because NCS 1.6 has a special cluster IP address. This cluster IP address is created when the clustering software is installed. ConsoleOne or NRM can be used to edit this address later.

NRM provides management of an NCS 1.6 cluster at any time from any place with access to the Internet.

Tuning for Improved Cluster Performance

Managing a cluster requires procedures and applications different from those used to manage a single-server system. Some maintenance and monitoring procedures must be performed on a regular basis to ensure effective management of the entire cluster.

Throughout the life of the cluster, the following activities must be performed:

- Upgrade the hardware components.
- Upgrade the software.
- Increase the storage capacity.
- Restructure the cluster groups.
- Back up the cluster data.

Managing Network Clients

When the cluster is initially brought into a production environment, it is essential to educate users on the effects a cluster will have on the users' information systems needs. Users might experience a temporary disruption of service and performance degradation during failover or at other times when the system is unavailable.

When a failover or failback event occurs, the user will be unable to access applications and data. When users have been properly forewarned of the effects of operating in a clustered environment, they will more readily recognize when a failover or failback event is occurring or has occurred.

Load Balancing

Manual load balancing can improve the performance of a cluster by moving applications or resources from one node to another. Manually fail over as many cluster groups as necessary to balance the load of each node in a cluster.

When scheduling load balancing, consider what type of groups should be moved and how many clients are using the group. The success of the modification depends on how well the secondary nodes are equipped to handle the increase in workload.

Managing a Cluster Without Interrupting Cluster Services

Some maintenance activities require nodes in the cluster to be powered off, potentially causing serious effects to clients. Properly managing clusters requires minimizing the interruption of services to users.

File and print services typically are not business-critical and are good candidates to move whenever load balancing is required. NSS volumes will fail over quickly, reducing downtime during the load balancing operation.

Other applications, such as databases, should not be moved from one node to another during peak processing periods. When a database is moved from one node to another, it must be shut down and restarted. During the time it takes to restart, users do not have access to the database. Only move a database group to another node during nonpeak hours.

Managing a Cluster in a Degraded Condition

Because of the high-availability feature of clustering, applications and network clients remain operational even when one or more of the cluster components have failed. When the cluster is in this degraded condition, use the following to regain the system with minimal disruption to users:

1. Understand what caused the degradation.
2. Determine whether the condition will continually worsen.
3. Determine how critical it is to repair the problem.
 - If the problem is not critical, wait until a non-peak time to service the problem.
 - If the problem is critical, fail over all clustered applications and resources to the other node before servicing the problem.

Learning Check

1. List at least three features that Compaq Insight Manager delivers to cluster monitoring.
.....
.....
.....
.....
2. If using Cluster Monitor queries to act as a filtering mechanism, how many clusters can the user see after running a query?
 - a. As many clusters as there are on the network
 - b. As many clusters as there are in the tree
 - c. As many clusters as there are in the container
 - d. Only the clusters defined in their scope
3. In trying to use ConsoleOne to manage and monitor a NetWare cluster, the appropriate cluster located on another interconnect cannot be accessed. What might be the problem?
 - a. The right version of ConsoleOne might not be employed.
 - b. The ConsoleOne workstation must reside on the same network as the cluster interconnect.
 - c. JAVA.NLM might be unloaded.

4. ConsoleOne is being used to check the status of servers and resources. The cluster object in the left pane has been highlighted and the View Menu selected. There are now three choices. Which one is used to check the status of servers and resources?
 - a. Cluster State View
 - b. Console View
 - c. Partition and Replica View
5. ConsoleOne is being used to work with cluster resources. The cluster object in the left hand pane has been highlighted and the view menu selected. There are three choices. Which one is used to create clustered volumes?
 - a. Cluster State
 - b. Console View
 - c. Partition and Replica View
6. In the Cluster State view in ConsoleOne, what would the color of the icon be if a resource is unassigned?
 - a. Green
 - b. Yellow
 - c. Red
 - d. Grey
 - e. There would be no icon if the resource is unassigned.

7. NetWare Remote Manager is accessed from where?
 - a. A web browser
 - b. Server console
 - c. SYS:\System
 - d. SYS:\PUBLIC
8. List three components of NCS Management and Administration.
.....
.....
.....
9. List some maintenance and monitoring procedures that must be performed on a regular basis to ensure effective management of the entire cluster.
.....
.....
.....
.....
.....

Objectives

After completing this module, you should be able to:

- Describe how cluster-aware and noncluster-aware applications communicate with Novell Cluster Services (NCS).
- Explain what happens to directly connected devices in a clustered environment.

Cluster- and Noncluster-Aware Applications

Before running any application on servers in an NCS cluster, the application must be installed and configured properly to take advantage of shared storage and to achieve high availability.

The application documentation is the best source for specific application installation and setup instructions.

INTERNET

For more detailed information on deploying applications within an NCS environment, visit:

<http://www.novell.com/documentation>

Cluster-Aware

A key component of any cluster solution is the deployment of applications on the cluster. Applications are considered cluster-aware if they have been written to take advantage of the cluster software application programming interfaces (APIs). A cluster-aware application presents a single system image to network clients.

GroupWise is an example of an application that has been written for a clustered environment.

Benefits of cluster-aware applications include:

- **Increased efficiency** — When restarting an application and its resources on another node after a failure of the primary node
- **Active/active configurations** — Multiple copies of an application running simultaneously on multiple nodes and acting as standby nodes to the other nodes in the cluster
- **Automatic creation and configuration of the cluster resources and groups** — Configuration of resource dependencies and failover policies that are more difficult to perform manually
- **Self-monitoring** — The ability to monitor both itself and the cluster for fault situations and to initiate failover more quickly than a noncluster-aware application
- **Separate instances** — The ability to run separate instances of the database on multiple nodes in a symmetric configuration with shared logical volumes

Deploying DHCP with NCS

A NetWare Dynamic Host Configuration Protocol (DHCP) server automatically assigns IP addresses and other configuration information to clients on request or when the clients are restarted. If for some reason the NetWare 6 DHCP server is not accessible, clients lose their ability to connect to the network because they cannot obtain an IP address.

Cluster Benefits

Configuring DHCP with NCS ensures that the IP address range required by users to connect to the network is highly available. This is possible because DHCP configuration information and IP address ranges are stored in the eDirectory, and the DHCP server is automatically started, stopped, and restarted on different servers in the cluster by NCS.

DHCP Installation and Configuration with NCS

NCS 1.6 must be installed before configuring DHCP. NCS 1.6 provides a DHCP resource template that facilitates configuring DHCP in a cluster environment.

Much of the DHCP cluster resource configuration is performed automatically by the DHCP resource template, which will be edited with specific information for the targeted environment.

The DHCP resource template is not available with NetWare Cluster Services (NWCS) 1.0 and is not installed during the upgrade from NWCS 1.0 to NCS 1.6. The DHCP resource template is included only in a new installation of NCS 1.6.

Before running DHCP with NCS, DHCP must be configured properly on one server in the cluster.

DHCP server software is included and installed automatically when NetWare 6 is installed, so configure DHCP after the NetWare 6 server is installed. Remember to configure DHCP on only one server in the cluster.


Because all DHCP configuration information and IP address ranges are stored in NDS, no shared storage is required to run DHCP with NCS.

Any server in the cluster that will run DHCP must have at least a read/write NDS replica. The replica allows NCS to modify DHCP-specific NDS objects regardless of what server in the cluster holds the objects.

After DHCP is properly configured on one server in the cluster, create and configure a DHCP resource in NCS. This process includes configuring DHCP load and unload scripts, setting DHCP failover and failback modes, and assigning the DHCP resource to specific servers in the cluster.

The DHCP resource template will be edited with specific information because the targeted environment performs much of the DHCP cluster resource configuration automatically.

Deploying Novell GroupWise 6 with NCS



GroupWise 6 is an electronic messaging system that is tightly integrated with NetWare 6. Both the operating system and GroupWise can be administered from ConsoleOne.

Enhancements and features in the components of GroupWise 6 and GroupWise 6 SP1 include:

- GroupWise Administration
- GroupWise Agents
- GroupWise Clients
- GroupWise Internet Agent
- GroupWise WebAccess
- GroupWise Monitor

INTERNET

For a complete list of all the new GroupWise functions and features, visit:
<http://www.novell.com/documentation/lg/gw6/index.html>

Setting Up GroupWise in a Cluster

Clustering allows GroupWise post office and domain directories to be more highly available under fault conditions. Without manual intervention in the event of a fault condition, GroupWise services are able to start, stop, and restart.

Before configuring NCS to use GroupWise, GroupWise must be installed and configured properly.

Ensure that the following GroupWise configuration requirements are met:

- Configure GroupWise in client/server mode. GroupWise will not fail over or fail back in a cluster environment without the client/server mode.
- Put all GroupWise post offices and domains on shared volumes if GroupWise is to fail over or fail back.
 - The Message Transfer Agent (MTA) links to the post office through a TCP/IP connection to the Post Office Agent (POA).
 - The POA and the MTA use the same IP address.
- Do **not** update the autoexec.ncf file during the GroupWise Agent installation process because the loading and unloading of the mail services is controlled by the load and unload script properties of the cluster.
- Install the GroupWise Agent on all nodes of the cluster that can host GroupWise.
- Edit the POA and MTA startup files to reflect the correct path to the home directory, which is /home-volume:\directory\.

Installing GroupWise

To install GroupWise in a cluster:

- Install Novell GroupWise 6 on each node of the cluster.
- Install Novell GroupWise 6 SP1 on each node of the cluster.
- Activate and mount Novell Software Support (NSS) volumes.
- Configure cluster resources.
- Configure load and unload scripts.
- Create an active/passive configuration.

INTERNET

For more information on installing GroupWise, visit:

<http://www.novell.com/documentation/lg/gw6/index.html>

Using Protected Memory Space with GroupWise for Greater Availability

Using protected memory for GroupWise components can improve high availability by allowing the server to restart a memory space if any component within that memory space abends.

Restarting the memory space creates a new memory space with the same name, and restarts all modules in the new space with the same load order and parameters of the original memory space. This feature allows for recovery without failing the entire server, which enhances up time and database integrity.

The trade-off with protected memory is higher memory usage and a slight performance penalty. Ensure that the cluster nodes have sufficient memory to handle the number of memory spaces that might reside on them (remember to account for failover paths and the preferred nodes).

Another reason to use protected memory is the ability to unload a single instance of a module, rather than all instances. Use protected memory if there is any possibility of the same GroupWise agent loading multiple times on any node in the cluster. Without protected memory, if the unload script unloads GroupWise POA, it would unload all instances of GroupWise POA, so instead of migrating one Post Office, all Post Office Agents on that cluster node would shut down.

**Note**

Load and Unload scripts provide for the use of protected memory space.

Noncluster-Aware

Applications that have not been written to take advantage of NCS programming interfaces cannot detect cluster software and will behave in the same way as on a stand-alone server. Most noncluster-aware applications can be configured to fail over using load and unload scripts.

These applications often can be clustered in an active/standby configuration only and might need to be installed on all the cluster nodes.

These applications do not have the ability to automatically initiate a failover after sensing an application fault. The application will fail over only if the operating system or node hardware fails.

In addition, noncluster-aware applications do not synchronize application status information between multiple application instances running on different nodes in the cluster. When the application restarts on another node after a failover, it does not know the status of the failed application and changes to the configuration of an application on one node must be manually replicated to the standby application. The behavior of a noncluster-aware application after a failure is similar to how it would respond on a single server if the server was powered off and then powered back on.



Important

It is important to test and determine how long it will take the specific application to recover after a failure. Recovery time directly affects how soon network clients will be able to access the application after a failover.

Application Sizing with Compaq ActiveAnswers



Compaq ActiveAnswers can be used to size applications for single-server environments.

INTERNET ActiveAnswers can be accessed at:
<http://www.compaq.com/activeanswers>

Remember to account for processor load and additional memory when planning for performance after failover. Remember also to use combined user counts when sizing with an active/active configuration.

Application Pairing

Application Type	CPU	Memory	Disk	Network
Database	High	High	Very High	Low
File/Print	Low	Low	High	Medium
Web Server	Medium	High	Low	Medium
Mail	Medium	Very High	High	Low
App Server	High	Very High	Low	Medium

When multiple applications are run from the same server, avoid having combinations of two or more applications that put stress on the same server/storage subsystem. Consider both pre- and post-failover conditions when planning for combinations.

Avoid multiple active/active application configurations for servers.

