

StorageWorks by Compaq

NAS B2000

Administration Guide

Part Number 292278-001

(First Edition) May 2002

Product Version: Version 1.0

This guide provides information on performing the administrative tasks necessary to manage the *StorageWorks*[™] NAS B2000 by Compaq. Overview information as well as procedural instructions are included in this guide.

COMPAQ

© 2002 Compaq Information Technologies Group, L.P.

Compaq, the Compaq logo, StorageWorks, Tru64, QuickFind, ActiveAnswer and ROMPaq are trademarks of Compaq Information Technologies Group, L.P. in the U.S. and/or other countries.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the U.S. and/or other countries.

UNIX is a trademark of The Open Group in the U.S. and/or other countries.

All other product names mentioned herein may be trademarks of their respective companies.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty..

StorageWorks by Compaq NAS B2000 Administration Guide

(First Edition) May 2002
Part Number 292278-001

Contents

About This Guide

Symbols in Text.....	xv
Important Safety Information	xv
Compaq Technician Notes.....	xvi
Where to Go for Additional Help	xvii
Integrated Management Log	xvii
Telephone Numbers	xvii

Chapter 1

System Overview

Product Definition and Information.....	1-2
Server Hardware Features	1-2
Software Features	1-3
Product Information	1-4
Deployment Scenarios	1-6
Environment Scenarios	1-8
User Interfaces	1-9
NAS B2000 Web-Based User Interface.....	1-10
NAS B2000 Microsoft Management Console.....	1-14

Chapter 2

Setup Completion and Basic Administrative Procedures

Setup Completion	2-2
Setting up Ethernet NIC Teams (Optional).....	2-3
Creating and Managing Users and Groups.....	2-18
Creating and Managing File Shares	2-19
Configuring Data Replication Software	2-20

Basic Administrative Procedures	2-21
Setting the System Date and Time.....	2-22
Powering Down and Restarting the Server.....	2-23
Viewing and Maintaining Audit Logs	2-24
Using Terminal Services.....	2-25
Setting up Email Alerts.....	2-26
Changing System Network Settings	2-27

Chapter 3

Storage Management Overview

Storage Management Process	3-2
Physical Storage Overview	3-4
Physical Hard Drives	3-4
Arrays	3-5
Logical Drives (LUNs)	3-6
Fault-Tolerance Methods.....	3-8
Physical Storage Best Practices	3-15
Virtual Storage Overview	3-16
Pools	3-16
Virtual Disks.....	3-17
Snapshots.....	3-18
VR Lifeguard Service.....	3-22
Virtual Storage Best Practices	3-24

Chapter 4

Storage Management Planning

Fundamental Storage Configuration Planning Issues	4-2
System Priorities.....	4-3
Array Configuration (Striping) Methods	4-4
Recommended System Configurations.....	4-12
Physical Storage Planning Issues	4-17
Hard Drive Sizes and Types	4-17
Use and Number of Spare Disks.....	4-20
LUN Sizing.....	4-21
Virtual Storage Planning Issues	4-23
Pool Considerations.....	4-23
Virtual Disk Considerations	4-24
Share Considerations	4-25

Storage Sizing Considerations	4-25
Storage Management Planning Scenarios.....	4-29
A Complete and Detailed Storage Planning Example.....	4-29
A Simple Sizing Comparison.....	4-40
An Example of a Storage Subsystem Using Different Array Configurations.....	4-41
Planning Worksheet	4-43
Migration Issues.....	4-47
Developing a Migration Plan	4-47
Performing the Migration.....	4-49
Storage Capacity Expansion Issues	4-51

Chapter 5

Physical Storage Management

Hard Drive Management	5-2
Defining Hard Drive LED Indicators.....	5-3
Replacing Failed Hard Drives	5-5
Moving Hard Drives	5-9
Moving Arrays	5-10
Array and LUN Management	5-11
Compaq ACU Overview	5-12
Accessing the ACU	5-13
Entering Controller Settings.....	5-16
Creating a New Array	5-19
Creating Logical Drives (LUNs).....	5-23
Expanding the Capacity of an Existing Array.....	5-26
Migrating an Existing LUN to a New RAID Level or Stripe Size.....	5-29

Chapter 6

Virtual Storage Management

Storage Management Wizard.....	6-3
Wizard Tasks.....	6-4
Virtual Disk Wizard Tasks.....	6-5
Snapshot Wizard Tasks	6-6
Pool Management (Details)	6-7
Creating a New Pool	6-9
Deleting a Pool.....	6-10
Viewing Pool Properties.....	6-10
Setting Pool Policies for a Specific Pool.....	6-11

Adding Storage Units to a Pool	6-14
Bringing a Pool Online and Offline.....	6-14
Virtual Disk Management (Details)	6-15
Creating a New Virtual Disk	6-17
Setting the Drive Letter of a Virtual Disk	6-19
Formatting a Virtual Disk.....	6-20
Deleting a Virtual Disk.....	6-21
Viewing Virtual Disk Properties	6-22
Growing a Virtual Disk	6-23
Backing Up (Creating a Scheduled Snapshot of) a Virtual Disk.....	6-25
Snapshot Management (Details)	6-29
Creating a New Snapshot.....	6-30
Deleting a Snapshot.....	6-33
Viewing Snapshot Properties.....	6-33
Setting the Drive Letter for a Snapshot	6-34
Creating a Snapshot Schedule	6-35
Displaying and Deleting a Snapshot Schedule	6-39
Scheduling Snapshot Deletions	6-41
Restoring a Virtual Disk from a Snapshot.....	6-43
Enabling Incremental Backup Support.....	6-47
Global Pool Policy Settings	6-49
Namespace Recovery	6-52
Drive Quotas	6-53
Managing Quotas Using the Quota Management Wizard.....	6-53
Managing Quotas (Details).....	6-54
Enabling and Disabling Quota Management on a Virtual Disk	6-55
Creating New Quota Entries for a User or Group	6-56
Deleting Quota Entries for a User or Group.....	6-58
Modifying Quota Entries for a User or Group.....	6-58

Chapter 7

User and Group Management

Domain Compared to Workgroup Environments	7-2
User and Group Name Planning.....	7-3
Managing User Names	7-4
Managing Group Names.....	7-4
Workgroup User and Group Management	7-5
Managing Users and Groups Using the Wizard	7-6
Managing Local Users (Details).....	7-10

Managing Local Groups (Details)..... 7-15

Chapter 8

Folder and Share Management

Folder Management	8-3
Navigating to a Specific Volume or Folder.....	8-4
Creating a New Folder	8-6
Deleting a Folder.....	8-7
Modifying Folder Properties	8-8
Creating a New Share for a Volume or Folder.....	8-9
Managing Shares for a Volume or Folder.....	8-11
Managing File-Level Permissions.....	8-12
Share Management	8-19
Defining Access Control Lists	8-20
Integrating Local File System Security into Windows Domain Environments	8-20
Comparing Administrative (Hidden) and Standard Shares.....	8-21
Planning for Compatibility between File-Sharing Protocols	8-21
Managing Shares Using the Shares Management Wizard	8-22
Managing Shares (Details).....	8-25
Protocol Parameter Settings.....	8-34

Chapter 9

UNIX File System Management

Network File System	9-3
Server for NFS	9-4
Authenticating User Access	9-4
Indicating the Computer to Use for the NFS User Mapping Server	9-5
Logging Events	9-7
Installing NFS Authentication Software on the Domain Controllers.....	9-8
NFS File Shares	9-9
NFS Protocol Properties Settings	9-15
NFS Client Groups	9-20
Adding a New Client Group.....	9-21
Deleting a client group.....	9-23
Editing Client Group Information	9-24
NFS User and Group Mappings.....	9-25
Types of Mappings.....	9-26
User Name Mapping Best Practices.....	9-28

Creating and Managing User and Group Mappings	9-29
Backing up and Restoring Mappings.....	9-37
NFS File Sharing Tests	9-39
Terminal Services, Telnet Service, and Remote Shell Service	9-40
Using Terminal Services.....	9-40
Using Telnet Service	9-41
Using Remote Shell Service	9-41
Password Synchronization	9-42
Password Synchronization Best Practices	9-43
Password Synchronization Requirements.....	9-44
Implementing Password Synchronization	9-44
Configuring Advanced Settings.....	9-45
Installing Password Synchronization.....	9-46
Customizing Password Synchronization	9-48

Chapter 10

NetWare File System Management

Installing Services for NetWare	10-2
Managing File and Print Services for NetWare	10-5
Creating and Managing NetWare Users.....	10-6
Adding Local NetWare Users.....	10-6
Enabling Local NetWare User Accounts.....	10-8
Managing NCP Volumes (Shares)	10-9
Creating and Managing NCP File Shares Using the WebUI.....	10-9
Creating and Managing NCP Shares using the MMC.....	10-13

Chapter 11

Remote Access Methods and Monitoring

Web-Based User Interface	11-2
Terminal Services.....	11-3
Remote Insight Lights-Out Edition Board	11-4
Features.....	11-4
Remote Insight Lights-Out Edition Board Configuration	11-6
Using the Remote Insight Lights-Out Edition Board to Access the NAS B2000 ..	11-7
Telnet Server	11-8
Enabling Telnet Server	11-8
Configuring Telnet Server	11-8
Remote Shell Daemon.....	11-10

Compaq Insight Manager	11-10
Compaq Insight Manager Console	11-11
Compaq Insight Manager Agent Web Interface.....	11-13
Enterprise Management Applications.....	11-14
HP OpenView (Windows-Based Operating System).....	11-14
Tivoli NetView (AIX).....	11-16
Installing the Management Software on the Client Machine	11-17

Appendix A

Backup Utility Management

Backup Solutions	A-2
System Environments.....	A-2
Hardware Options	A-2
Software Options.....	A-3
Best Practices.....	A-4
Regular and Reliable Backups	A-4
Automated Tape Libraries.....	A-5
Multiple Backup Devices	A-5
Backup Schedules	A-6
Media Rotation.....	A-7
Offsite Storage	A-7
Server Setup Information Archival	A-8
Snapshots and Quick Online Restores	A-8
Readiness Testing.....	A-9
Disaster Recovery	A-10

Index

List of Figures

1-1	Primary WebUI screen	1-10
1-2	Microsoft Management Console	1-14
2-1	Installing CPQTeam	2-4
2-2	CPQTeam installation complete.....	2-5
2-3	CPQTeam utility icon	2-5
2-4	CPQTeam Properties dialog box.....	2-6
2-5	NIC Properties, Teaming Controls tab, Fault Tolerant option	2-7
2-6	NIC Properties, Teaming Controls tab, Load Balancing option	2-9
2-7	CPQTeam dialog box	2-11
2-8	NIC Team Properties dialog box.....	2-13

2-9	NIC Team TCP/IP Properties dialog box	2-14
2-10	Updated CPQTeam Properties dialog box.....	2-15
2-11	Maintenance menu.....	2-21
2-12	Date and Time dialog box.....	2-22
2-13	Shutdown menu	2-23
2-14	Logs menu	2-24
2-15	Terminal Services session.....	2-25
2-16	Network menu	2-27
3-1	Storage Management process	3-3
3-2	Separate physical drive (P1, P2, P3) read/write (R/W) operations.....	3-5
3-3	Configuring the physical drives into an array (A1) dramatically improves read/write efficiency	3-5
3-4	RAID 0 (data striping) (S1-S4) of data blocks (B1-B12).....	3-6
3-5	Two arrays (A1, A2) and five logical drives (L1 through L5) spread over five physical drives	3-7
3-6	RAID 1 (drive mirroring) of P1 onto P2	3-10
3-7	RAID 5 (distributed data guarding) showing parity information (P)	3-11
3-8	RAID ADG (advanced data guarding) with two sets of parity data.....	3-13
4-1	System characteristics.....	4-3
4-2	Vertical Array configurations	4-6
4-3	Horizontal Array configurations.....	4-8
4-4	NSPOF Horizontal Array configuration.....	4-11
4-5	Recommended Configuration methods	4-14
4-6	How drive characteristics affect array performance	4-18
5-1	Hot-plug hard drive LED indicators	5-3
5-2	ACU Logical Drive view.....	5-14
5-3	ACU Physical view	5-15
5-4	Controller Settings dialog box	5-17
5-5	ACU main configuration screen	5-19
5-6	Create Drive Array screen	5-20
5-7	Example Array B	5-21
5-8	Example Array Logical Configuration view with two arrays.....	5-22
5-9	Create Logical Drive dialog box.....	5-23
5-10	Example array - Configuration View screen with two arrays	5-25
5-11	Array expansion example – Logical Configuration View screen.....	5-27
5-12	Expansion wizards - Logical Drive screen	5-28
5-13	Migrate RAID/Stripe Size screen	5-30
6-1	Storage Management wizard	6-3
6-2	Manage Pools dialog box	6-7

6-3	Create New Pool dialog box.....	6-9
6-4	Pool Properties summary display.....	6-10
6-5	Manage Policies dialog box.....	6-12
6-6	Manage Virtual Disks dialog box.....	6-15
6-7	Create a New Virtual Disk dialog box.....	6-17
6-8	New Virtual Disk Step 1 dialog box.....	6-18
6-9	Format Virtual Disk dialog box.....	6-20
6-10	Virtual Disk Properties screen.....	6-22
6-11	Schedule a Snapshot screen.....	6-25
6-12	Schedule Information screen.....	6-27
6-13	Virtual Disk Backup Summary screen.....	6-28
6-14	Manage Snapshots dialog box.....	6-29
6-15	Virtual Disks and Snapshots screen.....	6-31
6-16	New Snapshot Information screen.....	6-32
6-17	Snapshot Properties summary display.....	6-33
6-18	Set Drive Letter dialog box.....	6-34
6-19	Snapshot Wizard Welcome screen.....	6-35
6-20	Snapshot and Task Information screen.....	6-36
6-21	Schedule Information screen.....	6-37
6-22	Display/Delete Snapshot Schedule dialog box.....	6-39
6-23	Display/Delete the Snapshot Schedule wizard screen.....	6-40
6-24	Schedule Snapshot Deletion Task dialog box.....	6-41
6-25	User Information and Schedule Information screen.....	6-42
6-26	Restore Virtual Disk from Snapshot screen.....	6-45
6-27	Restore Virtual Disk screen.....	6-46
6-28	Backup Operation Information screen.....	6-47
6-29	Log File Information Screen.....	6-48
6-30	Manage Policies dialog box.....	6-50
6-31	Namespace Recovery screen.....	6-52
6-32	Disk Quota dialog box.....	6-54
6-33	Default Quota Dialog box.....	6-55
6-34	Quota Entries dialog box.....	6-56
6-35	New Quota Entry dialog box.....	6-57
6-36	Quota Entry dialog box for a user.....	6-59
7-1	User Management wizard.....	7-7
7-2	Local Users dialog box.....	7-10
7-3	Create New User dialog box.....	7-12
7-4	User Properties dialog box.....	7-14
7-5	Local Groups dialog box.....	7-15

7-6	Create New Group dialog box, General tab.....	7-16
7-7	Group Properties dialog box, General tab	7-17
7-8	Group Properties dialog box, Members tab	7-19
8-1	Volumes dialog box.....	8-4
8-2	Folders dialog box	8-5
8-3	Create a New Folder dialog box, General tab.....	8-6
8-4	Folder Properties dialog box, General tab	8-8
8-5	Create New Share dialog box, General tab.....	8-10
8-6	Security Properties dialog box for folder name NTSF Test	8-13
8-7	Access Control Settings dialog box for folder name NTSF Test, Permissions tab	8-14
8-8	User or Group Permission Entry dialog box for folder name NTSF Test	8-15
8-9	Access Control Settings, Auditing tab dialog box for folder name NTSF Test	8-16
8-10	Select User, Computer, or Group dialog box.....	8-17
8-11	Auditing Entry dialog box for folder name NTSF Test.....	8-17
8-12	Access Control Settings, Owner tab dialog box for folder name NTSF Test.....	8-18
8-13	Share Management wizard	8-23
8-14	Create a New Share dialog box, General tab	8-26
8-15	Share Properties dialog box, General tab	8-28
8-16	Share Properties dialog box, CIFS Sharing tab	8-30
8-17	Share Properties dialog box, NFS Sharing tab	8-31
8-18	Share Properties dialog box, NetWare Sharing tab	8-33
8-19	Sharing Protocols dialog box.....	8-35
9-1	MMC Server for NFS screen, User Mapping tab	9-6
9-2	MMC Server for NFS screen, Logging tab	9-7
9-3	Create a New Share dialog box, General tab	9-10
9-4	Share Properties dialog box, General tab	9-12
9-5	NFS Sharing tab.....	9-13
9-6	NFS Sharing Protocols menu.....	9-15
9-7	NFS Async/Sync Settings dialog box.....	9-17
9-8	NFS Locks dialog box	9-19
9-9	NFS Client Groups dialog box	9-21
9-10	New NFS Client Group dialog box	9-22
9-11	Client Groups dialog box.....	9-23
9-12	Edit NFS Client Groups dialog box.....	9-24
9-13	Mapping Server “ls -al” Command example	9-27
9-14	User and Group Mappings dialog box, General tab	9-30
9-15	User and Group Mappings dialog box, Simple Mapping tab	9-32
9-16	User and Group Mappings dialog box, Explicit User Mapping tab	9-33
9-17	User and Group Mappings dialog box, Explicit Group Mapping tab.....	9-35

9-18	MMC User Name Mapping screen, Map Maintenance tab.....	9-37
9-19	MMC Password Synchronization screen	9-42
9-20	MMC Password Synchronization screen, Advanced Settings dialog box	9-45
10-1	Local Area Connection Properties page, Install option.....	10-3
10-2	Installing File and Print Services for NetWare	10-4
10-3	File and Print Services for NetWare screen	10-5
10-4	MMC New User dialog box	10-7
10-5	NetWare Services tab.....	10-8
10-6	Create a New Share dialog box, General tab.....	10-10
10-7	Share Properties dialog box, General tab	10-11
10-8	Share Properties dialog box, NetWare Sharing tab	10-12
10-9	Create Shared Folder dialog box	10-13
10-10	NetWare Basic Share Permissions dialog box	10-14
10-11	Customize Permissions dialog box, Share Permissions tab	10-15
10-12	Customize Permissions dialog box, Security tab	10-16
11-1	Telnet Server interface screen.....	11-9
11-2	Device Information window.....	11-12
11-3	Compaq Insight Manager Agent Web interface.....	11-13
11-4	Web Enabled interface	11-15

List of Tables

3-1	Summary of RAID methods.....	3-14
4-1	Vertical Carving Disk Use per RAID Level	4-7
4-2	Horizontal Configuration Disk Use per RAID Level.....	4-9
4-3	Suggested Storage Enclosure Configurations	4-13
4-4	Table 4-6 Legend	4-19
4-5	Example Storage Need Worksheet.....	4-32
4-6	Example Array Configuration Requirements Worksheet	4-36
4-7	Example Drives Required Worksheet	4-38
4-8	Example Enclosures Required Worksheet	4-39
4-9	Example Usable Space Using Different Configurations	4-42
4-10	Usable Storage Need Worksheet.....	4-44
4-11	Array Configuration Storage Needs Worksheet.....	4-45
4-12	Drive and Enclosure Requirements.....	4-46
5-1	Hard Drive LED Combinations.....	5-4
5-2	Storage Enclosure Drive Bay Configuration.....	5-7
5-3	Optimum Stripe Sizes for Different Environments	5-24
6-1	Pool Policy Default Settings	6-13

6-2	Pool Policy Default Settings (duplicated table).....	6-51
7-1	Group Name Examples.....	7-5
9-1	Command Line Interface Command Prompts	9-41

About This Guide

This guide provides step-by-step instructions for managing the *StorageWorks™* NAS B2000 by Compaq.

Symbols in Text

These symbols may be found in the text of this guide. They have the following meanings.



WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or loss of life.



CAUTION: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or loss of information.

IMPORTANT: Text set off in this manner presents clarifying information or specific instructions.

NOTE: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Important Safety Information

Before installing this product, read the *Important Safety Information* document provided.

Compaq Technician Notes



WARNING: Only authorized technicians trained by Compaq should attempt to repair this equipment. All troubleshooting and repair procedures are detailed to allow only subassembly/module-level repair. Because of the complexity of the individual boards and subassemblies, no one should attempt to make repairs at the component level or to make modifications to any printed wiring board. Improper repairs can create a safety hazard.



WARNING: To reduce the risk of personal injury from electric shock and hazardous energy levels, do not exceed the level of repairs specified in these procedures. Because of the complexity of the individual boards and subassemblies, do not attempt to make repairs at the component level or to make modifications to any printed wiring board. Improper repairs can create conditions that are hazardous.



WARNING: To reduce the risk of electric shock or damage to the equipment:

- Disconnect power from the system by unplugging all power cords from the power supplies.
 - Do not disable the power cord grounding plug. The grounding plug is an important safety feature.
 - Plug the power cord into a grounded (earthed) electrical outlet that is easily accessible at all times.
-



CAUTION: To properly ventilate the system, you must provide at least 7.6 cm (3.0 in.) of clearance at the front and back of the server.



CAUTION: The computer is designed to be electrically grounded (earthed). To ensure proper operation, plug the AC power cord into a properly grounded AC outlet only.

NOTE: Any indications of component replacement or printed wiring board modifications may void any warranty.

Where to Go for Additional Help

In addition to this guide, the following information sources are available:

- User documentation
- *Compaq Service Quick Reference Guide*
- Service training guides
- Compaq service advisories and bulletins
- Compaq *QuickFind*[™] information services
- *Compaq Insight Manager*[™] software

For additional copies, visit the Compaq website:

www.compaq.com

Integrated Management Log

The server includes an integrated, nonvolatile management log that contains fault and management information. The contents of the Integrated Management Log (IML) can be viewed with Compaq Insight Manager.

Telephone Numbers

For the name of your nearest Compaq authorized reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.

For Compaq technical support:

- In the United States and Canada, call 1-800-OK COMPAQ.
- For Compaq technical support phone numbers outside the United States and Canada, visit the Compaq website:

www.compaq.com

System Overview

The *StorageWorks*[™] NAS B2000 by Compaq can be used in many types of computing environments, from basic Microsoft Windows workgroups to complicated multiprotocol domains using CIFS, NFS, NCP, AppleTalk, FTP, and HTTP. The corresponding varieties of clients that can be serviced include any Windows, UNIX, Linux, Novell, or Macintosh variant.

This chapter provides an overview of these environments and deployments and includes brief descriptions of system user interfaces, applications, and options.

- Product Definition and Information
 - Server Hardware Features
 - Software Features
 - Product Information
- Deployment Scenarios
- Environment Scenarios
- User Interfaces
 - NAS B2000 Web-Based User Interface
 - NAS B2000 Microsoft Management Console

Product Definition and Information

The NAS B2000 is a business-class NAS solution that provides reliable performance, manageability and fault tolerance.