Example

A combination of a web service and file/print service running on the same server is acceptable. However, combining a database service with other major application services that stresses the same subsystem does not lead to the best performance.

Compaq Application Partner Software

Compaq ProLiant Clusters provide a complete solution. Compaq supports clustered applications on ProLiant Cluster platforms through:

- Testing and validation of specific clustered applications on ProLiant Cluster platforms.
- Technical documentation.

ISV Partner Program

The Compaq High-Availability Independent Software Vendor (ISV) Partner Program links customers with Compaq approved ISV partners. The ISV Partner Program:

- Provides optimized, integrated solutions.
- Minimizes customer risk.
- Facilitates rapid application development.
- Minimizes customer time-to-success.
- Maximizes customer return on investment.
- Addresses the issue of solution life cycle.
- Provides customer service through a worldwide support infrastructure.

INTERNET

For the most recent list of high availability ISV partners, see:
www.compaq.com/enterprise/solutions/highavailability.html

Directly Connected Devices

Other factors to consider are the devices that are directly connected to each node. When a cluster resource or node failure occurs, only the devices configured in the cluster service can fail over. Devices that are directly connected to the node generally cannot fail over to other nodes, unless there are alternate means of accessing the devices.

Example

One example of a directly connected device is a printer. The node is providing print services to users and the printer is directly connected to the parallel port of the node. Although the print queue and spooler can be configured to switch over, there is no way to switch the physical connection to the other node.

In the case of clustered print services, the printer must support direct network connections rather than a direct connection to the node.

Example

Another example is a directly connected mainframe interface. If the first node is directly connected to the mainframe, such as through a synchronous data link control (SDLC) card in the node, there is no way to switch the physical connection.

In this case, use the client network to access the mainframe using TCP/IP. Because TCP/IP addresses can be configured to fail over, reestablish the connection after a failover using a media access control address (MAC). MAC addresses cannot be configured to fail over.

Other directly connected devices include:

- Modems
- Fax interfaces
- Customized input devices (bank card reader)

Directly connected devices must be examined closely on each node, and alternate solutions should be provided outside of what the cluster hardware and software can accomplish. These devices can be considered single points of failure because the cluster components might not be able to provide failover capabilities for them.

Learning Check

1. List three benefits of a cluster-aware application.
.....
.....
.....
2. Noncluster-aware applications cannot detect cluster software and behave in the same way as on a stand-alone server. Which of the following tasks will create an active/standby configuration?
 - a. Contact the vendor for a cluster plugin.
 - b. Go to download.novell.com for the latest application patch.
 - c. Write load and unload scripts.
 - d. Manually create a cluster resource in NDS.
3. If a printer is directly connected to the parallel port of a node in a cluster environment, what happens to printing services on that printer if the node fails?
 - a. The printer will be out of service.
 - b. The physical connection will be serviced by the other nodes.
 - c. Nothing, the printer will keep printing.
4. When configuring GroupWise post offices and domains to be on shared volumes for failover and failback, should the shared volumes for GroupWise be cluster-enabled to function with NCS?
.....
5. Explain why the `autoexec.ncf` file should **not** be modified to automatically start the `grpwise.ncf` file when launching GroupWise.
.....
.....

Cluster Maintenance and Troubleshooting

Module 9

Objectives

After completing this module, you should be able to:

- Describe how to add and troubleshoot shared storage in a Novell Cluster Services (NCS) cluster.
- Describe how to add and replace a cluster node.
- List and describe troubleshooting tools included with NetWare for use with NCS.
- Describe troubleshooting tools available from Compaq for use with NCS.

Adding and Troubleshooting Shared Storage in an NCS Cluster

Add shared storage to a Compaq ProLiant cluster as necessary to increase shared storage capacity.

The SYS volume must be installed on the local hard disk and not on a shared storage system. To avoid the possibility of mistakenly installing the sys volume on a shared storage system, power off all shared storage systems before performing this procedure.

Adding Shared Storage to the RA4100

To configure an RA4100:

1. Install and connect the hardware components.
2. Configure the logical arrays.
3. Create and cluster-enable Novell Storage Services (NSS) volumes.

Installing the Hardware

Before installing the storage, shut down and power off all cluster servers. Then install and connect any necessary hardware components:

- Drives
- Gigabit Interface Converters (GBICs)
- Interconnects
- Cables

After installing the hardware, verify that the firmware on the RA4100 controller in the new storage subsystem matches the firmware on the other RA4100 controllers.

Troubleshooting the RA4100 Shared Storage

Problems can arise from the use of the RA4100 storage subsystem as shared storage in a clustered environment. This section does not address problems that are specific to the storage subsystem itself or to the use of the storage subsystem in a stand-alone server configuration.

Verify that the array controllers and the cluster nodes are properly connected to all paths of the Fibre Channel Arbitrated Loop (FC-AL) or Fibre Channel Switched Fabric (FC-SW).

Using FC-AL or FC-SW Connectivity

1. Verify that a physical connection exists from each node to the storage hub or switch.
 - a. Verify that the GBICs are properly seated.
 - b. Verify that all Fibre Channel cables are properly connected to their GBICs.
2. Verify that all nodes can access the shared storage.

On each node, enter *List Devices* to verify that each node can see the same shared storage devices.
3. Confirm that NSS volumes and cluster volumes have been created and are online.
4. Verify that the `mount all` command is not present in the `autoexec.ncf` file.
5. Verify that the following command line is included in the `autoexec.ncf` file and that it is the only NSS command line in the file:

```
nss/ autodeactivate=all
```

This line must precede the `ldnfs.ncf` line in the file.

Adding Shared Storage to the MA8000 and EMA12000

Add shared storage to a ProLiant cluster using MA8000/EMA12000 storage as necessary to increase shared storage capacity. The tasks are basically the same as the steps performed to configure an RA4100 storage solution.

These steps include:

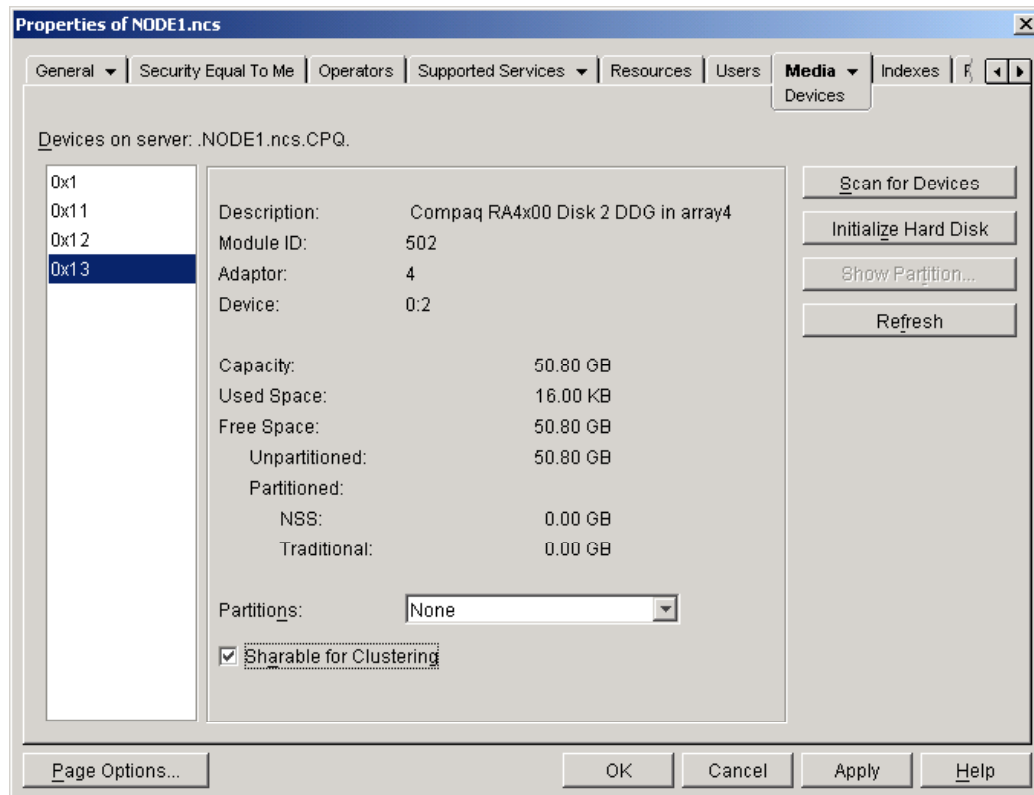
1. Add drives to the new storage subsystem.
2. Use the Run Config utility or StorageWorks Command Console (SWCC) to add disks to the configuration.
3. Create the necessary storagesets.
4. Initialize the storagesets.
5. Assign unit numbers to the storagesets.
6. Set the cache policy for each storageset.
7. Configure the HSG80 controller ports for the appropriate topology.
8. Change the connection operating system name to *NetWare*.
9. Launch ConsoleOne and ensure that any new logical unit numbers (LUNs) are accessible.

Verifying Configuration

After the configuration is complete, check to ensure that the cluster servers recognize the same drives and free space on the storage subsystem and that the disk partitions are valid.

Perform this check by using ConsoleOne or the NRM:

1. Start ConsoleOne.
2. Browse to a server object that is connected to the storage device.
3. Right-click the server *Object* and select *Properties*.
4. Select the *Media Devices* tab.
5. Confirm that the array configuration is on the resulting display.



Verifying Details

```

Novell RConsole: SERVER1
Server Screens: System Console (active)  << >> Sync Activate Disconnect Help
Thu Sep 23 09:23:44 1999
RCONAG6 123.45.67.254:1102 Remote console connection granted

SERVER1:nwconfig
Loading module NWCONFIG.NLM
NetWare Configuration Utility
Version 3.25   June 22, 1999
Copyright 1999 Novell, Inc. All rights reserved.
Auto-loading module NWI.NLM
NetWare Install (NWI) Module
Version 1.04   January 19, 1999
Copyright 1999 Novell, Inc. All rights reserved.

Retrieve Client/Server Database
(C) Copyright 1988-1993, 1996, Novell Inc.
All Rights Reserved.
SERVER1:list devices
0x0001: [U500-A0-D0:0] Compaq 53C876 Slot 0 Port 1 ID 0 COMPAQ MAB3091S
0x000A: [U025-A2-D1:0] COMPAQ CD-224E rev:8.0J
0x000C: [U502-A3-D0:0] Compaq RA4000 in Slot 1 of Orion I
0x000D: [U502-A3-D0:0] Compaq RA4000 Disk 0 MIR in Orion I
0x000E: [U502-A3-D0:1] Compaq RA4000 Disk 1 MIR in Orion I
0x000F: [U502-A3-D0:2] Compaq RA4000 Disk 2 MIR in Orion I
SERVER1:

```

View details about attached storage devices by entering the `List Devices` command at the system console prompt to display a screen similar to the one shown in the preceding graphic.

Troubleshooting the MA8000 and EMA12000 Shared Storage

Problems can arise from the use of the MA8000 or EMA12000 storage subsystem as shared storage in a clustered environment. This section does not address problems that are specific to the storage subsystem itself or to the use of the storage subsystem in a stand-alone server configuration. For these issues, refer to the documentation provided with the MA8000 or EMA12000 storage subsystem.

Verifying Connectivity to a Redundant FC-AL or FC-SW

Verify that the array controllers and the cluster nodes are properly connected to all paths of the FC-AL or FC-SW as described in this section.

Using FC-AL or FC-SW Connectivity

To ensure that all connections along the FC-AL or FC-SW paths are correct, verify the following:

1. Host bus adapters (HBAs) have been installed in the same slot locations across all cluster nodes.
2. For each HBA pair installed in the server, verify that the upper (or leftmost) HBA is connected by fiber optic cable to one storage hub or switch and that the bottom (or right) HBA is connected to the second storage hub or switch.
3. For the array controller pair installed in each storage subsystem, verify that the two ports on each array controller are connected to different storage hubs.
4. If a storage subsystem has three or fewer disk enclosures, verify that a dual-bus I/O module with two SCSI cables is installed in each disk enclosure. Also verify that 14 or fewer disk drive units are installed in each disk enclosure.
5. If a storage subsystem has four to six disk enclosures, verify that a single-bus I/O module with one SCSI cable is installed in each disk enclosure. Also verify that 12 or fewer disk drive units are installed in each disk enclosure.

Connecting the MA8000 or EMA12000 to FC-AL or FC-SW Connection Paths

To ensure that an array controller is connected to the correct I/O path, open a Command Line Interface (CLI) window and issue the command `show this_controller`. If the port 1 topology is in an offline state, refer to the StorageWorks documentation about the HSG80 array controller software, and to the CLI help documentation for more information about how to configure an MA8000 and EMA12000 storage subsystem.

The topology for port 1 on each array controller should be:

- Loop Up state (not the Standby state)
- LOOP_HARD

Connecting Server Nodes to FC-AL or FC-SW Connection Paths

To determine servers are connected to all FC-AL or FC-SW paths, open a CLI window and issue the command `show connection`. The output of this command shows status information about present and previous connections to all FC-AL or FC-SW paths including:

- Connection name
- Controller
- Controller port
- Adapter ID address
- Status
- Unit offset

The status entry indicates that the port is either online (connected) or offline (disconnected).

Rename the connection name field to reflect the configuration of the array controllers.

Example

Enter `VR1TOP1` for node 1, top controller, port 1.

This makes it easier to identify the FC-AL or FC-SW connections.

Determining Why a Node Cannot Connect to the Shared Drives

If one of the cluster nodes cannot communicate with the shared drives in the MA8000 or EMA12000 storage subsystem, follow these steps:

1. Verify that a physical connection exists from each node to the storage hub or switch.
2. Verify that the GBICs are properly seated.
3. Verify that all fiber optic cables are properly connected to their GBICs.
4. Verify that all nodes can access the shared storage. On each node, use ConsoleOne to verify that the same shared disk resources are seen from every node in the cluster.
5. Confirm that NSS volumes and cluster volumes have been created and are online.
6. Verify that the `mount all` command is not present in the `autoexec.ncf` file.
7. Verify that the following command line is included in the `autoexec.ncf` file and that it is the only NSS command line in the file:

```
nss/ autodeactivate=all
```

This line must precede the `ldnccs.ncf` line in the file.

Using Common Fibre Channel Shared Storage Troubleshooting Techniques

The following techniques are used to troubleshoot common problems on the RA4100, MA8000, and EMA12000 storage systems.

Verifying the Hardware

Verify that the following components are powered on:

- Cluster nodes
- Storage subsystems
- Storage hub or switch
- Fiber Channel hub or switch for the cluster interconnect and public LAN

The correct power-on sequence is:

1. Storage hubs or switches (Power is applied when the AC power cord is plugged in.)
2. Storage subsystems
3. Ethernet hubs or switches
4. ProLiant servers



Note

The Fibre Channel interconnect devices and some network hub devices do not have on/off buttons; they receive power when the power cord is plugged in.

If intermittent loop errors are occurring, upgrade to the latest firmware.

If the operating system does not recognize storage devices:

1. Verify that the proper driver is loaded.
2. Verify the proper start order of devices.

Because a cluster uses many hardware components, physical connections between the components are used for communication. If redundant connections exist, ensure that the redundant paths are cabled correctly. Verify that all cable connections are properly made between:

- Cluster nodes and the cluster interconnect hub or switch.
- Cluster nodes and the storage hub or switch.
- Cluster nodes and the client LAN hub or switch.
- Storage subsystems and the storage hub or switch.

Determining Why Devices on One I/O Connection Path Cannot Be Seen by the Cluster Nodes

If the cluster node cannot see all the storagesets that are configured on the I/O paths, follow these steps:

1. Verify that the array controllers on the unseen path are set for multibus failover mode. If an array controller has not been set to multibus failover mode, the cluster nodes will only see the storagesets that are configured to the other array controller in the storage subsystem.
2. If the array controllers have already been set to multibus failover mode, verify the integrity of the components along the unseen I/O path.
 - a. Check the status LEDs on each HBA for error indications.
 - b. Verify that the Fibre Channel cable has been properly installed between each HBA and its storage switch.
 - c. Check the status LEDs on the storage switch ports for error indications.
 - d. Verify that the Fibre Channel cables have been properly installed from the storage switch to both ports on each array controller.
 - e. Use the CLI to obtain status information about the array controllers. Verify that both port 1 and port 2 on each array controller are active.

Determining Why RAID Volumes in the Storage Subsystem Are Not Recognized by All Servers

If shared disk resources are not recognized on a node, follow these steps:

1. Restart the affected nodes.
2. Ensure that the Fibre Channel HBA driver is installed and running on all cluster nodes.

INTERNET For installation information, refer to the *Compaq StorageWorks Fibre Channel Host Adapter Installation Guide* at: www.compaq.com

3. If the volumes are still not recognized, start investigating the problem from the lowest level of drive configuration.
 - a. Run the cpqonlin utility to discover if any of the drives are offline (degraded or failed).
 - b. The cpqonlin utility is installed from the Compaq Software Support Diskette for Novell NetWare (contained on the Compaq SmartStart and Support Software CD) or from a diskette pack. If this utility is run offline, shut down the servers and restart with either the SmartStart CD in the CD slot or diskette number 1 in the A: drive.
4. If all drives are not recognized by the cpqonlin utility, a physical connection problem probably exists.
 - a. Verify that the lights on the HBA, storage hub or switch, and array controller are green.
 - b. Verify that the GBICs are properly seated in the HBAs, the array controller, and the storage hub or switch.
 - c. Verify that all fiber optic cables are properly connected to their GBICs.

INTERNET For details on how to connect the GBICs and the Fibre Channel cables, refer to the *Compaq StorageWorks RAID Array 4000 User Guide* or the *Compaq StorageWorks RAID Array 4100 User Guide* at: www.compaq.com

- d. If using cascaded FC-AL switches, verify that the same port number on the master and slave switches is used for the interswitch link. Then use the Switch Management Utility to verify that the correct ports are selected for cascading.

5. If all drives are recognized by the `cpqonlin` utility and are configured correctly, use ConsoleOne and verify that all drive volumes are present.

INTERNET

If problems continue, refer to the *Compaq Fibre Channel Troubleshooting Guide* for additional corrective actions for the RA4100 at:
www.compaq.com

Determining Why A Node Does Not Recognize a Volume

If a cluster node does not recognize a volume, follow these steps:

1. Try to mount the volume on another server. If that fails, a physical problem with the volume might exist.
2. Run the `List Devices` command to ensure all the RAIDsets appear on the volume.
3. If one or more RAIDsets do not display in the List Devices output, run *Scan for New Devices*.
4. If the node still does not recognize the volume, the volume might not be properly configured.
5. If the volume has been properly configured, check the SSP setting for the affected node. If SSP is enabled on a node for a particular logical drive, only nodes with SSP enabled for that drive would recognize the drive. Any node with SSP disabled for that drive will not recognize the drive. Disable SSP on all nodes for all logical drives.

Determining Why a Mirrored Cluster Services Partition Cannot Be Synchronized

If a mirrored cluster services partition cannot be synchronized or is corrupted (because of an array failure, for example), use the following procedure to create a new cluster services partition.

1. Use ConsoleOne to delete both the original Cluster Services partition and the mirror.
2. Shut down the cluster servers.
3. Enter the following command to run the System Bus Driver utility to create a new cluster services partition and mirror:

```
sbd.nlm install
```
4. Follow the on-screen prompts.
5. Restart the cluster servers.

Verifying Interconnects

If using the recommended method of using static IP addresses, ensure that the hosts file is properly set up.

A backup cluster interconnect is provided by the Fibre Channel link to a shared disk drive used for quorum arbitration, a scheme whereby each cluster server can check on the availability of the other servers.

In a configured dedicated cluster interconnect:

1. Ensure that the cluster interconnect cable is connected to the cluster interconnect adapters for a two-node cluster, **not** the client LAN adapters.
2. Ensure that the cluster management client running ConsoleOne is on the cluster interconnect network, not the public network. This configuration is required because NCS determines the operational status of the cluster nodes by monitoring information sent over the cluster interconnects. If the management workstation is connected to the public network and not the cluster interconnects, cluster management is not possible.

Adding or Replacing a Cluster Node

Servers can be added to an existing cluster, or a server node can be added back to a cluster to which it had previously been a member.

Installing NetWare on a Server to Be Added to an Existing Cluster

Always run the NCS installation program when adding new nodes to an existing cluster.

Run NWDEPLOY.EXE from the root of the CD to launch the NetWare Deployment Manager. Accept the license agreement and then complete the installation steps.

The NetWare 6 installation program automatically copies NCS files to every NetWare 6 server. Even though NCS files might already exist on each NetWare 6 server, run the NCS installation program to configure and set up the cluster nodes.

Leaving the Skip File Copy check box unchecked will cause existing NCS files to be copied over, but will not otherwise affect the installation.



Note

If the existing cluster is using the two-node license that ships with NetWare 6, the cluster must be upgraded with more licenses before additional nodes can be added.

When adding a server to a cluster, NCS automatically detects the IP address of the server. If the server being added has more than one IP address, a prompt to select the IP address for NCS to use displays.



Important

Have at least 15 to 25MB of free space that is not part of an NSS partition on one of the shared disk drives to create the cluster partition. If no free space is available, the shared disk drives cannot be used by NCS.

Adding a Node to a Cluster to Which it Had Been a Member

Use the following steps to add a server back to a cluster to which it had previously been a member, but had been removed for some reason.

1. Remove the cable from the Fibre Channel card on the server.
2. Install NetWare on the node using the same node name and IP address.
3. If the cluster node object for the server is still present, use ConsoleOne to delete the object. Do this by going to the cluster container, selecting the node in the right frame, and pressing *Delete*.
4. Reconnect the server to the shared storage by reconnecting the cable to the Fiber Channel card.
5. Run NCS install from NWDEPLOY. The node will take on its former identity.
6. Review and update the preferred node lists in resources that use this server.

NetWare Troubleshooting Tools and Tips

The following software is required on each cluster node:

- NetWare 6
- NCS 1.6
- Server cluster licenses and cluster user access licenses

For the Operating System

After the physical connections are validated, ensure that communication is occurring between all cluster components.

Ensure that all servers in a cluster are in the same NDS tree and that all servers are on the same TCP/IP subnet.

Use the TCP/IP ping utility for the Ethernet cluster interconnect and client LAN connections. If the cluster interconnect and the client LAN are bound to the same network interface card (NIC) as recommended, ping the name from each cluster node and then the address of the network connections located in the other nodes.

In a configured, dedicated cluster interconnect, from each cluster node ping the name, then the address of the cluster interconnects located in the other nodes.

Similarly, from each cluster node, ping the name, then the address of the client LAN connections located in the other nodes.

For NCS

Before installing NCS:

1. Ensure that the client being installed from has a connection to each server that will be added to the cluster.
2. Ensure that all mapped drives on the client are operational and are connected to the client (a faulty connection is indicated by a red “crossed out” symbol on the display). A faulty drive connection could cause the client to lock-up during NCS installation.
3. Ensure that all servers are on the same subnet mask. A server on a different subnet mask cannot be added to the cluster.
4. Leave the RAID volumesets on shared storage as free space. Do not create NetWare or NSS partitions. If no free space is present, an NCS partition cannot be created.

If NCS does not load:

1. Check the Membership property. If set to more server objects than are currently available to the cluster, NCS will not load.
2. Check the Timeout parameter. If set too low, the cluster might not start because fewer nodes than are defined in the Membership field are fully up and operational at the expiration of the timeout period.

If the timeout period expires, NCS should load and advertise resources using whatever nodes are available.

For Licenses

Install a Cluster Server License for every server in the cluster and a User Access License for the number of clients connected to the cluster. NCS will not function without the required licenses installed.

Tools

Troubleshooting tools are necessary for locating problems and repairing performance degraded clusters and communication problems between nodes. NetWare provides tools useful for troubleshooting clusters. Because NCS is an integrated part of the NetWare operating system, the same tools used to troubleshoot the operating system can be used to troubleshoot NCS.

Some troubleshooting tools provided by Novell are:

- **iMonitor** — Is now used for DSRepair (synchronizes master and slave NDS replicas) and DSTrace.



Note

Do **not** use *Rebuild* without first running a *Verify* and running the scan tool against the data.

- **NetWare Logger screen**
- **Conlog.nlm** — Captures console messages.
- **Config.nlm** — Captures server configuration information.

Tips

Additional troubleshooting tips for NetWare clusters include:

- Use the system error log file, sys:system\sys\$log.err.
- Comment out Compaq survey.nlm in the autoexec.ncf.



Note

The abend.log contains abend messages and resides in the sys:system directory.