Server Hardware Features

The following features are included in the NAS B2000 server:

- Intel Pentium III 1.4-GHz FC-PGA2 processor, with 512-KB L2 cache
- 512-MB PC 133MHz registered ECC SDRAM memory
- 64-bit I/O technology (66/33 MHz)
- Two 18.2-GB, 10,000-rpm hot-pluggable hard drives for the NAS operating system
- Three 72.8-GB hot-pluggable hard drives for data
- Remote Insight Lights-Out Edition board
- Two embedded 10/100 network interface controllers (NICs)
- IDE CD-ROM drive
- Redundant hot-plug power supplies and fans
- Two open PCI hot-plug controller slots

Optional Features

The following features are optional for the NAS B2000 server:

- Additional memory
- Smart Array 5i controller
- Smart Array 5300 controller
- StorageWorks 4300 Family storage enclosures
- Network interface cards (NICs)

NOTE: Refer to the *StorageWorks by Compaq NAS B2000 Quick Reference Guide* for a complete list of options.

Software Features

Advanced features included and supported by the NAS B2000 include:

- Array Configuration Utility (ACU)
- Compaq Insight Manager performance monitoring
- Microsoft Services for Macintosh
- Microsoft Services for Netware
- Microsoft Services for UNIX (SFU)
- NAS Web-Based User Interface (WebUI)
- RAID 0, 1, 1+0, 5
- StorageWorks Data Copy
- Virtual Replicator (VR)
- Windows-powered OS plus Service Pack 2
- Third-party supported software, including:
 - Backup software
 - Management software
 - Quota management
 - Virus protection

For specific software product recommendations, go to the Compaq website:

www.compaq.com

Product Information

The NAS B2000 provides significant performance gains over general-purpose servers by integrating optimized hardware components and specialized software. Integrating NAS devices into the network improves the performance of existing servers because NAS devices are optimized for file-serving tasks.

Product Manageability

The NAS B2000 ships with the following utilities and features that ease the administration tasks associated with managing the system:

- The Rapid Startup Utility is a user-friendly configuration utility that ensures easy configuration.
- The WebUI is a simple, graphical user interface (GUI) that helps with administration tasks.
- The Remote Insight Lights-Out Edition board also provides remote access, sends alerts, and performs other management functions, even if the host server operating system is not responding or the server has lost power.
- Compaq Insight Manager is a comprehensive tool designed to be a key component in the systems management environment. It monitors the operations of Compaq servers, workstations, and clients. Compaq Insight Manager provides system administrators more control through visual interface, comprehensive fault and configuration management, and industry-leading remote management.

Product Redundancy

The NAS B2000 is specifically designed to perform file-serving tasks for networks. Using industry standard components, redundancy of power supplies, NICs, and fans ensures reliability.

Other industry-standard features, such as redundant array of independent drives (RAID) and remote manageability, further enhance the overall dependability of the NAS B2000.

The server contains dual 18.2-GB hard drives preconfigured with the NAS operating system so that the active system volume is mirrored (RAID 1) to the second drive. If one of the internal drives fails, the integrity of the system is preserved, because the system will use the copy of the operating system on the remaining healthy drive. The drives in the server are hot pluggable, so the failed drive can be replaced while the system is running. When the failed drive is replaced, the system automatically uses the version of the operating system on the healthy drive to rebuild the replacement.

The NAS B2000 also contains three 72.8-GB hard drives for data storage. These hard drives are not configured, allowing maximum configuration options. These drives may be configured to RAID levels 0+1 and 5.

NOTE: RAID 1 requires an even number of drives.

Power supplies can be replaced while the server is running. To ensure redundancy, it is important to connect each power supply to a separate power source. If one power source fails, the server remains operational through the second power source.

Through a seamless, hardware-based, graphical remote console, the Remote Insight Lights-Out Edition board provides the administrator with full control of the server from a remote location. Using a client browser, the administrator can remotely power up, power down, and operate the console. A built-in processor, combined with an external power supply, makes the board independent of the server and the operating system.

Product Scalability

The NAS B2000 offers optimized performance for a growing environment. Storage capacity can increase as a business grows without downtime or compromised performance. Internally the NAS B2000 can grow up to four data drives. With four 72.8 disk drives storage capacity, it can grow up to 291.2 GB of raw storage. Externally the NAS B2000 can attach with the embedded SCSI port to nine StorageWorks 4300 Family storage enclosures with the use of two Smart Array 5300 controllers. This increases the maximum storage capacity to nine terabytes with one hundred twenty six 72.8-GB drives.

Deployment Scenarios

The default shipping configuration contains two 10/100 integrated network interface controller (NIC) ports for client data access. These data ports also allow access to the Web user interface (WebUI) that accompanies the product. It is from the WebUI that most management and administrative procedures can be accomplished. An additional management port for remote console and diagnostics is provided off the Remote Insight Lights-Out Edition (RILOE) board. Compaq recommends that this connection be placed on a management LAN separate from the corporate infrastructure.

The NAS B2000 supports the use of NIC teaming. NIC teaming provides failover and load balancing of network ports of the NAS B2000. NIC teaming requires the network cables to be installed on the same subnet to enable it to work. However, it is not recommended to assign IP addresses to the ports that will be teamed or load balanced prior to the installation and setup of NIC teaming. For this reason, Compaq recommends that you set all network ports to DHCP.

Typical deployment scenarios include:

- **File server consolidation**

As businesses continue to expand their information technology (IT) infrastructures, they must find ways to manage larger environments without a corresponding increase in IT staff. Consolidating many servers into a single NAS device decreases the number of points of administration and increases the availability and flexibility of storage space.

- **Multiprotocol environments**

Some businesses require several types of computing systems to accomplish various tasks. The multiprotocol support of the NAS B2000 allows it to support many types of client computers concurrently.

- **Protocol and platform transitions**

When a transition between platforms is being contemplated or planned, the ability of the NAS B2000 to support most popular file-sharing protocols allows companies to continue to invest in file storage space without concerns about obsolescence. For example, an administrator planning a future transition from Windows to Linux can deploy the NAS B2000 with the confidence that it can support both CIFS and NFS simultaneously, thereby assuring not only a smooth transition, but also a firm protection of their investment.

- **Remote office deployment**

Frequently, branch offices and other remote locations lack dedicated IT staff members. An administrator located in a central location can use the WebUI of the NAS B2000, Microsoft Terminal Services, and other remote administration methods to configure and administer all aspects of the NAS B2000.

Environment Scenarios

The NAS B2000 is deployed into one of two modes:

- Workgroup
- Domain (Windows NT Domain or Active Directory Domain)

The NAS B2000 uses standard Windows user and group administration methods in each of these environments. For procedural instructions on managing users and groups, see Chapter 1 later in this guide.

Regardless of the deployment, the NAS B2000 integrates easily into multiprotocol environments, supporting a wide variety of clients. The following protocols are supported:

- Common Internet File System (CIFS)
- Network File System (NFS)
- Novell Core Protocol (NCP)
- Hypertext Transfer Protocol (HTTP)
- File Transfer Protocol (FTP)
- AppleTalk for Macintosh (AFP, also called MAC)

Workgroup

In a workgroup environment, users and groups are stored and managed separately, on each member server of the workgroup. Workgroups are typical for very small deployments where little or no computing environment planning is required.

Domain

When operating in a Windows NT or Active Directory domain environment, the NAS B2000 is a member of the domain and the domain controller is the repository of all account information. Client machines are also members of the domain and users log on to the domain through their Windows-based client machines. The domain controller also administers user accounts and appropriate access levels to resources that are a part of the domain. Additional information about planning for domain environments can be found at the Compaq *ActiveAnswers*[™] website in the eBusiness Infrastructure solutions area:

www.compaq.com/ActiveAnswers/

The NAS B2000 obtains user account information from the domain controller when deployed in a domain environment. The NAS B2000 itself cannot act as a domain controller.

User Interfaces

There are several user interfaces that administrators can use to access and manage the NAS B2000. Two of these interfaces are:

- NAS B2000 WebUI
- NAS B2000 Microsoft Management Console (MMC)

Each interface contains the same or similar capabilities, but presents them in a different manner. Each of these interfaces are illustrated in the following sections.

NAS B2000 Web-Based User Interface

The WebUI provides for system administration, including user and group management, share management, and local storage management. In addition to dialog boxes, the WebUI includes wizards to guide the administrator while performing repetitive tasks such as creating shares.

To access the WebUI, launch a Web browser and enter the following in the address field:

`http://<your NAS machine name or IP Address>:3201/`

Extensive online help for the WebUI is available by clicking Help on the primary WebUI screen.

The primary screen of the WebUI is illustrated in Figure 1-1.

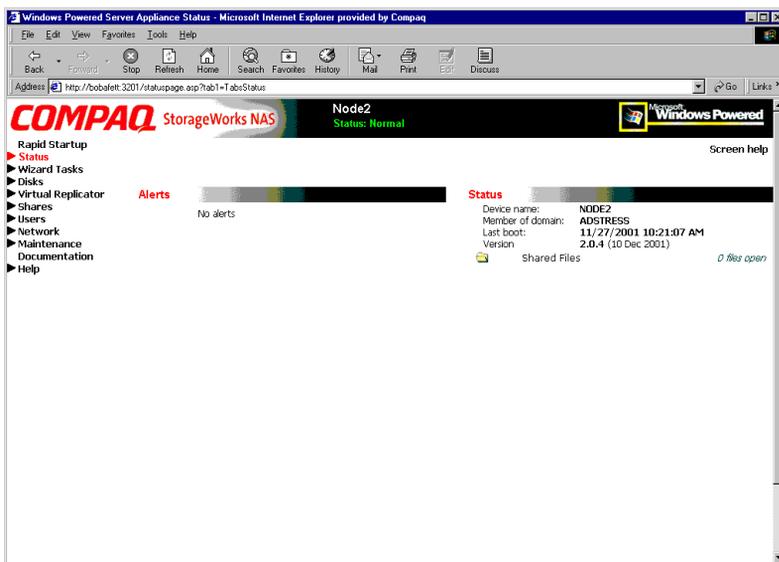


Figure 1-1: Primary WebUI screen

As shown in Figure 1-1, the following areas are administered through this interface:

- Rapid Startup
- Status Information
- Wizard Tasks
- Disks
- Virtual Replicator
- Shares
- Users
- Network
- Maintenance
- Help

Each of these administration areas is briefly discussed in the following paragraphs.

Rapid Startup

Use this utility to enter system setup and configuration information.

Status

The Status option displays system information, including disk status data and system information

Wizard Tasks

Wizards ease administration by guiding the administrator through common, repetitive system tasks. Primary administration tasks of the NAS B2000 include setting up and managing shares, users, and groups. A wizard is provided for each of these management tasks.

Disks

Use this option to manage disks, volumes, and disk quotas. The Compaq ACU is included in this menu option. The ACU is used to create and manage RAID arrays and LUNs.

See Chapter 5 for procedural information regarding managing the arrays and LUNs using ACU.

Virtual Replicator

VR provides advanced, logical-storage management capabilities in Windows-based computing environments. Innovative storage management features simplify storage configuration and management, and enhance availability and scalability.

VR uses pools, virtual disks, and snapshots to manage storage. See Chapter 6 for information on VR.

Shares

The administrator creates folders and shares to control access to files. When a share is created, the administrator indicates the protocols that can be supported by that share as well as the users and groups of users that have access. Protocol parameters are entered in this Shares option. See the Chapter 8 for additional information.

Users

When deployed, the administrator uses this option to manage local users and groups. Local users and groups are discussed in Chapter 7.

Network

The Network option contains system settings, including system identification, global settings, interfaces settings, administration settings, Telnet settings, and SNMP settings.

Maintenance

Some of the maintenance tasks include setting date and time, performing system restarts and shutdowns, viewing audit logs, accessing Terminal Services, and setting up Email alerts.