Compaq Troubleshooting Tools and Techniques

Compaq has developed tools and techniques to assist in troubleshooting ProLiant clusters. These tools and techniques are for:

- QuickFind 2000
- Fibre Channel solutions
 - Compaq Fibre Channel Fault Isolation utility (FFI)
 - SANworks Secure Path for NetWare
- Device drivers
- NLMs

Using Compaq QuickFind 2000

Compaq QuickFind 2000 is a self-help information retrieval tool providing access to the most comprehensive and current CD-ROM-based collection of information on Compaq products and services. The Compaq QuickFind four-CD set contains:

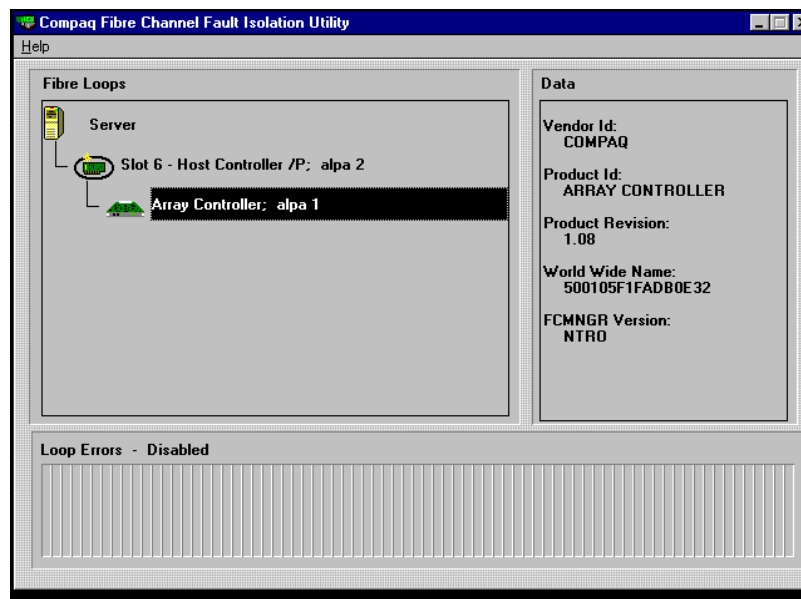
- **TIM Viewer** — TIM Viewer enables navigation through either a hierarchical system or using a search tool. QuickFind 2000 information collections can reside locally or remotely and the TIM Viewer provides viewing, navigating, and searching of these collections.
- **Information collection CD-ROMs** — These CDs provide an integrated set of information consisting of customer advisories, product bulletins, maintenance guides, parts lists, and other information for Compaq hardware products.

The Compaq QuickFind 2000 CD-ROM subscription service provided as a benefit with several Compaq service provider programs, such as:

- Compaq Authorized Service Providers
- Compaq System Service Providers
- Compaq Authorized Service Engineers
- Compaq Self-Maintainers

Using Fibre Channel

Using the Compaq Fibre Channel Fault Isolation Utility



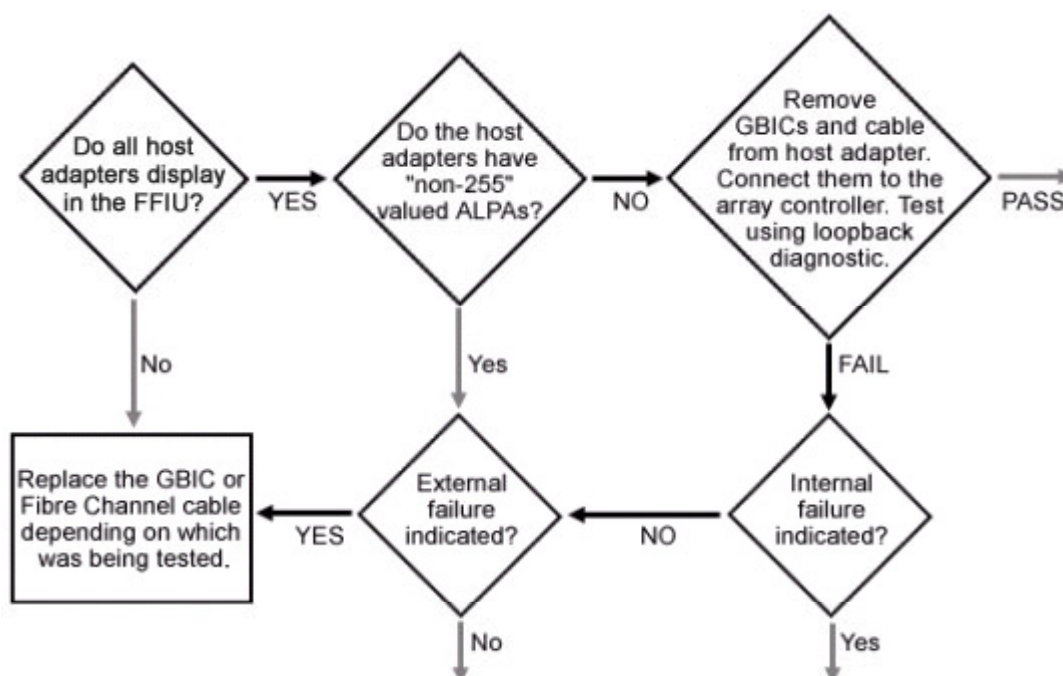
The FFI utility, when run on one of the cluster nodes, verifies the integrity of a new or existing FC-AL installation. The utility provides fault detection and helps in locating a failing device on the FC-AL. FFI is an offline utility that can be run from the SmartStart and Support Software CD.



Note

The FFI utility does not support the FC-SW topology, modular data routers, or MA8000/EMA12000 storage devices.

FFI displays all devices that are attached to the FC-AL and tests for link errors within the loop. A link is that portion of the loop between the Fibre Channel HBA and a storage hub, or between a storage hub and the RAID Array.



The FFI utility analyzes Fibre Channel components including HBAs, array controllers, and tape controllers. Fibre Channel storage hubs are logically transparent to the operations of the FC-AL, but even a failed hub can be detected when the FFI utility is used in combination with troubleshooting flow charts like the sample shown in the graphic.



Note

Information on how to use this utility and the troubleshooting flow chart can be found in the *Fibre Channel Troubleshooting Guide*, PN 297877-006, that ships with the storage hubs. It is also available at:

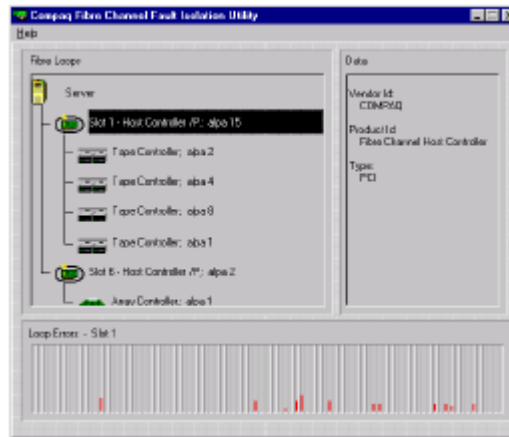
<ftp://ftp.compaq.com/pub/products/storageworks/ebs/297877e2.pdf>

Use the FFI utility to:

- Verify controller firmware revisions.
- Determine whether a loop segment is faulty.
- Isolate a faulty loop segment.

Each device on the FC-AL has an Arbitrated Loop Physical Address (ALPA). The ALPA is allocated dynamically and can change with each power-up or as new devices are added to the loop.

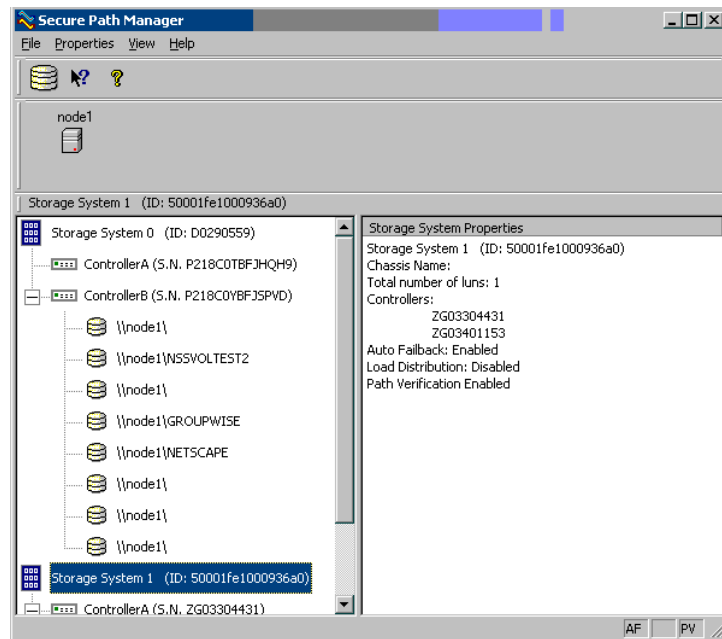
The FFI utility indicates all devices that are active on the FC-AL. It dynamically updates the screen to show the current ALPA of each device. If a device has a default ALPA of 255, it is not initialized.



Loop Error Histogram Showing a Single, Three Second–Error Interval

A Loop Error Histogram, measuring the frequency with which errors occur on the logical FC-AL, displays when specific servers or components are highlighted to further analyze the health of individual components.

Using SANworks Secure Path for NetWare



Secure Path network connectivity troubleshooting information includes:

- Client/agent considerations
- Network considerations



Note

Two versions of Compaq SANworks Secure Path for NetWare are now available:

- **SANworks Secure Path for NetWare for the RA4100** — Comes with the HA/N200 cluster kit
- **SANworks Secure Path for NetWare** — Comes with the HA/N500 cluster kit

Troubleshooting Clients and Agents

The following client and agent considerations can be useful in troubleshooting network connection problems:

- Use the Agent Configuration utility to add the system and client name of each client to the list of authorized clients, then set the password in the Password Dialog box. When the modifications are made, stop and restart the Secure Path Agent to update the database using the Services applet from the Control Panel.
From the server console enter:
 - a. `unload cpqspagt`
 - b. `load cpqspagt`
- Each name used must be mapped to its network IP address using either the HOSTs file (server and client name mapped to IP) or the Domain Name System (DNS) with a fully qualified domain name.
- In cluster configurations, ensure that the password chosen is common for both agents in the cluster.
- Secure Path does not use Novell eDirectory to authenticate and authorize clients. Client authentication and authorization is handled for each Agent using name-to-IP address resolution and password verification from the Secure Path configuration database.

Troubleshooting Network Connections

The following network considerations can be useful in troubleshooting network connection problems:

- Client names up to 15 letters without a dot or period can be resolved by NetBIOS broadcast resolution, if the client and agent nodes are located on the same subnet. If the client and agent are located on different subnets, use the LMHOSTs file, HOSTs file, WINS, or DNS to resolve the address.
- If using the LMHOSTs file, ensure that the Enable LMHOSTs Lookup box is checked in the TCP/IP protocol properties of the client system.

On the client system, enter the NETBIOS name and the IP address of the Agent to connect with in the LMHOSTs file and save it.

Click the *Import LMHOSTs* button to specify the location of the LMHOSTs file. The LMHOSTs and HOSTs files typically are located in the `\system32\drivers\etc` subdirectory.

- From a command prompt, issue the `nbtstat -r` command to purge and reload the remote name table.
- Client names that exceed 15 letters or carry a dot require an entry for that name in the HOSTs file or resolution by a DNS server. It is also possible to have DNS resolve NetBIOS names as long as the DNS server is updated with the appropriate information. Check the *Enable DNS for Novell NetWare Resolution* box in the TCP/IP protocol properties of the client system.
- If using DNS for host name-to-IP resolution, the DNS database on the DNS server must be updated with the appropriate information.
- For the best network connection results, use fully qualified domain names with DNS.
- For production environments where management and security are a concern, use fully qualified names with DNS name resolution.
- For test and evaluation environments, it is usually easier to add the name of the server to the HOSTs file of the client and the name of the client to the HOSTs file of the server.
- Verify that the Secure Path host can be pinged locally and from a remote host using the host name, not the IP address.

Using Device Drivers

Although the Compaq drivers and utilities are installed during the installation, new and updated drivers and utilities from the Compaq Support Paq (CSP) for Novell NetWare must also be installed. The CSP is contained on the Compaq SmartStart and Support Software CD.

Drivers and utilities must be reinstalled after a Support Paq installation ensure that the correct versions were not overwritten when the Support Paq was installed.