Help

This option contains help information for the WebUI.

NAS B2000 Microsoft Management Console

The NAS B2000 includes a console screen that shows only those items required for NAS administration.

The functionality of the MMC is similar to the WebUI, but the contents and procedures are arranged and performed differently. The MMC is shown in Figure 1-2.

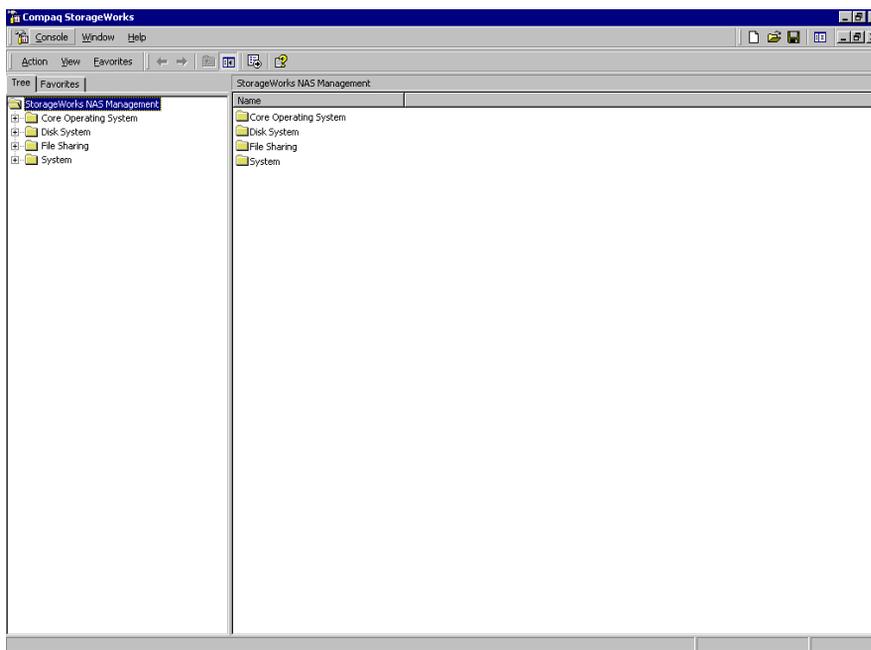


Figure 1-2: Microsoft Management Console

To access the MMC from the WebUI, select **Maintenance, Terminal Services**. A session in Terminal Services is opened, and the NAS device desktop is displayed. Click the icon for the MMC.

From within the MMC, the following folders are available:

- Core Operating System
- Disk System
- File Sharing
- System

Core Operating System

Use this folder to manage local users and groups, access performance logs and alerts, and manage the Event Viewer.

Disk System

Storage management functions are contained in this folder. These functions are managed by the Virtual Replicator. Virtual Replicator is used to configure storage pools, virtual disks, and snapshots.

File Sharing

The File Sharing folder contains modules for the configuration of file sharing exports. CIFS and UNIX file shares are managed through this folder.

System

The System folder contains system summary information.

Setup Completion and Basic Administrative Procedures

This chapter continues the process of setting up the system that was started using the *StorageWorks by Compaq NAS B2000 Quick Start Guide* by discussing additional setup procedures and options.

Basic system administration functions are also included in this chapter.

Unless otherwise instructed, all procedures are performed using the NAS Web-Based User Interface (WebUI).

The following topics are included in this chapter:

- Setup completion
 - Setting up Ethernet NIC teams (optional)
 - Managing system storage
 - Creating and managing users and groups
 - Creating and managing file shares
 - Configuring data replication software

- Basic administrative procedures
 - Setting the system date and time
 - Powering down and restarting the server
 - Viewing and maintaining audit logs
 - Using terminal services
 - Setting up email alerts
 - Changing system network settings

Setup Completion

After the NAS device is physically set up and the basic configuration is established, additional setup steps must be completed. Depending on the deployment scenario of the NAS device, these steps may vary.

Additional setup steps may include:

- Setting up Ethernet NIC teams (optional)
- Managing system storage
- Creating and managing users and groups
- Creating and managing file shares
- Configuring data replication software

Each of these setup steps is discussed in the following sections.

Setting up Ethernet NIC Teams (Optional)

The NAS B2000 is equipped with the Compaq Network Teaming and Configuration (CPQTeam) utility. The CPQTeam utility allows administrators to configure and monitor Ethernet network interface controllers (NIC) teams in a Windows-based operating system. These teams provide options for increasing fault tolerance and throughput.

Fault tolerance provides automatic redundancy. If the primary NIC fails, the secondary NIC takes over. Load Balancing provides the ability to balance transmissions across NICs.

NOTE: The NAS B2000 does not ship with NIC teaming configured.

IMPORTANT: Installing NIC teaming requires a restart of the server.

Procedures include:

- Installing the CPQTeam utility
- Opening the CPQTeam utility
- Adding and configuring NICs in a team
- Configuring the NIC team properties
- Checking the status of the team

Installing the CPQTeam Utility

Before using the CPQTeam utility, it must be installed. To do this:

1. From the WebUI, use Terminal Services to go to the NAS B2000 desktop. Double click the **CPQTeam Setup** icon on the desktop.

If the CPQTeam icon is not displayed, enter the following command:

```
c:\winnt\bin\nicteam\cpqsetup.exe
```

2. When the following message box is displayed, click **Install**.

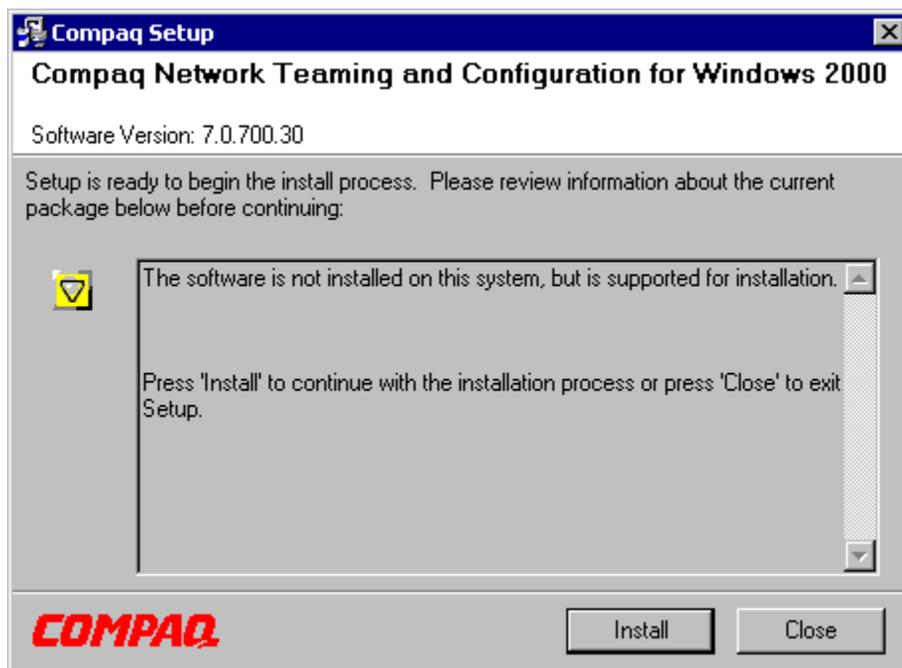


Figure 2-1: Installing CPQTeam

3. When the installation process is complete, the following screen is displayed. Click **Close**.

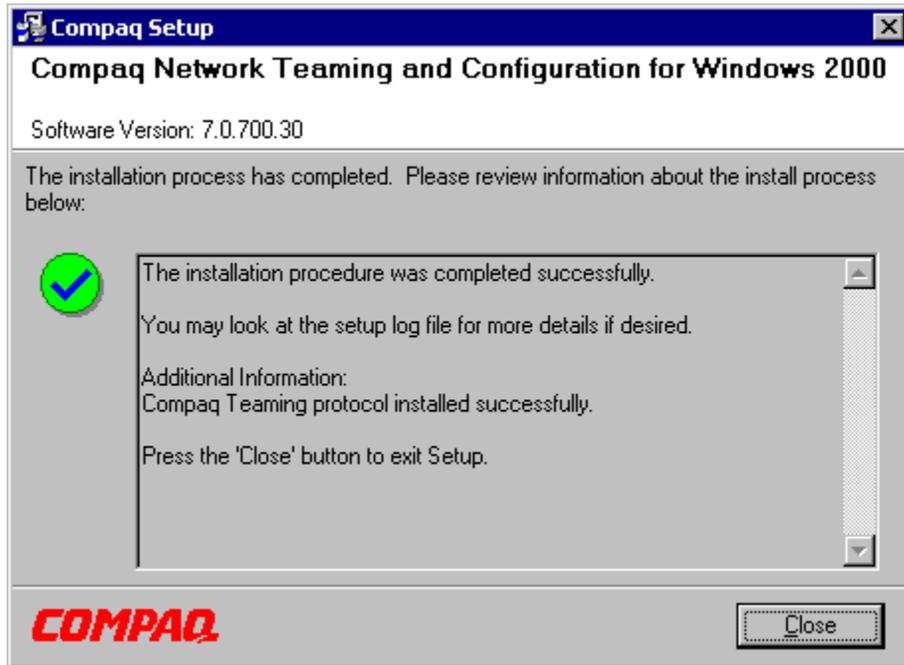


Figure 2-2: CPQTeam installation complete

4. Restart the system.

Opening the CPQTeam Utility

The CPQTeam utility is now accessible from the Windows toolbar at the bottom of the NAS B2000 desktop. To open the CPQTeam utility, click the Tray icon.

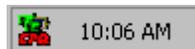


Figure 2-3: CPQTeam utility icon

Adding and Configuring NICs in a Team

Before a NIC is teamed, verify the following:

- The NICs must be on the same network.
- The NICs must be DHCP-enabled and the DNS server address must be left blank.

NOTE: The teaming utility becomes unstable if static IP addresses, subnets, and DNS addresses are set before teaming.

- Duplex and speed settings must be set to use the default values.

To team the NICs:

1. Open the CPQTeam utility. The Network Teaming and Configuration dialog box is displayed. Included in the screen display is the type of NIC and the slot and port used.

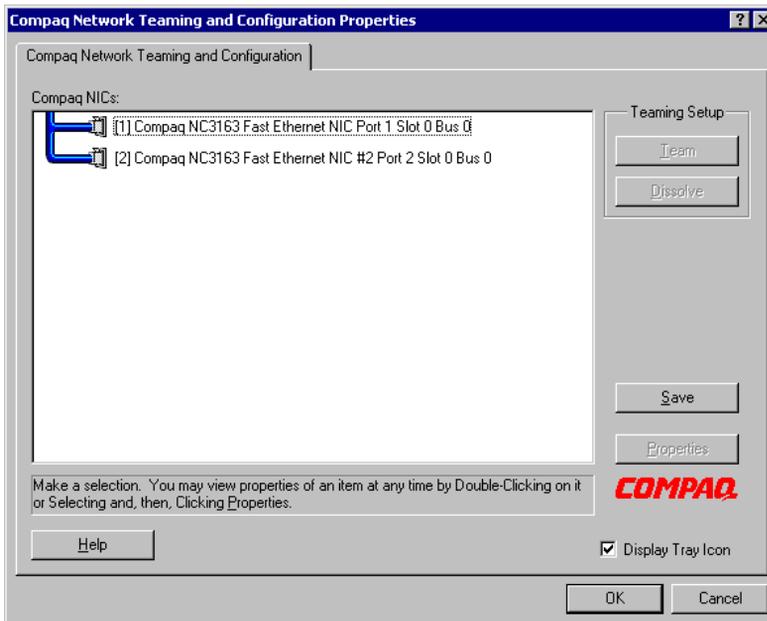


Figure 2-4: CPQTeam Properties dialog box

2. Highlight the NICs to team.

IMPORTANT: Do not select NIC Bus 0.