Note

Installing Compaq Support Paq for Novell NetWare requires administrative access to the server. Before installing any products:

- Back up the server.
 - Apply the latest recommended minimum support patches or Enhancement Packs from Novell.
 - Perform the installation when system utilization is low.
-

Ensure that the same driver versions are used on each cluster node and that the same firmware versions are used on each storage subsystem in the cluster.

Minimum driver and firmware requirements for all ProLiant Cluster Solutions for NetWare are:

- Cpqfc.ham driver (2.52 or greater) for the 64-bit, 66MHz PCI-to-Fibre Channel adapter — The SLOT parameter is not required unless there is a preference for initialization order. It provides no SMP support and the threads run on PO only. This process is called funneling.
- Cpqraid.ham 2.01 driver for the Smart Array 53xx controller, Advanced Data Guarding (ADG), and Compaq hot-plug tape drives.
- Firmware — Multiple SAN switches must have matching array controller firmware.

Depending on the configuration, all firmware might have to be of the same revision level. If a hardware device is part of a path that is accessible by more than one cluster node, more than one shared storage subsystem, or more than one client computer, the firmware level of the hardware devices must be identical. The interconnect firmware for the switches must also be the same.

INTERNET

The latest firmware level is specified in the *NetWare Cluster Certification Matrix* available at: <http://www.compaq.com/highavailability>

Use the NetWare Configuration Utility to install the latest drivers and utilities from the SmartStart CD. The directory path on the SmartStart CD for the Compaq Software Support Diskettes for Novell NetWare is `cpqsmstxxx:\cpqsupsw\nwesp`. Choose to either *Install Selected Files* or *Install All Files* during the installation, depending on the needs.

Install the following files from the Compaq NWPA Storage menu:

- `cpqfc.ham`
- `cpqfc.ddi`
- `cpqshd.cdm`
- `cpqshd.ddi`

**Note**

If another instance of the Compaq Software Support Utility for NetWare is running, exit that instance before performing this or the utility will not load properly.

Using NLMs

Install the Compaq Online Configuration Utility (cpqonlin.nlm) to configure SMART-2 Array controllers and network controller fault tolerance.

Optional utilities include:

- **cpqiml.nlm** — Views and manages server events stored in the Integrated Management Log (IML)
- **cpqpower.nlm** — Displays redundant power subsystem status
- **cpqhthl.nlm** — Enables the server management features of the ProLiant servers
- **cpqsbd.nlm** — Controls the system hardware required to perform hot-plug functions
- **cpqnssu.nlm** — Updates Novell SSD drivers on a NetWare server automatically



Important

Install the driver versions that are designed specifically for use with NCS. For example, check the version number of cpqfc.ham by entering the command `modules cpqfc.ham`.

For current driver versions, see the *Compaq ProLiant Cluster for NetWare 5 Certification Matrix* available at: <http://www.compaq.com/highavailability>

When the installation of the utility is complete, reset the server (enter `reset server` at the command prompt). Repeat this procedure for each cluster server that might be used to manage the shared storage array.

Comparing Software Components

Shared SCSI Solutions

Compaq	Novell
Compaq Insight Manager 7	NetWare 6
CPQSSCSI.ham for CR3500	NSS (supports manual setting of shareable cluster flag)
CR3500 Configuration Utility	NetWare Client (for access to NetWare volumes through NCP)
Offline ACU	NCS 1.6 (up to 12 nodes supported; greater than 12 nodes supported by Professional Services)
CPQONLIN	Cluster Migration Utility
	NRM

Fibre Channel Solutions

Compaq	Novell
Compaq Insight Manager 7	NetWare 6
CPQFC.ham	NSS (supports manual setting of shareable cluster flag)
SWCC or CLI, used with RA8000, MA8000, ESA12000, and EMA12000	NetWare Client (for access to NetWare volumes through NCP)
Array Controller Software (ACS) 8.5 or 8.6 used with RA8000, MA8000, ESA12000, and EMA12000	NCS 1.6 (up to 12 nodes supported; greater than 12 node support provided by Professional Services)
SANworks Secure Path	Cluster Migration Utility
Offline ACU	NRM (Formerly NetWare Portal Management)
CPQONLIN	

Learning Check

1. Shared storage is being added to an RA4100. What three tasks are required to increase the shared storage capacity?
.....
.....
.....
2. After increasing the shared storage for the cluster, what two NetWare utilities are used to ensure that the cluster servers recognize the same drives and free space?
 - a. ConsoleOne
 - b. NWDEPLOY
 - c. MONITOR
 - d. NWADMIN.EXE
 - e. Novell Remote Manager
3. Where is thenwdeploy.exe file located?
 - a. SYS:\PUBLIC\mgmt
 - b. SYS:\SYSTEM
 - c. SYS:\NOVELL
 - d. The NetWare 6 Product CD

4. A node on a cluster has failed and needs to be replaced. NetWare has been reinstalled and the node has been given a new unique name and IP address to ensure that there are no NDS conflicts. When running NCS to install the node, it will not retake its former place in the cluster. Why not?
 - a. The node name and IP address must be the same as the replaced node.
 - b. Go into DSREPAIR and rediscover the MAC address of the network cards in the new server.
 - c. From Partition Manager, go in and delete and recreate any partitions on the server.
 - d. The server must contain at least an RW partition of [ROOT] to be recognized as a member of the cluster.
5. List four tools provided by Novell to troubleshoot and repair performance problems between nodes.

.....

.....

.....

.....
6. List three Compaq troubleshooting tools to be used with ProLiant clusters.

.....

.....

.....

7. There are client and agent problems. After adding a list of client names to the agent's list of authorized clients, and setting the password in the Password Dialog box, what is done next?
 - a. Start and stop the Fibre Channel interconnect device.
 - b. Stop the node, ensure it fails over, and restart the node.
 - c. Stop and restart NODE.NLM.
 - d. Unload and reload cpqspagt.
8. There are problems with a shared storage solution using Fibre Channel. The cluster nodes, storage subsystem, and storage hubs or switches are powered off. In what order should the devices be powered on?
 - a. ProLiant server, storage subsystem, Ethernet hub, storage hub
 - b. Storage hub or switch, storage subsystems, Ethernet hubs or switches, ProLiant servers.
 - c. Storage subsystem, storage hub or switch, ProLiant server Ethernet hub
9. Where will the SYS: volume be located when installing a new NetWare 6 server that is going to be part of a cluster solution?
 - a. On the SCSI drive with the lowest LUN in the remote storage subsystem.
 - b. On any drive, local or remote, as long as the drive uses Ultra SCSI technology or better.
 - c. On a local drive.
 - d. On any drive.

Learning Check Answers

Module 1 Learning Check

1. Which availability level (AL) allows users to stay online although current transactions might need restarting and there is a possibility of performance degradation?

AL3

2. Name the two classes of critical applications.

a. Mission-critical

b. Business-critical

3. What methods of high availability should you attempt before clustering should be implemented?

a. Fault prevention

b. Fault tolerance

4. Name the three basic cluster models.

a. Shared-nothing

b. Shared-disk

c. Shared-everything

5. A company has two servers. One server is busy running all the applications that the company uses for business. The other server has only a few, rarely used applications. If the servers are clustered with application failover capabilities, the system administrator can manually fail over some of the applications to the other server to redistribute the workload. This is an example of what type of cluster advantage?

Load balancing

6. Which component is the least likely to fail in a network environment?
Hardware
7. Name three platforms supported by Compaq ProLiant Clusters.
 - a. Windows 2000 Cluster service
 - b. Novell Cluster Services (NCS)
 - c. LifeKeeper for Linux
 - d. UnixWare NonStop Clusters
 - e. Oracle Parallel Server (OPS) for Oracle 8i
 - f. Compaq Parallel Database Cluster (PDC) for Oracle 9i Real Application Cluster (RAC)
8. Where can you find technical information on all industry-standard, high-availability solutions for Compaq clusters?
<http://www.compaq.com/highavailability>

Module 2 Learning Check

1. List five things that should be considered when planning a cluster.
 - a. Type of business using the cluster
 - c. Business function provided by the cluster and the impact on the business of any individual server failing
 - d. Number of users accessing the system and whether this number will increase or decrease in the future
 - e. Storage capacity required
 - f. Physical layout of the cluster
 - g. Percentage of downtime users can accept and still perform core business functions
 - h. Primary tasks of the servers, such as administration, storage, or messaging
 - i. Aggregate resource requirements for the servers
 - j. Workload failover requirements (can the user separate the workload that relies on failover from the workload that can tolerate failure) and load balancing capabilities
 - k. Enterprise backup/restore options
2. Which of the following ProLiant servers would be appropriate for a clustering solution?
 - a. ProLiant DL320
 - b. ProLiant DL380 G2
 - c. ProLiant DL580
 - d. ProLiant ML350
 - e. ProLiant ML510
 - f. ProLiant ML570

3. List five features of the Compaq Intelligent Fault Resilience.
 - a. Hot-pluggable hard drives
 - b. PCI Hot Plug slots
 - c. Error checking and correcting (ECC) memory
 - d. Online spare memory
 - e. Redundant power supplies
 - f. Redundant processor power modules
 - g. Redundant cooling fans
 - h. Easy installation through SmartState
 - i. Manageability through Compaq Insight Manager
4. In a cluster, what does shared storage provide?
 - a. Larger capacity storage devices
 - b. Ease of backup
 - c. 99.999% uptime
 - d. Continued access to critical information if one of the servers in the cluster fails
 - e. Fault tolerant storage
5. Which RAID level provides the lowest risk of data loss?
 - a. RAID 0
 - b. RAID 1+0
 - c. RAID 4
 - d. RAID 5
 - e. RAID 3/5



Note

RAID 1+0 provides the lowest risk of data loss, however 50% of the storage capacity is lost providing RAID capabilities.

6. SCSI priorities are based on SCSI ID's. Which SCSI ID has the highest priority?
 - a. 0
 - b. 1
 - c. 15
 - d. 8
 - e. 7
7. When ranking SCSI priorities, which application should receive the highest priority?
 - a. NetWare Loadable Modules, executable files and system files
 - b. Log files used to rebuild transactions after a critical failure
 - c. Database files
 - d. Temporary files and TTS files.
8. What are the three main components of a Novell Cluster Services cluster?
 - a. Servers
 - b. Storage System
 - c. Interconnection between servers and storage system
9. Which of the following are factors to use when deciding the amount of shared storage to provide in a NCS cluster?
 - a. Client operating systems
 - b. Printer configurations
 - c. Clustered applications and dependencies
 - d. Type of interconnection



Note

The factors to be considered when planning the required amount of shared storage disk space include the amount of space required for all clustered applications and their dependencies, as well as the level of data protection required for the types of data used by the clustered applications.

10. When configuring cluster interconnects, they must be configured to provide maximum data availability. How should this be done?

There should be two distinct data paths. The first should run from the Fibre Channel host bus adapter (HBA) to the Fibre Channel interconnect device, and the second should run from the Fibre Channel interconnect device to the Fibre Channel array.

11. How should a tape backup solution running on a server that is not a member of a cluster be configured?
- a. The tape backup software should be configured to point to the IP address of the shared storage device.
 - b. The tape backup server should be configured to back up the cluster drives as a logical network drive accessed through a file share, pointing to a physical drive.
 - c. The tape backup software should be connected through a cluster volume name or share. The software is pointed to a virtual cluster drive resource and not a physical drive.
 - d. Through a drive configured using the MAP command in a NetWare login script.
12. Which fiber switch is only supported in the RA4100 platform?
- b. Fibre Channel Storage Hub 7/12
 - c. Compaq FC-AL Switch
 - d. Compaq SAN Switch 8-EL
 - e. Compaq SAN Switch 8

Module 3 Learning Check

1. Novell Cluster Services (NCS) 1.6 supports up to _____ nodes. The Compaq ProLiant Cluster Solutions for NetWare currently supports up to _____ nodes.
 - a. 32 and 12
 - b. 12 and 32
 - c. 24 and 8
 - d. 8 and 24
2. Which three of the following are features of the Compaq ProLiant Cluster Solutions for NetWare?
 - a. High availability at a minimum cost
 - b. Operating system independence
 - c. Support for the entire line of ProLiant servers
 - d. Integrated hardware and software solutions
 - e. Protection of investment
 - f. Five-year on-site 7/24 warranty with 4 hour response time guaranteed.
3. The Compaq ProLiant DL380 G2 Packaged Cluster Solution can house up to _____ Ultra 3 hot-plug hard drives.
 - a. 32
 - b. 12
 - c. 16
 - d. 14
 - e. 8
4. The Compaq ProLiant DL380 G2 Packaged Cluster Solution can support _____ shared storage space.
 - a. 750MB
 - b. 750GB
 - c. 900GB
 - d. 1TB
 - e. Over 1TB

5. One key difference between the ProLiant Cluster HA/N100 Fibre Channel solution and the ProLiant Cluster HA/N200 solution is:
 - a. The HA/N200 supports up to the full 32 nodes.
 - b. The HA/N100 offers full redundancy.
 - c. The HA/N200 has dual Fibre Channel array controllers and dual HBAs.
 - d. The HA/N100 only supports NetWare 6.
6. The highest level of data availability in the ProLiant cluster family is provided by the:
 - a. ProLiant Cluster HA/N200
 - b. ProLiant Cluster HA/N500
 - c. ProLiant Cluster HA/F200
 - d. ProSignia Cluster HA/F500



Note

The ProLiant Cluster HA/N500 can be configured so there is no single point of failure in the storage subsystem. SANworks Secure Path software has to be used to enable support for redundant 64-bit/66MHz PCI-to-Fibre Channel HBAs in each server.