- Click the **Team** button. The **Teaming Controls** tab of the Properties dialog box is displayed.

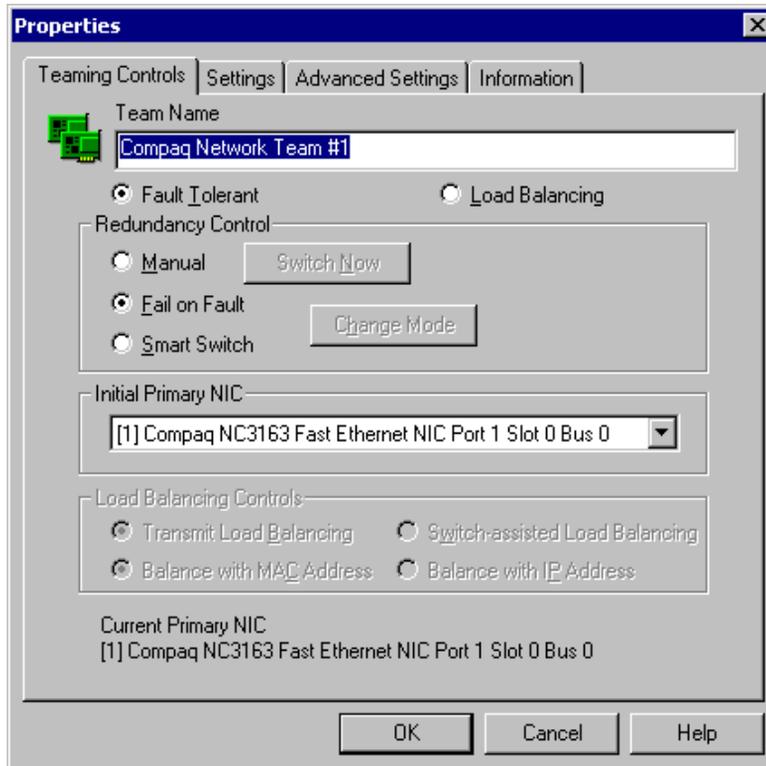


Figure 2-5: NIC Properties, Teaming Controls tab, Fault Tolerant option

- Configure the team by choosing either **Fault Tolerance** or **Load Balancing**. The fault tolerance and load balancing options are discussed in the following sections.

Fault Tolerance

The Fault Tolerance teaming option provides three redundancy control options:

- **Manual**—This setting lets the user change from a Primary NIC to a Secondary NIC only when the user clicks **Switch Now**.

NOTE: The Switch Now option is disabled until the user selects **Manual** and then clicks **Apply**.

- **Fail on Fault**—This setting automatically switches from a primary NIC to a secondary NIC when the primary NIC fails.
- **Smart Switch**—This setting lets a member of a team be selected as the preferred Primary Smart Switch NIC. As long as this NIC is operational, it is always the active NIC. If the NIC fails and it is eventually restored or replaced, it automatically resumes its status as the active NIC.

NOTE: Smart Switch is the recommended choice for fault tolerance.

Detailed information about configuring teams for fault tolerance can be found in the CPQTeam utility help.

Load Balancing

The Load Balancing teaming option provides four load balancing control options:

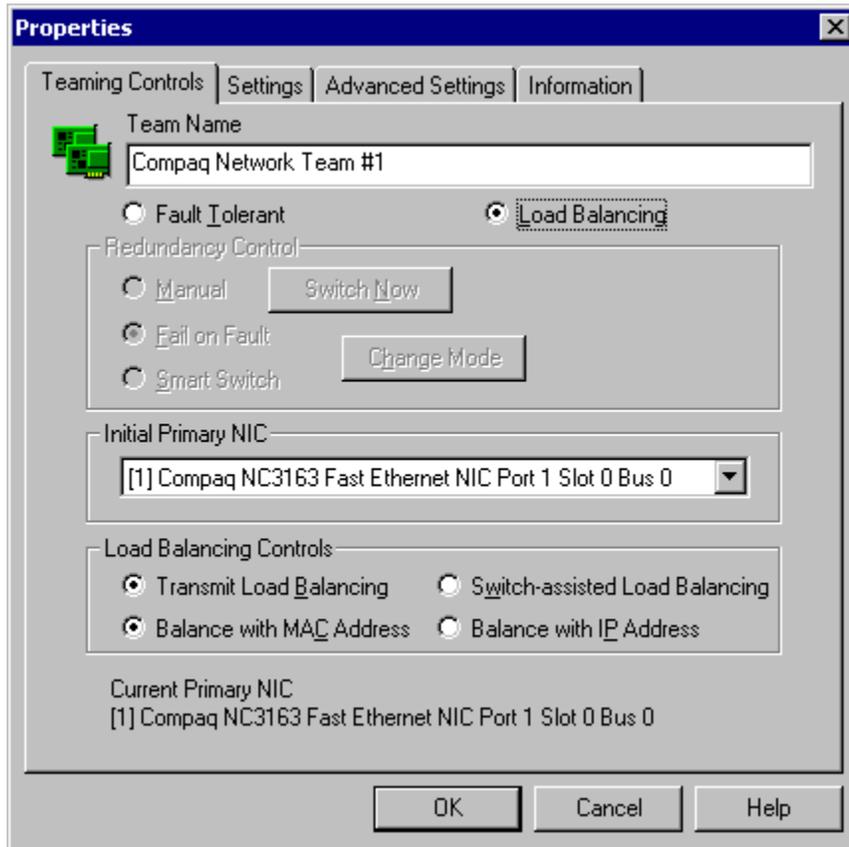


Figure 2-6: NIC Properties, Teaming Controls tab, Load Balancing option

Detailed information about these four load-balancing teaming options can be found in the CPQTeam help.

- **Adaptive Load Balancing (ALB)**—Creates a team of NICs to increase a server transmission throughput and works with any 100Base-TX or Gigabit switch. With ALB, NICs can be grouped into teams to provide a single, virtual NIC with increased transmission bandwidth. To use ALB, link at least two teamed Compaq NICs in the server to the same network switch.
- **Cisco Fast EtherChannel (FEC)**—This feature creates a team of NICs to increase transmission and reception throughput. Unlike ALB, Cisco FEC can be configured to increase both transmission and reception channels between the server and switch. For example, a Cisco FEC team containing four Compaq Fast Ethernet NICs configured for full-duplex operation provides an aggregate maximum transmit rate of 400 megabits per second (Mbits/s) and an aggregate maximum receive rate of 400 Mbits/s, resulting in a total bandwidth of 800 Mbits/s.
- **Balance with MAC Address**—This feature allows load-balancing of IP packets among the teamed NICs using the last four bits of the MAC Address. (See Note.)
- **Balance with IP Address**—This feature allows load-balancing of IP packets among the teamed NICs using the last four bits of the IP Address. (See Note.)

NOTE: The teaming utility can load-balance IP packets among the teamed NICs installed in a server. The primary NIC in the team receives all incoming packets. The choice is available to load-balance with the source MAC address (the address transmitted from the workstation) or the source IP address.

Using the last four bits of either source address, the teaming driver algorithm assigns this source address to the port of one of the NICs in the team. This port is then used to transmit all packets destined for that source address. If there are four NICs in the team, the packets are received by the primary NIC on the team. The packets are retransmitted through one of the four ports.

- Click **Yes** to continue. The following screen is displayed, indicating that there are additional procedures to perform in the NIC teaming process.

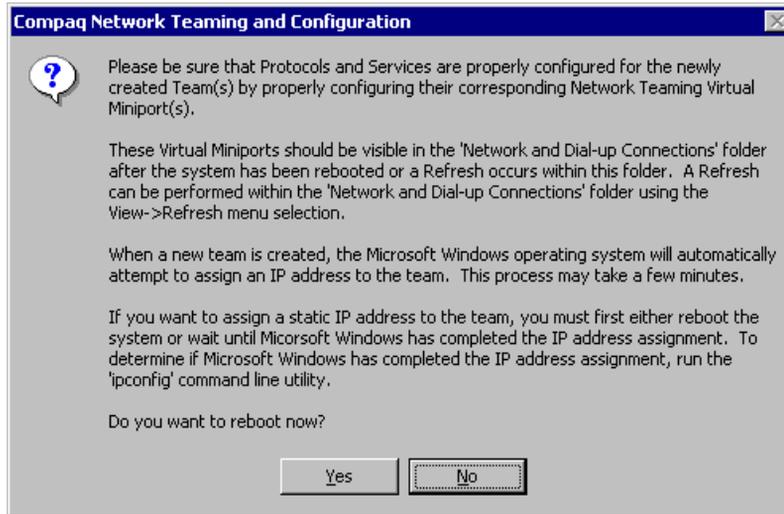


Figure 2-7: CPQTeam dialog box

- Click **OK** to continue.

Configuring the NIC Team Properties

At this point, the NICs are teamed but are not completely configured. Additional procedures include:

- Renaming the teamed connection
- Selecting the option to show an icon on the taskbar
- Configuring TCP/IP on the new team

Renaming the Teamed Connection

The assigned name for the new NIC team connection is “Local Area Connection X,” where *X* represents the next available connection number generated by the system. Compaq recommends changing this name to a more meaningful name, such as “NIC Team.”

To change the name of the connection:

1. From the desktop, right-click the **My Network Places** icon, then click **Properties**. The Network and Dial-up Connections screen is displayed.
2. Move the cursor over each connection icon to view the pop-up box of the icon's name. Locate "Compaq Network Teaming Virtual Miniport."
3. Right-click the connection icon for "Compaq Network Teaming Virtual Miniport," and select **Rename**. Enter a name that is more descriptive than "Local Area Connection X," such as "NIC Team."

Selecting the Option to Show a Connection Icon on the Taskbar

To show a connection icon:

1. In the **Network and Dial-up Connections** screen, double-click the NIC Team connection, and then click **Properties**.
2. At the bottom of the screen, select **Show icon in task bar when connected**, and then click **Close**.

Configuring the TCP/IP Protocol on the New Team

After teaming the NICs, a new virtual network adapter for the team is automatically created. However, the team's IP address, subnet mask, and DNS server address information must be entered.

To enter the TCP/IP address information for the team:

1. From the desktop, go to the **Network and Dial-up Connections** screen and click **Properties**. Right-click the NIC Team icon and then select **Properties**. A screen similar to the following is displayed.