7. The storage management software NSS 3.0 will support which of the following options?
 - a. Management of traditional NetWare volumes, including the ability to resize volumes and make the volume smaller
 - b. Remote management of file system security settings using Active Directory.
 - c. Manual setting of Shareable for Clustering flag
 - d. Backward compatibility to NetWare 3.2
8. Which software does **not** come from Compaq?
 - a. Compaq Insight Manager 7
 - b. SWCC
 - c. NWAdmin
 - d. ACS
 - e. CLI

9. What is the most direct interface to the HSG80 controller?
- a. ACU
 - b. SWCC
 - c. CPQONLIN
 - d. RAID
 - e. CLI

Module 4 Learning Check

1. Describe how shared storage provides a high-availability solution.
It uses an automatic failover procedure. If one server in the cluster should lose access to the shared storage and the network, the other server assumes the responsibilities of the downed server.
2. List four features of the Compaq ProLiant Cluster Solution for NetWare?
 - a. Supports up to 12 nodes
 - b. Supports full FC-SW and FC-AL SANs
 - c. Is based on an architecture in which clustered servers share access to a common set of hard drives
 - d. Requires all shared data be stored in an external storage system for the data to be readily accessible from each cluster node
3. The _____ utility found on the SmartStart CD is used to configure the shared storage array after the installation process.
cpqonlin
4. What is the minimum amount of suggested free space that should be allocated for the Cluster Services partition?
 - a. 500MB to 1GB
 - b. 500GB to 1TB
 - c. 15MB to 25MB
 - d. 100MB to 500MB
5. Put the following steps in the order that you would start a Fibre Channel Storage System.
 - b. Start storage hubs or switches.
 - c. Fibre Channel arrays.
 - a. Start servers.

6. List the steps for configuring the ProLiant DL380 Packaged Cluster to be sent to another location.
 - a. Unpack the cluster
 - b. Install the internal or external options
 - c. Cable the system
 - d. Configure the servers by installing the operating system
 - e. Configure the shared storage
 - f. Install the cluster software
 - g. Verify the installation
 - h. Pack the cluster for shipment
 - i. Deploy in a rack
7. What utility is used to configure the ProLiant DL380 Packaged Cluster?
 - a. ORCA
 - b. cpqonlin
 - c. Console One
 - d. Install.NLM
 - e. ACU
8. A complete failure in the shared storage subsystem can be catastrophic. List the two areas where failures can occur.

Failure can occur in either the cluster node interconnect or network connection, or in the path between the cluster nodes and the storage subsystems.
9. Which of the following components have redundant options in the ProLiant DL 380 G2 Packaged Cluster?
 - a. Processor
 - b. Memory
 - c. Power Supply
 - d. Fan
 - e. Shared Storage Controller

10. List two ways to increase the availability of the I/O path to the cluster shared storage in a high-availability Fibre Channel cluster configuration.
 - a. Add redundant Fibre Channel Host Bus Adapters (HBAs) and array controllers.
 - b. Add a separate, redundant arbitrated loop or switched fabric.
11. What are the three main components to SANworks SecurePath?
 - a. ACU
 - b. Client-based agent
 - c. Server-based agent
 - d. Enhanced CPQFC.HAM driver
 - e. SPM
 - f. CLI
12. If a system in the Secure Path storage profile does not respond, what is the path state set to in the Physical Path view?
 - a. Down
 - b. Inactive
 - c. Unresponsive
 - d. Failed
 - e. Off line

Module 5 Learning Check

1. Which of the following best describes the supported network cards for Novell Cluster Services 1.6?
 - a. Any NIC certified for NetWare 6
 - b. All token ring cards
 - c. All Compaq network cards
 - d. Any Ethernet card
 - e. Any NIC that has a LINUX driver
2. List the two features of Compaq Advanced Network Services that can be configured to minimize downtime caused by a NIC failure.
 - a. Two Compaq Ethernet adapters
 - b. Two ports on a single adapter, with one port running in hot standby mode
3. Choose two network card drivers that work with the Compaq Advanced Network Services driver (CPQANS.LAN).
 - a. NE2000.LAN
 - b. N1000.LAN
 - c. TOKEN.LAN
 - d. N10.LAN
 - e. N100.LAN
4. How many NICs can be used to provide automatic redundancy with the Compaq Advanced Network Services?
 - a. 1
 - b. 2
 - c. 2 to 4
 - d. 2 to 8

5. What is ALB?
 - a. Automatic Link Bridging
 - b. Automatic Load Balancing
 - c. Automatic Load Bridging
 - d. Automatic Link Balancing
6. What is Network Adapter Teaming?

Network adapter teaming combines two or more physical NICs into a single logical NIC with multiple IP addresses.
7. Which of the following are types of NIC teaming supported by Compaq ProLiant Ethernet network adapters? Select all that apply.
 - a. SFT
 - b. NFT
 - c. ALB
 - d. TLB
 - e. SLB
8. What are the two most common reasons for network adapter failover?
 - a. Dust
 - b. Physical link failure
 - c. Heartbeat failure
 - d. Static electricity
 - e. Power surge
9. Which communication protocol is used to transport clustering data?
 - a. TCP/IP
 - b. IPX
 - c. SPX
 - d. RIP
 - e. NLSP

Module 6 Learning Check

1. List three important features of NCS that help ensure and manage the availability of network resources.
 - a. Scalability
 - b. Online cluster software updates
 - c. NSS 3.0
 - d. Improved volume mount protection
 - e. ConsoleOne client
2. What are the objects that belong to the cluster container?
 - a. Cluster resources
 - b. Resource templates
 - c. Volume resources
 - d. Cluster nodes
 - e. Master IP address
3. Cluster resources are specifically defined as either applications or _____.
Services.
4. You are using either load or unload scripts and notice that Cluster Resource Manager is generating warning messages when a resource joins or leaves a cluster. What could be a possible cause for these warning messages?
Load and unload scripts can have no more than 600 characters each. Check to ensure that you have not exceeded this limit.

5. A cluster resource is assigned to nodes 1, 2, and 3 in that order. Nodes 1 and 2 are offline. Node 2 comes back online. Node 3 is not automatically failing back to node 2. Why not?

The resource will fail back automatically only to the first node listed on the Nodes Property page.

6. What is the process for creating a cluster resource template?
- Set preferred node assignments
 - Set automatic or manual failover and failback mode
 - Configure Load and Unload scripts
7. What are the three eDirectory objects created for a cluster-enabled volume?
- Cluster Volume
 - Cluster Volume Resource
 - Cluster Virtual Server
8. You have set the quorum membership field to 3. This is the number of nodes currently defined in the cluster. When bringing up the servers in the cluster, NCS reads this field, waits until the servers are up, and attempts to load resources. It cannot do so. Explain why.

There must be four servers up in the cluster before any resource will load and start.

9. What is the default time between transmits for all nodes in the cluster except the master?
- 8
 - 6
 - 1
10. What property specifies the amount of time the master node waits for all other nodes in the cluster to signal that they are alive?
- Master watchdog
 - Slave watchdog
 - Heartbeat
 - Tolerance

11. When changes are made to cluster objects in eDirectory, the Cluster Configuration Library updates the local version of the configuration on:
 - a. The master node, which sends the updates to slave nodes
 - b. The master and slave nodes at the same time
 - c. Only the master node, which does not forward the updates to the slave nodes
 - d. Its own local drive. It does not send updates anywhere.
12. The Cluster Management Agent (CMA.NLM) acts as a proxy for the ConsoleOne management application, allowing ConsoleOne to perform what functions?
 - a. View the current location of cluster resources
 - b. Add, move, or remove cluster resources
 - c. Create or modify Load and Unload scripts, and script timeout parameters
 - d. Take resources offline or bring them online
 - e. Specify failover and failback policies for cluster resources
 - f. Invoke manual failover and failback
13. In restoring an SBD partition, ConsoleOne has been used to delete the old SBD. All NCS NLMs have been unloaded using `uldncs.ncf`. What is the last step to be done before issuing the SBD install command?

You must load clstrlib and vll manually.
14. List the six main steps to install NCS.
 - a. Ensure that you use the latest Compaq Support Paq (cpqdpjoy.nlm).
 - b. Configure the shared storage with ConsoleOne or NRM.
 - c. Install the NCS software.
 - d. Add NSS-based pools and volumes to the cluster.
 - e. Create or select cluster resource templates.
 - f. Create or modify the cluster resources.
15. When upgrading to NetWare 6, the drivers are loaded, eDirectory is upgraded, and one other step is automated before the files are copied to the server. What is that step?

NetWare 6 information is added to the autoexec.ncf and startup.ncf files.

Module 7 Learning Check

1. List at least three features that Compaq Insight Manager delivers to cluster monitoring.
 - a. Device management
 - b. Event management
 - c. Intelligent monitoring
 - d. Failure and pre-failure alerts
 - e. Remote maintenance
 - f. Visual control
 - g. Cluster monitoring
 - h. Software deployment
2. If using Cluster Monitor queries to act as a filtering mechanism, how many clusters can the user see after running a query?
 - a. As many clusters as there are on the network
 - b. As many clusters as there are in the tree
 - c. As many clusters as there are in the container
 - d. Only the clusters defined in their scope
3. In trying to use ConsoleOne to manage and monitor a NetWare cluster, the appropriate cluster located on another interconnect cannot be accessed. What might be the problem?
 - a. The right version of ConsoleOne might not be employed.
 - b. The ConsoleOne workstation must reside on the same network as the cluster interconnect.
 - c. JAVA.NLM might be unloaded

4. ConsoleOne is being used to check the status of servers and resources. The cluster object in the left pane has been highlighted and the View Menu selected. There are now three choices. Which one is used to check the status of servers and resources?
 - a. Cluster State View
 - b. Console View
 - c. Partition and Replica View
5. ConsoleOne is being used to work with cluster resources. The cluster object in the left hand pane has been highlighted and the view menu selected. There are three choices. Which one is used to create clustered volumes?
 - a. Cluster State
 - b. Console View
 - c. Partition and Replica View
6. In the Cluster State view in ConsoleOne, what would the color of the icon be if a resource is unassigned?
 - a. Green
 - b. Yellow
 - c. Red
 - d. Grey
 - e. There would be no icon if the resource is unassigned.

7. NetWare Remote Manager is accessed from where?
 - a. A web browser
 - b. Server console
 - c. SYS:\System
 - d. SYS:\PUBLIC
8. List three components of NCS Management and Administration.
 - a. Managing network clients, and load balancing
 - b. Managing clusters without interrupting cluster services
 - c. Managing a cluster in a degraded condition
9. List some maintenance and monitoring procedures that must be performed on a regular basis to ensure effective management of the entire cluster.
 - a. Upgrade the hardware components.
 - b. Upgrade the software.
 - c. Increase the storage capacity.
 - d. Restructure the cluster groups.
 - e. Back up the cluster data.