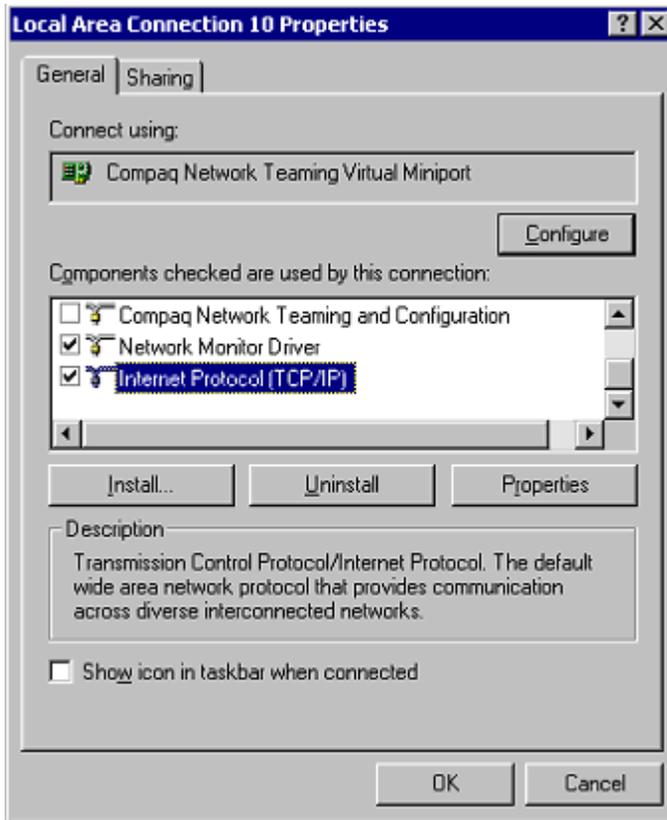


Figure 2-8: NIC Team Properties dialog box

2. Use the arrows and the scroll bar on the right of the screen to scroll through the **Components** list.

3. Click **Internet Protocol (TCP/IP)** and then click **Properties**. The following screen is displayed:

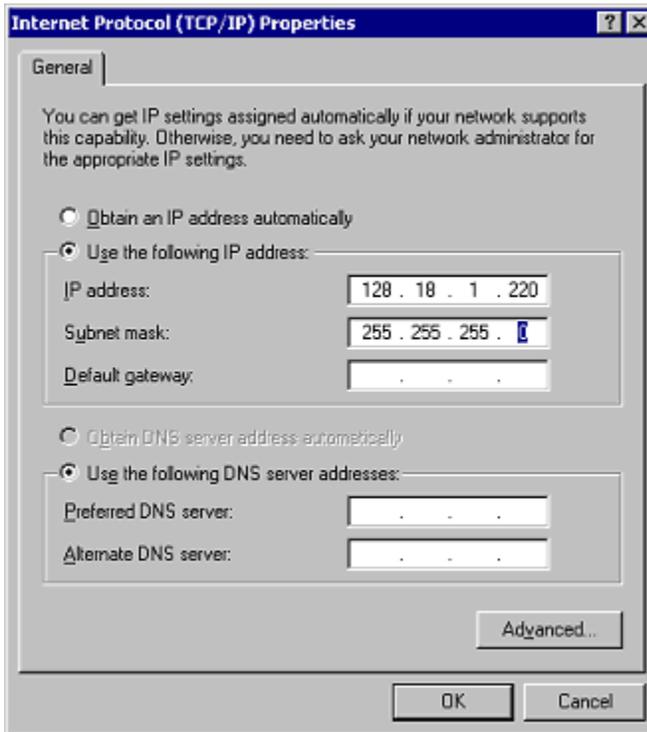


Figure 2-9: NIC Team TCP/IP Properties dialog box

IMPORTANT: If a NIC is teamed, do not modify the TCP/IP settings for the individual NIC ports.

4. *For teamed NICs*, select **Obtain an IP address automatically**.
For un-teamed NICs, select **Use the following IP address**, and enter the IP address and subnet mask. If desired, enter the default gateway.
5. Click **OK**. The Ethernet Team should be working.

Checking the Status of the Team

To check the status of the Ethernet Team, open the CPQTeam utility. The Configuration Properties screen is displayed, showing the teamed NICs.

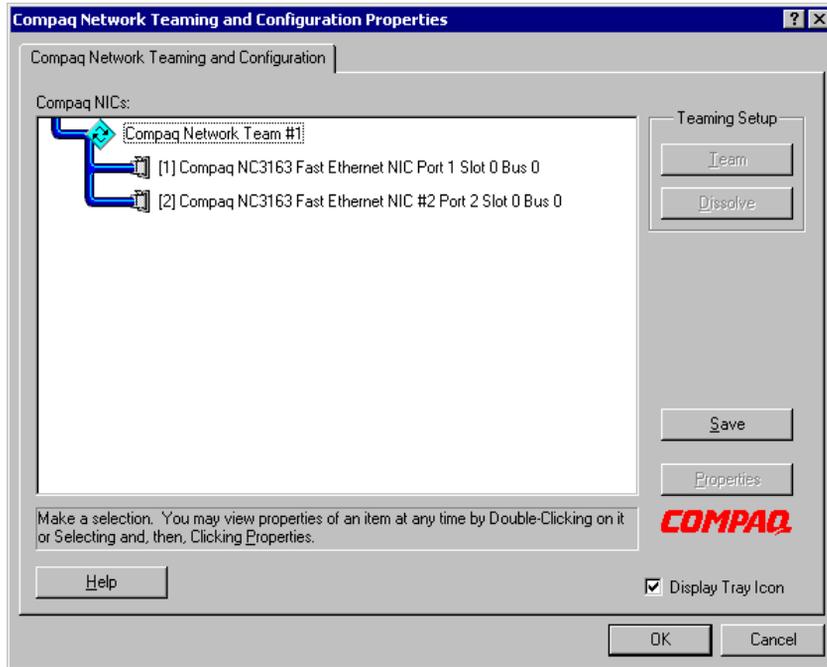


Figure 2-10: Updated CPQTeam Properties dialog box

NIC Teaming Troubleshooting

Problems with the NIC teaming feature are diagnosed by the connection icons displayed in the Compaq Network Teaming and Configuration dialog box. The following table lists the error icons for RJ-45 NICs.

RJ-45	Description
	Active OK —The NIC is operating properly. The driver is installed in the registry and is loaded. If the NIC is a member of a team, the NIC is active.
	Installed inactive —The NIC is installed and is OK, but is not active.
	Cable fault —The driver is installed in the registry and is loaded. The broken cable indicator means that the cable is unplugged, loose, broken, or the switch or hub is not operating properly. If this icon is displayed, check all network connections and make sure the hub/switch is working properly. When the connection is restored, this icon will change.
	Inactive cable fault —A cable fault has occurred while the NIC was inactive.

continued

continued

RJ-45	Description
	<p>Hardware failure—The driver is installed in the registry and is loaded. The driver is reporting a hardware problem with the NIC. This indicates a serious problem. Contact your Compaq authorized service provider.</p>
	<p>Unknown—The server is unable to communicate with the driver for the installed NIC. The NIC is installed in the registry, but the driver is not. This error occurs when the NIC has been installed but the server has not been restarted. If this problem persists after the server has been restarted, the driver has not been loaded or the Advanced Network Control utility is unable to communicate with the driver</p> <p>Note: Only NICs assigned as members of a team are displayed as Unknown. If a teamed NIC is turned off, it displays as Unknown.</p>
	<p>Disabled—The NIC has been disabled through the Device Manager or NCPA.</p>

For more advanced problems with NIC teaming, refer to the help section in the CPQTeam utility.

Creating and Managing Users and Groups

User and group information and permissions determine whether a user can access files. If the NAS device is deployed into a workgroup environment, this user and group information is stored locally on the device. By contrast, if the NAS device is deployed into a domain environment, user and group information is stored on the domain.

To enter local user and group information, see Chapter 7.

The following information is included in this chapter:

- Domain compared to workgroup environments
- Workgroup user administration
 - Adding new users
 - Deleting users
 - Modifying user properties
- Workgroup group administration
 - Adding new groups
 - Deleting groups
 - Modifying group properties — including group members

Creating and Managing File Shares

Files shares must be set up, granting and controlling file access to users and groups. See Chapter 8 for complete information on managing file shares.

The following information is included in this chapter:

- Managing folders
 - Creating new folders
 - Deleting new folders
 - Modifying folder properties
- Managing shares
 - Creating new shares
 - Deleting shares
 - Modifying share properties
- Managing sharing protocol settings
 - Enabling and disabling protocols
 - Modifying protocol settings for Common Internet File System (CIFS), File Transfer Protocol (FTP), Hypertext Transfer Protocol (HTTP)

UNIX-specific information is discussed in the “UNIX File System Management” chapter.

Configuring Data Replication Software

Data replication is the process of making a copy of system data. *StorageWorks* NAS Data Copy is a real-time data replication and failover software product that augments existing data protection and tape backup strategies. This product is not intended to replace regular tape backups.

Using NAS Data Copy, mission-critical data and data that must be protected is marked. NAS Data Copy replicates this data in real-time from the production machine (source) to a backup machine (target). The target machine can be either on-site or off-site. After the initial copy-out, NAS Data Copy monitors any changes to the specified data files and sends only the changes to the target machine.

NAS Data Copy can operate in many different system environments, including:

- **Single machine**—Source and target components are loaded on the same machine, allowing data to be replicated from one location to another on the same machine.
- **One-to-one**—One target machine, having no production activity, is dedicated to support one source machine. An alternative one-to-one scenario is when each machine acts both as a source and a target, actively replicating data to each other.
- **Many-to-one**—Many source machines are protected by one target machine.
- **One-to-many**—One source machine sends data to multiple target machines. The target machines may or may not communicate with each other.
- **Chained**—One or more source machines send replicated data to a target machine that in turn acts as a source machine and sends selected data to a final target machine.

NAS Data Copy is supported for all deployments of the NAS B2000.

A temporary license of NAS Data Copy is included in the NAS B2000 software. To access a permanent user license, order the NAS Data Copy kit from Compaq. Further information and user instructions about NAS Data Copy can be found in the *NAS Data Copy Users Guide*, located in the online help menus and at the Compaq website.

Basic Administrative Procedures

Basic administrative procedures include:

- Setting the system date and time
- Viewing and maintaining audit logs
- Using terminal services
- Changing the language
- Shutting down or restarting the sever
- Backing up or restoring the operating system
- Setting up email alerts

These functions are performed in the **Maintenance** menu of the WebUI.

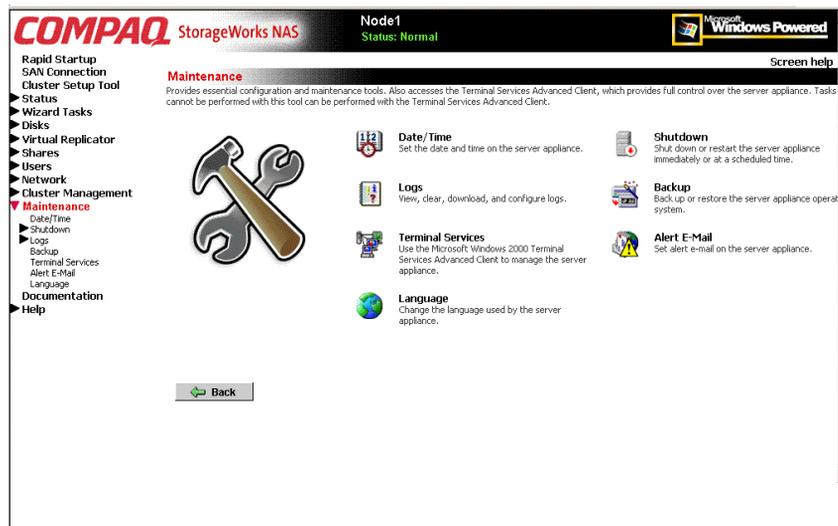


Figure 2-11: Maintenance menu

Setting the System Date and Time

To change the system date or time:

1. From the WebUI, select **Maintenance** and **Date/Time**. The **Date and Time Settings** dialog box is displayed.
2. Enter the new values and then click **OK**. The **Maintenance** menu is displayed.

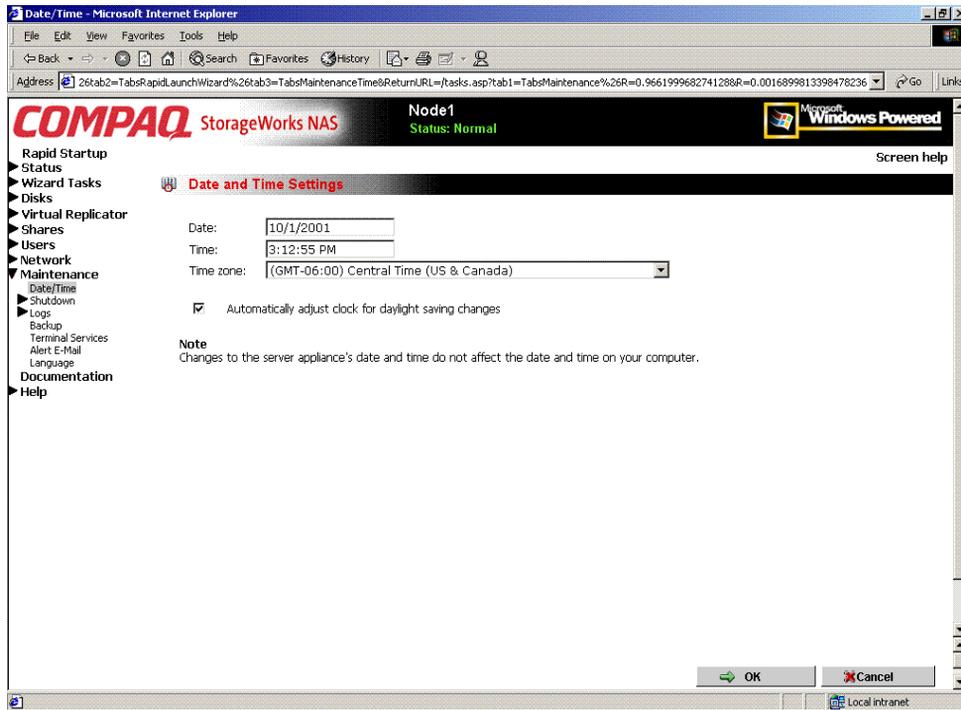


Figure 2-12: Date and Time dialog box

Powering Down and Restarting the Server

IMPORTANT: Notify users before powering down the system. Both UNIX and Windows NT users can be drastically affected if they are not prepared for a system power-down.

1. From the NAS B2000 WebUI, select **Maintenance, Shutdown**. Several options are displayed: Shutdown, Restart, and Scheduled Shutdown.

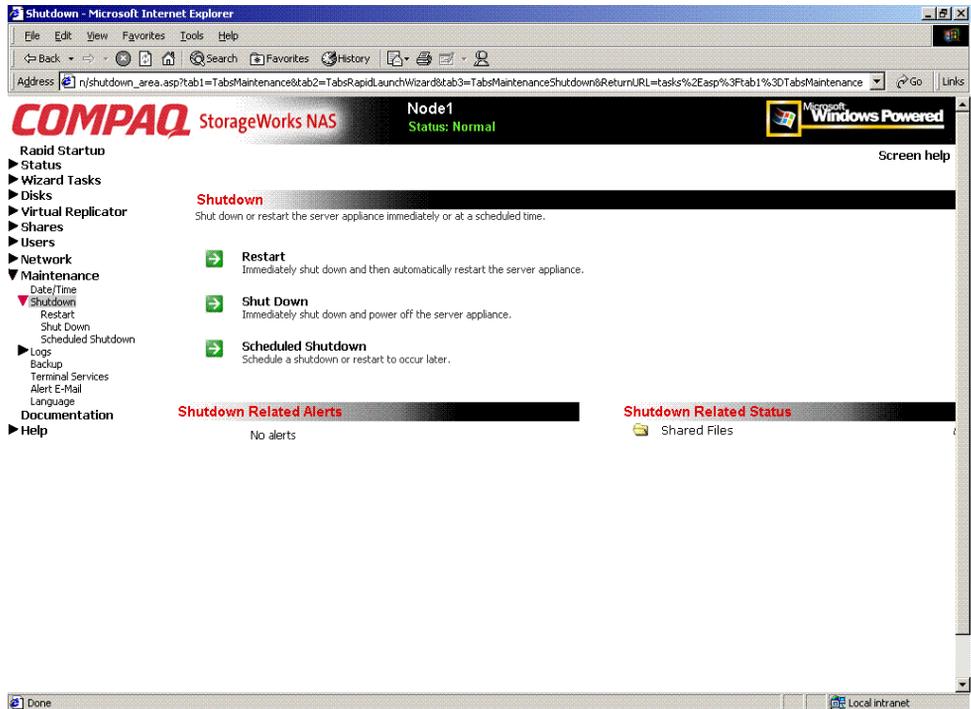


Figure 2-13: Shutdown menu

- a. To shut down and automatically restart the server, click **Restart**.
 - b. To shut down and power off the server, click **Shutdown**.
 - c. To schedule a shutdown, click **Scheduled Shutdown**.
2. Regardless of the choice, a confirmation prompt is displayed. After verifying that this is the desired action, click **OK**. Several status messages are displayed during the shutdown process.

Viewing and Maintaining Audit Logs

A variety of audit logs are provided on the NAS B2000. System events are grouped into similar categories, representing the eight different logs.

To access the logs from the WebUI, select **Maintenance, Logs**. The Logs menu is displayed.

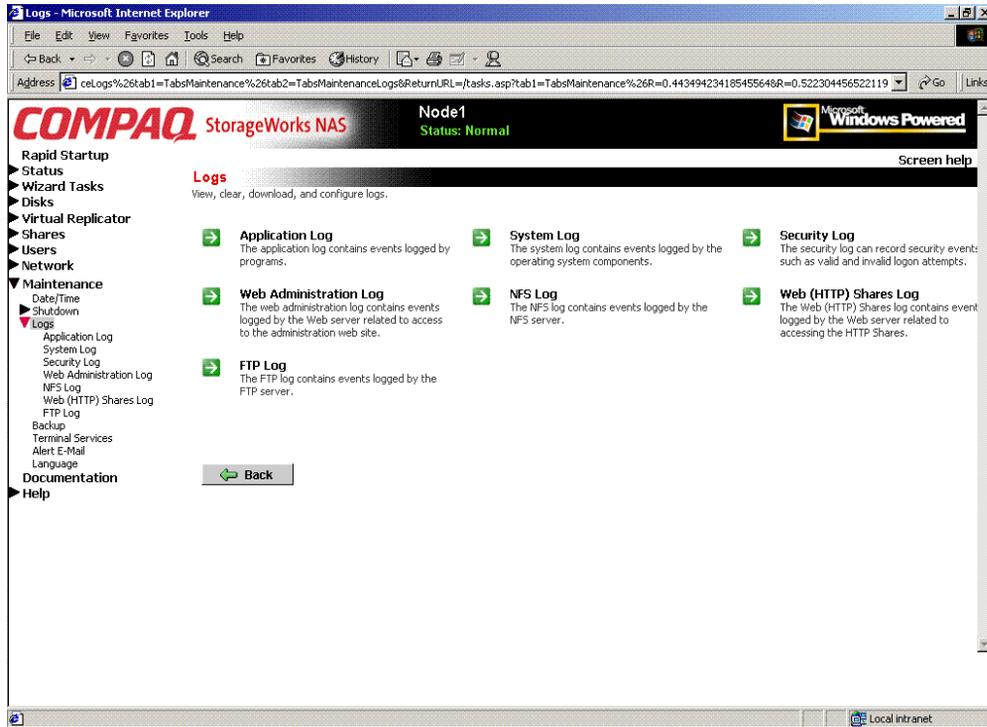


Figure 2-14: Logs menu

A variety of logs are available and are listed in Figure 2-14.

Each log has viewing, clearing, printing, and saving options.

Using Terminal Services

Terminal Services is provided in the WebUI to allow for additional remote system administration and the use of approved third-party applications. SecurePath, backup software, and antivirus programs are examples of approved applications.

In addition, Terminal Services is used to access the Microsoft Management Console (MMC) of the NAS device.

To open a Terminal Services session from the WebUI, select **Maintenance, Terminal Services**. A Terminal Services session is opened. Enter the appropriate password to log on to the server.

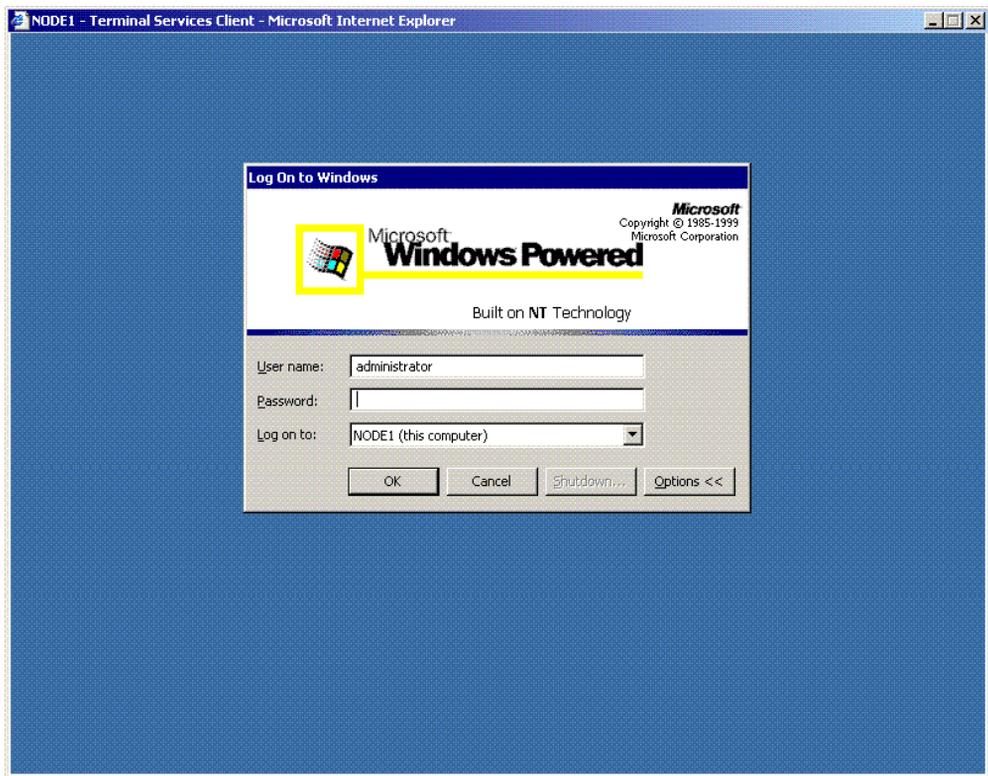


Figure 2-15: Terminal Services session

IMPORTANT: Two open sessions of Terminal Services are allowed to operate at the same time. After completing an application, close that session of Terminal Services by clicking the "X" at the top of the Terminal Services window.

Microsoft Terminal Services has a limitation of two sessions per machine. If a user closes a terminal services session without logging off, the next user to attempt to log on could be locked out of a session.

To unlock the sessions:

1. Go to the Remote Insight Lights-Out Edition board and logon to the console.
2. From the console, click **Start, Administrative tools**, then **Terminal Services Manager**.
3. When the Terminal Services Manager comes up, right-click the user and select **Disconnect**.

Setting up Email Alerts

If desired, the system sends emails of system events to a specified email account. When activated, this feature sends an email whenever system alerts occur.

To activate this option:

1. From the WebUI, select **Maintenance, Alert Email**. The Set Alert E-Mail dialog box is displayed.
2. Select **Enable alert email**.
3. Indicate the types of messages to be sent.
 - Critical alerts
 - Warning alerts
 - Informational alerts

4. Enter the desired e-mail address in the appropriate boxes.
5. To send test e-mails, click **Test**.
6. After all settings have been entered, click **OK**. The Maintenance menu is displayed again.