Module 8 Learning Check

1. List the benefits of a cluster-aware application.
 - a. Increase efficiency
 - b. Active/active configurations
 - c. Automatic creation and configuration of cluster resources and groups
 - d. Self-monitoring
 - e. Ability to run separate instances of a database on multiple nodes
2. Noncluster-aware applications cannot detect cluster software and behave in the same way as on a stand-alone server. Which of the following tasks will create an active/standby configuration?
 - a. Contact the vendor for a cluster plug-in.
 - b. Go to download.novell.com for the latest application patch.
 - c. Write load and unload scripts.
 - d. Manually create a cluster resource in NDS.
3. If a printer is directly connected to the parallel port of a node in a cluster environment, what happens to printing services on that printer if the node fails?
 - a. The printer will be out of service.
 - b. The physical connection will be serviced by the other nodes.
 - c. Nothing, the printer will keep printing.
4. When configuring GroupWise post offices and domains to be on shared volumes for failover and failback, should the shared volumes for GroupWise be cluster-enabled to function with NCS?

No, cluster-enabling shared volumes is not recommended and not necessary for GroupWise to function with NCS.
5. Explain why the autoexec.ncf file should **not** be modified to automatically start the grpwise.ncf file when launching GroupWise.

GroupWise must be launched from within a cluster resource load script.

Module 9 Learning Check

1. Shared storage is being added to an RA4100. What three tasks are required to increase the shared storage capacity?
 - a. Install and connect the hardware.
 - b. Configure the logical arrays.
 - c. Create and cluster-enable NSS volumes.
2. After you have increased the shared storage for the cluster, what NetWare utilities can you use to ensure that the cluster servers recognize the same drives and free space.
 - a. ConsoleOne
 - b. NWDEPLOY
 - c. MONITOR
 - d. NWADMIN.EXE
 - e. Novell Remote Manager (NRM)
3. Where is thenwdeploy.exe file located?
 - a. SYS:\PUBLIC\mgmt
 - b. SYS:\SYSTEM
 - c. SYS:\NOVELL
 - d. The NetWare 6 Product CD

4. A node on a cluster has failed and needs to be replaced. You reinstall NetWare and give the node a new unique name and IP address to ensure that there are no NDS conflicts. When you try to run NCS to install the node, it will not retake its former place in the cluster. Why not?
 - a. The node name and IP address must be the same as the replaced node.
 - b. Go into DSREPAIR and rediscover the MAC address of the network cards in the new server.
 - c. From Partition Manager, you have to go in and delete and recreate any partitions on the server.
 - d. The server must contain at least an RW partition of [ROOT] to be recognized as a member of the cluster.
5. List four tools provided by Novell to troubleshoot and repair performance problems between nodes.
 - a. iMonitor
 - b. NetWare logger screen
 - c. Conlog.NLM
 - d. Config.NLM
6. List three Compaq troubleshooting tools to be used with ProLiant clusters.
 - a. Compaq Fibre Channel Fault Isolation Utility (FFI)
 - b. SANworks Secure Path for NetWare
 - c. QuickFind 2000
7. There are client and agent problems. After adding a list of client names to the agent's list of authorized clients, and setting the password in the Password Dialog box, what is done next?
 - a. Start and stop the Fibre Channel Interconnect device.
 - b. Down the node, make sure it fails over, and restart the node.
 - c. Stop and Restart NODE.NLM.
 - d. Unload and reload cpqspagt.

8. There are problems with a shared storage solution using Fibre Channel. The cluster nodes, storage subsystem, and storage hubs or switches are powered off. In what order should the devices be powered on?
 - a. ProLiant server, storage subsystem, Ethernet hub, storage hub
 - b. Storage hub or switch, storage subsystems, Ethernet hubs or switches, ProLiant servers
 - c. Storage subsystem, storage hub or switch, ProLiant server Ethernet hub
9. Where will the SYS: volume be located when installing a new NetWare 6 server that is going to be part of a cluster solution?
 - a. On the SCSI drive with the lowest LUN in the remote storage subsystem.
 - b. On any drive, local or remote, as long as the drive uses Ultra SCSI technology or better.
 - c. On a local drive.
 - d. On any drive.

Compaq ProLiant CL380 Packaged Cluster

Appendix A

ProLiant CL380 Packaged Cluster



The ProLiant CL380 is a packaged two-node cluster that simplifies clustering for business-critical applications in branch offices, remote locations, or departmental computing—wherever space is scarce, technical support minimal, and simplicity of operation important.

The ProLiant CL380 uses:

- Two ProLiant server nodes.
- Dedicated Ethernet interconnect.
- Shared SCSI storage subsystem in a single cabinet.

Configure and manage these clustering solutions using Compaq installation and systems management utilities.

INTERNET For ProLiant CL380 technical specifications, go to:
<http://www.compaq.com/products/servers/platforms.html>

The ProLiant CL380 supports a maximum of two nodes. Additional nodes cannot be added.



ProLiant CL380 Cluster

The ProLiant CL380 offers

- Support for
 - NetWare Cluster Services (NWCS) for NetWare 5.x
 - Novell Cluster Services (NCS) for NetWare 6
 - Microsoft Windows 2000 Advanced Server Cluster service
 - Microsoft Windows NT Server 4.0 Enterprise Edition
 - Red Hat Linux 3.x or later
- Shared SCSI
 - More than 100GB of SCSI internal shared storage
 - Expandable with six 1-inch hot-pluggable Ultra2 or Ultra3 hard drives
 - RAID-protected disk subsystems
 - Intel Pentium III processors
 - Redundant Shared Storage RAID CR3500 controller (one ships standard)

- Compaq software
 - Compaq Insight Manager (all versions)
 - SmartStart
 - Array Configuration Utility (ACU) with embedded array controllers
 - CR3500 Configuration Utility
- Integrated keyboard, monitor, and mouse (KMM)-switch

Compaq ProLiant Server Nodes

The server nodes used in ProLiant CL380 clusters are based on modified ProLiant DL380 servers. Although the general characteristics are similar, these nodes are not ProLiant DL380 servers.

Shared SCSI Storage Subsystem

The shared storage subsystem of the ProLiant CL380 potentially encompasses up to six internal shared drives, an optional external storage subsystem, and up to two CR3500 RAID controllers in a shared SCSI bus configuration.



Note

No cluster kits are required for the ProLiant CL380 packaged cluster solution.

Shared SCSI Bus Configuration

Unlike the other ProLiant Cluster solutions for Novell, the ProLiant CL380 clusters use a shared SCSI configuration for their shared storage subsystems.

The shared SCSI storage subsystem in the ProLiant CL380 supports:

- Fourteen 1-inch Wide Ultra2 and Wide Ultra3 SCSI drives.
- Two CR3500 RAID controllers in the shared storage area.


The physical location of the hot-pluggable drive determines the SCSI ID.

Three configurations for the shared storage are possible when an optional StorageWorks Enclosure Model 4214 is attached to the system. The first two configurations use a single-bus Enclosure Model 4214, and the third configuration uses an Enclosure Model 4214 with a 4200 dual-bus option.

Shared Storage Capacity and External Storage

The ProLiant CL380 supports up to 14 shared storage drives.

- The internal shared storage area of the ProLiant CL380 supports up to six 1-inch Wide Ultra2 and Wide Ultra3 SCSI drives.
- Server nodes in the ProLiant CL380 include two half-height 5.25-inch removable media bays and a hot-pluggable boot drive cage that supports up to two 1-inch Wide Ultra2 and Wide Ultra3 SCSI drives.



ProLiant CL380 with External Storage

- Connect a member of the Compaq StorageWorks Enclosure 4200 family to the ProLiant CL380 shared storage area to host up to eight more 1-inch Wide Ultra2 and Wide Ultra3 SCSI drives in a single-bus configuration for a total of 14 drives.

The addition of the Model 4200 Enclosure provides a possible 20 physical drive bays. However, only 14 drives can be configured at a time because of the number of available SCSI IDs.

- Using a dual-bus configuration, the external storage enclosure supports up to seven additional drives for a total of 13 drives. Because the 4200 dual-bus service uses SCSI ID 15, this ID is no longer available for a drive. Distribution of the remaining SCSI IDs is similar to the single-bus configuration.

CR3500 RAID Controllers

The ProLiant CL380 supports up to two CR3500 RAID controllers in the shared storage area. The CR3500 RAID controller manages the shared storage subsystem and contains a read/write cache.

If a controller that is not part of a redundant pair is operating normally, removing the controller will cause any data contained in the cache to be lost.

One CR3500 RAID controller ships with the system and an optional CR3500 is available. Adding a second RAID controller makes them both hot-pluggable and redundant, but not hot-replaceable.

The additional CR3500 also allows for an active/active redundant controller configuration with automatic controller failover and failback and helps eliminate single points of failure.

The configurations supported by the CR3500 RAID controller include:

- Single-device arrays
- Striped arrays (RAID 0)
- Mirrored arrays (RAID 1)
- Striped mirrored arrays (RAID 1+0)
- Striped parity arrays with a fixed parity disk (RAID 4)
- Striped parity arrays with a floating parity disk (RAID 5)

Data Protection with ROC

An optional feature of the ProLiant CL380 is the Integrated Smart Array Controller, also called *RAID on a Chip* (ROC). This is an intelligent array controller for entry-level, hardware-based fault tolerance with support for up to six Wide Ultra2 and Wide Ultra3 SCSI internal hard drives.

ROC provides a cost-effective alternative to software-based RAID. It offers worry-free data protection for all server internal storage requirements of the ProLiant CL380.

ROC allows increased logical drive capacity inside the server without:

- Taking production systems offline.
- Powering down to install additional capacity.

ROC reduces reconfiguration downtime by allowing for the changing of RAID levels without:

- Taking data protection levels offline.
- Backing up to tape.
- Rebuilding.

CR3500 Configuration Utility

The CR3500 RAID controller manages the shared storage subsystem on the ProLiant CL380. To configure this device, Compaq provides the CR3500 Configuration Utility on the Compaq SmartStart and Support Software CD.

Network Adapters Supported

The network adapters supported for the ProLiant CL380 are:

- NC3163 Fast Ethernet NIC Embedded (standard)
- NC3123 Fast Ethernet NIC PCI 10/100 (standard)
- NC3134 Fast Ethernet NIC 64 PCI Dual Port 10/100 (optional)
- NC6136 Gigabit Server Adapter, 64-bit/66MHz, PCI, 1000 SX (optional)
- NC7131 Gigabit Server Adapter, 64-bit/66MHz, PCI, 10/100/1000-T (optional)

Configuring the ProLiant CL380 for High Availability

The ProLiant CL380 supports up to two CR3500 RAID controllers in the shared storage area. One CR3500 RAID controller ships with each system and an optional CR3500 can be added.

Eliminating Significant Single Points of Failure

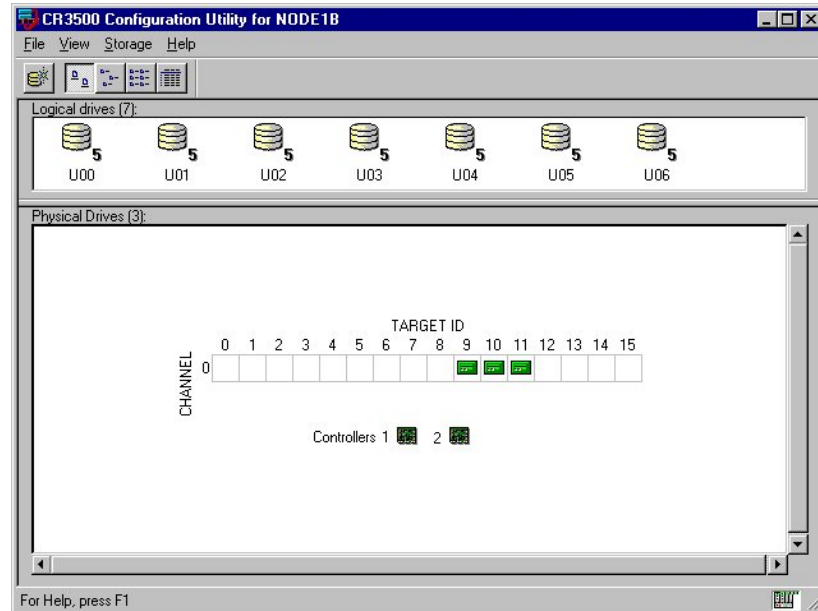
Eliminate significant points of failure in the shared storage area of the ProLiant CL380 by installing and configuring a second CR3500 RAID controller. This addition makes both RAID controllers hot-pluggable and redundant. It also allows for an active/active configuration with automatic controller failover and failback.



Important

The CR3500 RAID controller has a read/write cache. If a controller that is not part of a redundant pair is operating normally, removing the controller will cause any data contained in the cache to be lost.

CR3500 Configuration Utility



The CR3500 Configuration Utility is a graphical user interface (GUI) based utility that allows for easily created logical drives and for setting RAID fault-tolerance levels for shared storage hard drives. The CR3500 Configuration Utility also configures one or two CR3500 Shared Storage RAID controllers.

A separate system running a Microsoft Windows operating system is required to use the CR3500 Configuration Utility. The procedure uses a serial cable to connect to the CR3500 RAID controller through a connection in the back of the shared storage area.

Use the CR3500 Configuration Utility to:

- Initially configure the array controller.
- Reconfigure the array controller.
- Add additional disk drives to an existing configuration.
- Expand capacity.

The CR3500 Configuration Utility:

- Displays the controller configuration in an easy-to-understand graphical format.
- Contains configuration wizards for the configuration process.
- Provides detailed information about the physical drives configured in the storage subsystem.
- Provides information about the array, including:
 - Power
 - Temperature
 - Fan state



Note

The CR3500 Configuration Utility is unique to the shared SCSI storage used by the ProLiant CL1850 and CL380 clusters. The Compaq Array Configuration Utility used with other Compaq cluster solutions is for Fibre Channel storage and will not work with a ProLiant CL1850 or CL380 cluster.

The ProLiant CL380 cluster can be configured with a StorageWorks 4214 Storage Enclosure for additional storage capacity.

Upgrading from a Single Controller to Redundant Controllers

To upgrade the system from a single CR3500 RAID controller to a redundant CR3500 RAID controller configuration:

1. Use the Save Configuration function of the CR3500 Configuration Utility to save the existing single controller configuration.
2. Suspend all I/O and power off the shared storage subsystem.
3. Install the redundant controller in the bottom controller slot of the shared storage subsystem.
4. Power on the shared storage subsystem.
5. Restore the controller configuration settings saved in step 1.


Configuration of the CR3500 RAID controller must be performed through a serial connection for ProLiant CL380 clusters running NetWare.



Note

When restoring configuration settings using a serial connection, both controllers must be attached to the host machine.

Serial Connection



Serial Connection Locations on ProLiant CL1850 and CL380 Clusters

Configure the shared storage subsystem with a separate computer running Windows 95/98, Windows NT, or Windows 2000. Connect the Windows system to the port labeled C-1 on the back of the shared storage area using a serial cable. If using redundant controllers, connect another serial cable from the Windows system to the port labeled C-2.

Updating Firmware

Firmware updates for the CR3500 RAID controller can be obtained from the Compaq website. To update the firmware in the storage subsystem:

1. Download the latest version of firmware from the Compaq website.
2. Connect to the CR3500 RAID controller through the serial connection with HyperTerminal, a program that ships with Windows operating systems.



Important

Any I/O must be stopped during the firmware update.

Upgrading Firmware in a NetWare Environment

Follow these instructions to perform a firmware upgrade on a CR3500 RAID controller:

1. Connect the controller to a system running Windows 95 or Windows NT through the serial connector on the back of the shared storage subsystem.
2. Use a terminal program, such as HyperTerminal, to connect.
3. Use the following settings to connect:
Bit per Second = 9600, Data Bit = 8, Parity = None, Stop Bits = 1, Flow Control – XON / XOFF
4. Power on the shared storage area and press *Ctrl-C* to stop the process.
5. Select *option 2* to change the baud rate.
6. Select *option 3* to set baud rate to 38400 baud.
7. Change the baud rate on the terminal to match the 38400 baud settings of the controller.
8. Disconnect, then reconnect the terminal service.
9. Press *Enter* to reestablish communication with the controller.
10. Select *option 1* and press *Enter*.
11. Transfer the firmware file as a text file from the terminal service. The update process takes approximately 15 minutes at 38400 baud.
12. After the firmware update, select *option 2* and change the baud rate to 9600 by selecting *1*.
13. Change the baud rate to 9600 on the terminal service.
14. Disconnect, then reconnect the terminal service.
15. Press *Enter* to reestablish communication with the controller.
16. Select *option 9* to restart the controller.

Comparing Hubs and Switches

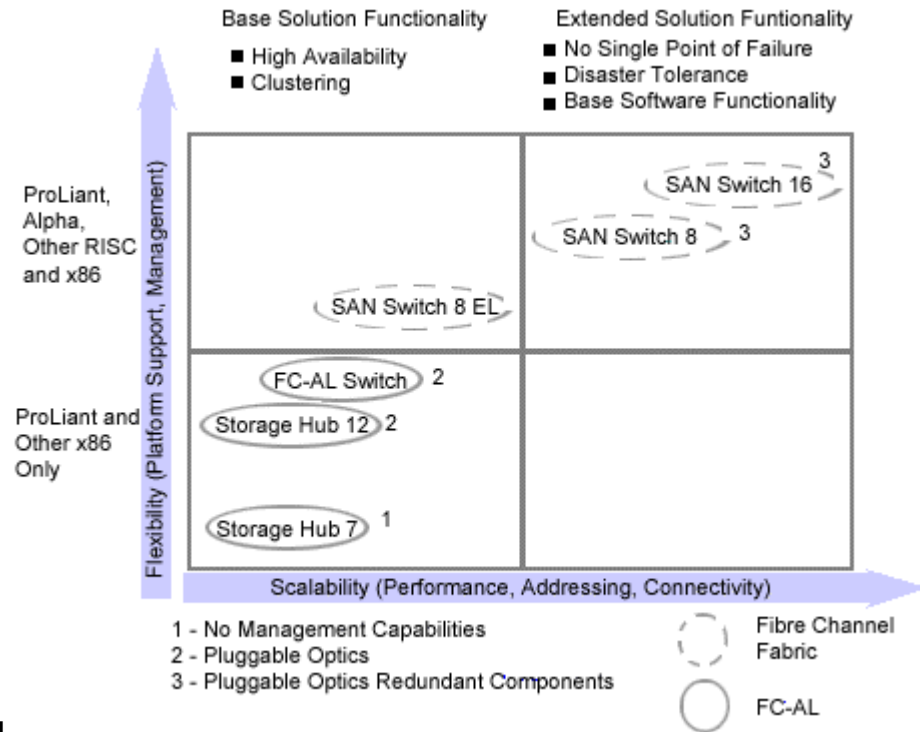
Each of the two main interconnect devices in a Fibre Channel environment has its own advantages and environment in which the hub or switch best meets the requirements of the customer. The following table lists the differences between hubs and switches in a Fibre Channel environment.

Hubs	Switches
Implement an FC-AL topology with shared bandwidth among all devices.	Implement a FC-SW topology, with ports for fabric devices or arbitrated loop devices
Provide connection to a shared 100MB/s FC-AL transfers for single 100MB transfers at a time.	Provide 100MB/s point-to-point connection between devices for multiple 100MB transfers at a time with no multiple device arbitration time.
Decrease performance as nodes are added because of additional arbitration time.	Have no performance reduction as nodes are added.
Allow for nodes on the Fibre Channel loop to see all traffic going between other nodes.	Allow nodes on the Fibre Channel system to see only data destined for themselves.
Provide more complex error recovery because of loop initializations.	Isolate individual nodes from reconfiguration and error recovery of other nodes within the SAN.
Have more chances for traffic disruption on reconfiguration and for errors because of loop initializations.	Provide fewer chances for traffic disruption on reconfiguration and fewer errors because control features manage loop initialization events.
Have optional management features of Compaq 12-port hubs.	Provide management control of the Fibre Channel infrastructure.
Provide a seven-step loop initialization process (LIP). If a device encounters an error, it initiates a LIP. The loop is unavailable for data transfer during the error recovery.	Provide greater stability during error recovery.

The following table lists the features and recommendations for each interconnect device.

Interconnect Device	Features and Recommendations
Hub	Recommended for direct-attached storage and homogeneous clusters where lowest cost per port is important.
FC-AL switch	Recommended for easy-to-deploy, affordable, high-performance, x86-based SANs (heterogeneous operating system). Best suited for workgroup and departmental environments. Enable future connectivity to a fabric switch. Provide low cost per port.
FC-SW	Recommended for larger, highest-performance, heterogeneous (x86 and RISC) fabric-based SANs with additional levels of business continuance (mirroring and disaster tolerance). Have a higher cost per port than the FC-AL switch, but with a higher level of scalability and functionality.

SAN Infrastructure



Positioning

When selecting a storage solution, consider current and future storage needs. Storage and networking requirements determine which interconnect device is best suited for this system.

The following factors influence the interconnect device requirements:

- Scalability
- Flexibility
- Base solution functionality
- Extended functionality
- Budget

Identify the budget and performance criteria, platform support criteria, and scalability and addressing criteria before selecting an interconnect.

The goal is to provide workgroups, departmental, and enterprise environments a choice of interconnect solutions that will meet overall application and future growth needs. Affordability is an important purchase consideration in departmental applications.

References and Additional Help

Refer to the following Compaq documents for more information regarding Compaq Clusters for NetWare:

- User guides for the clustered ProLiant servers
- Installation poster for the clustered ProLiant servers
- User guides for the arrays
- User guides for the Fibre Channel host bus adapters (HBAs)
- Installation Guide for the network interface card (NIC) of choice
- Compaq SmartStart Installation poster
- Compaq Insight Manager Installation poster
- *Novell Cluster Services Overview and Installation*
- *Guide to Installing and Configuring the ProLiant Cluster for NetWare*
- *Certification Matrix for the Compaq ProLiant Cluster for NetWare*

INTERNET These documents are available from the Compaq website at:
<http://www.compaq.com>

Other Related Courses

The following courses provide supplemental information on topics covered within this course:

- Compaq High-Availability Full-Line Technical WBT
- Compaq StorageWorks Full-Line Technical Training
- Designing and Implementing Compaq StorageWorks Solutions on Windows NT and NetWare Platforms
- Designing and Implementing the StorageWorks Enterprise Backup Solutions
- Designing and Implementing the StorageWorks Enterprise Backup Solutions WBT
- Implementing Compaq SAN Solutions

Obtaining Help

If encountering problems at any time while following the procedures in this document, assistance is available from the following.

Compaq Technical Support

Technical support service is available 24 hours a day, 7 days a week. Calls might be monitored for quality assurance.

INTERNET

Telephone numbers for worldwide Technical Support Centers are listed on the Compaq website at: **<http://www.compaq.com>**

Compaq Authorized Reseller

For the name and location of the nearest Compaq authorized reseller, refer to the Compaq website.

INTERNET

The Compaq website has information on products and related high-availability topics at:
<http://www.compaq.com/solutions/enterprise/highavailability.html>

Web Resources

In addition to hardware and software products, Compaq also provides information to help stay current on the latest developments and assistance in making deployment decisions. These information products range from those with no specific operating system focus to those that address specific operating system issues and answers.

The following table lists Compaq resources on the Web.

Item	Web Location
Compaq ActiveAnswers — Gives the benefit of Compaq experience to help manage the system and reduce the time, risks, and complexity associated with deploying solutions.	http://www.compaq.com/activeanswers
Compaq ActiveUpdate — Offers proactive notification and delivery of the latest software updates eliminating wasted time spent searching the web. Compaq ActiveUpdate subscribers receive automatic delivery of software updates for Compaq servers, desktops, workstations, and portables.	http://www.compaq.com/products/servers/management/activeupdate/index.html
Compaq Insight Manager 7 — Leverages the power of the Internet to provide web-based systems management for Compaq servers, and any HTTP, SNMP MIB2, or DMI 2.0-compliant device.	http://www.compaq.com/manage
Compaq SmartStart for Servers — Provides everything needed to get servers up and running with full Compaq support.	http://www.compaq.com/products/servers/SmartStart/index.html
Customer Advisories — Describes any known problems and workarounds related to a Service Pack release.	http://www.compaq.com/support/techpubs/Customer_advisories/index.html
Press Releases and Communiqués — Announces the availability of new products and versions.	http://www.compaq.com/newsroom/pr
Novell Partnership site — Contains partnership information, white papers, event listings, and so forth.	http://www.compaq.com/novell
Server Software Download Center — Contains a complete listing of software available for download. This database is searchable by product and operating system or by category and operating system.	http://www.compaq.com/support/files/server/us/index.html
White Papers (complete listing) — Explains ways to optimize the environment and obtain the maximum benefit from software enhancements.	http://www3.compaq.com/support/reference_library/selectproduct.asp