Changing System Network Settings

Network properties are entered and managed from the **Network** menu. Most of these settings are entered as part of the Rapid Startup process. Settings made from this menu include adding the NAS B2000 to a domain.

Excellent online help is available for these settings. Figure 2-16 is an illustration of the **Network** settings menu.

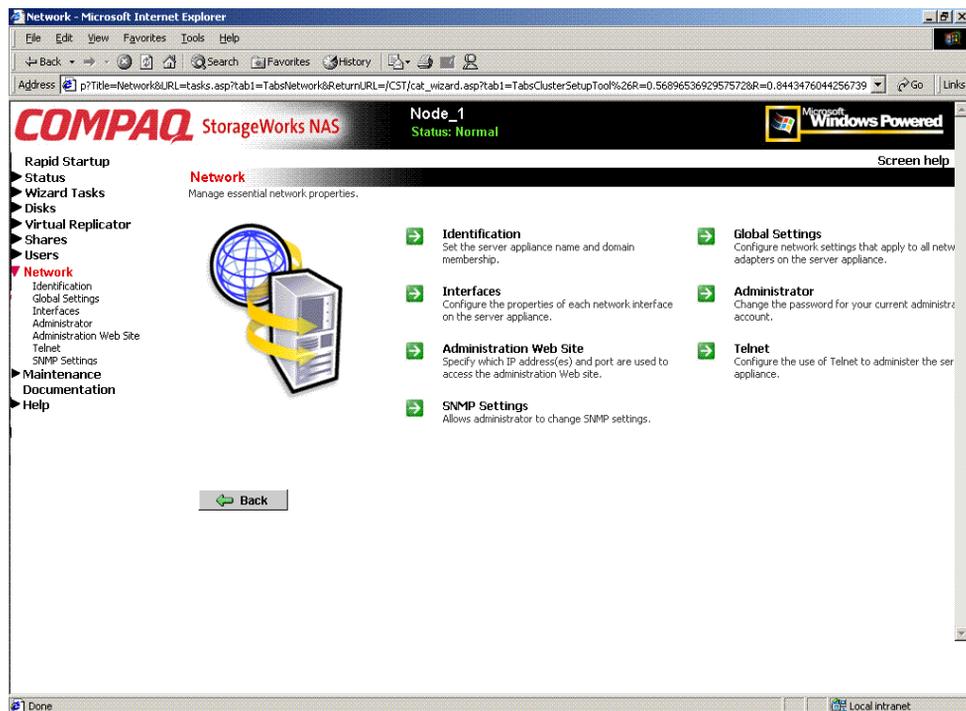


Figure 2-16: Network menu

Storage Management Overview

With the *StorageWorks* NAS B2000 the administrator has complete control over all storage issues. The NAS administrator uses the Compaq Array Configuration Utility (ACU) to manage the hardware storage and uses Virtual Replicator (VR) software to manage the virtual storage.

The NAS B2000 is configured at the factory with default system settings and with the NAS operating system installed. Storage, however, is not pre-configured, allowing the NAS administrator to tailor the organization and configuration of the storage to specific environmental needs. Refer to the “Storage Management Process” section later in this chapter for more information.

The NAS administrator uses the Compaq Array Configuration Utility (ACU) to manage the hardware storage and uses Virtual Replicator (VR) software to manage the virtual storage.

This chapter defines and discusses both physical and virtual storage concepts on the *StorageWorks* NAS B2000, including:

- Storage Management Process
- Physical Storage Overview
 - Physical Hard Drives
 - Arrays
 - Logical Drives (LUNs)
 - Fault-Tolerance Methods
 - Physical Storage Best Practices

- Virtual Storage Overview
 - Pools
 - Virtual Disks
 - Snapshots
 - VR Lifeguard Service
 - Virtual Storage Best Practices

Additional storage management information is included in the following chapters:

- Chapter 4 — Discusses planning decisions in detail.
- Chapter 5 — Discusses hard drive, array, and LUN management procedures.
- Chapter 6 — Discusses pool, virtual disk, and snapshot management procedures.
- Chapter 8 — Discusses folder and share management procedures.

Storage Management Process

The lowest level of storage management occurs at the physical drive level. Physical drives are grouped into arrays for better performance and fault tolerance.

The arrays are then configured with RAID fault tolerance and presented to the operating system as logical drives or units called LUNs.

At the virtual level of storage, VR software is used to take the LUNs and group them together into large-capacity storage pools. These pools are then segmented into virtual disks and re-presented to the operating system as the disks on which to write the data.

Folders, subfolders, and file shares are created on the virtual disks to organize, store, and give access to the system data.

For organizational and documentation purposes, this administration guide separates physical storage from virtual storage. While this chapter provides an overview of storage components and concepts, additional chapters discuss storage management planning and storage management procedures.

See Figure 3-1 for an illustration of these storage management elements.

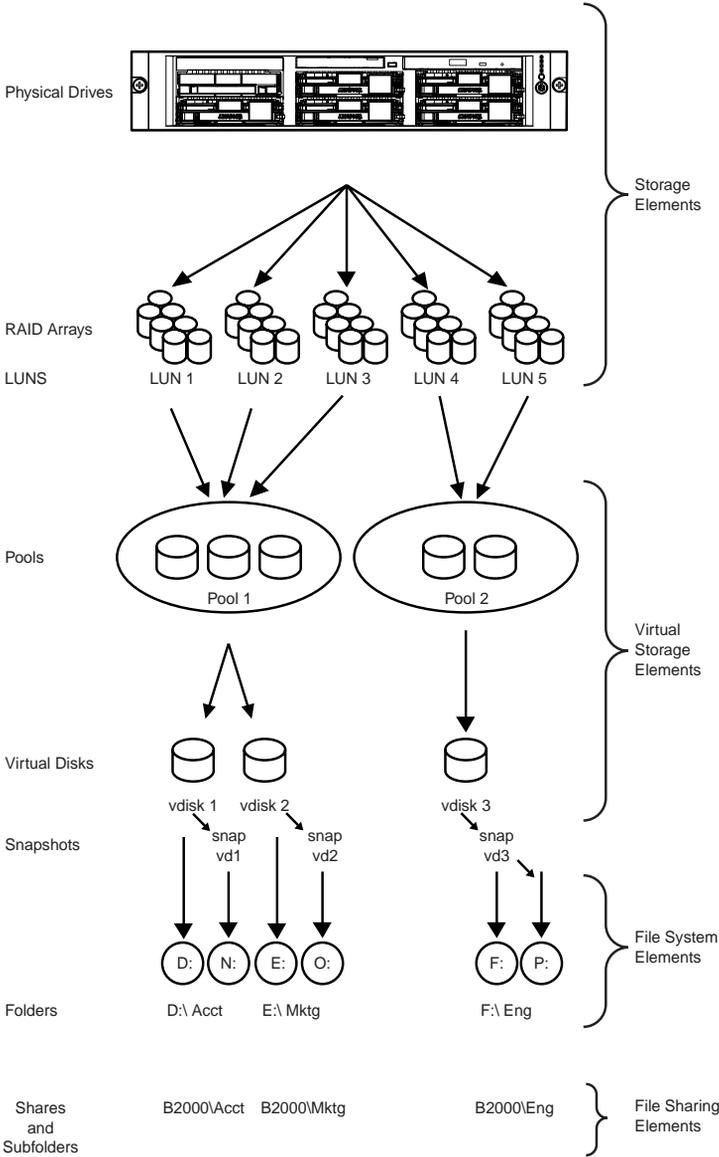


Figure 3-1: Storage Management process

Physical Storage Overview

The NAS B2000 is specifically designed for file serving and can support up to 9.3 terabytes of raw storage capacity, spanning one hundred twenty six 72-GB hard drives spread over 9 StorageWorks 4300 Family storage enclosures. (Each fully populated StorageWorks 4300 Family storage enclosure supports 14 hard drives. This provides a maximum raw capacity of 1 terabyte when using 72-GB hard drives.)

Preliminary physical storage management tasks involve managing:

- Physical Hard Drives
- Arrays
- Logical Drives (LUNs)

Drive array concepts and data protection methods, including fault tolerance options are discussed in this section. This information will help guide decisions on how to best configure the arrays.

Physical Hard Drives

For personal or small business use, the capacity and performance of a single hard drive is adequate. However, larger businesses demand higher storage capacities, higher data transfer rates, and greater security from data loss if drives fail.

Merely adding extra drives to the system increases the total storage capacity, but has little effect on the system efficiency, because data can only be transferred to one hard drive at a time.

Figure 3-2 illustrates the read/write process with separate physical hard drives.

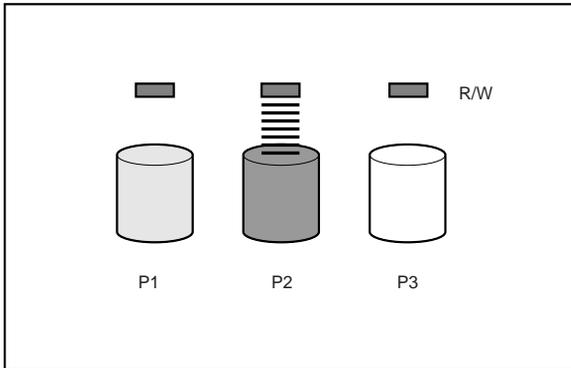


Figure 3-2: Separate physical drive (P1, P2, P3) read/write (R/W) operations

Arrays

With an array controller installed in the system, the capacity of several physical drives can be logically combined into one or more logical units termed **arrays**. When this is done, the read/write heads of all the constituent physical drives are active simultaneously, dramatically reducing the overall time required for data transfer.

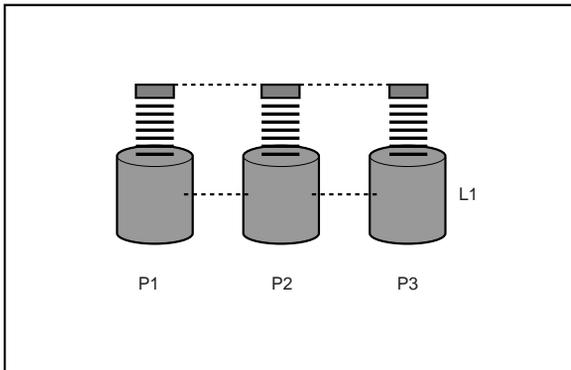


Figure 3-3: Configuring the physical drives into an array (A1) dramatically improves read/write efficiency

Because the read/write heads are active simultaneously, the same amount of data is written to each drive during any given time interval. Each unit of data is termed a **block**. The blocks form a set of data **stripes** over all the hard drives in an array, as shown in Figure 3-4.

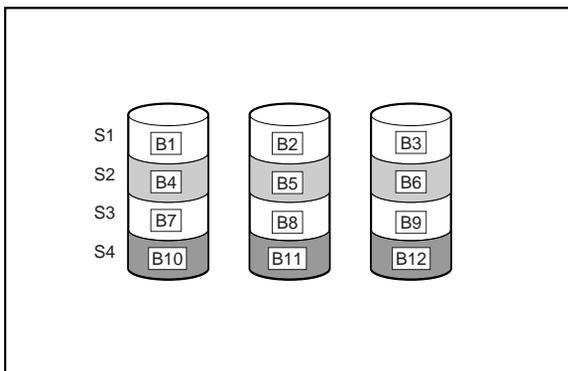


Figure 3-4: RAID 0 (data striping) (S1-S4) of data blocks (B1-B12)

For data in the array to be readable, the data block sequence within each stripe must be the same. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each hard drive in a given array will contain the same number of data blocks. If one hard drive has a larger capacity than other hard drives in the same array, the extra capacity is wasted because it cannot be used by the array.

Logical Drives (LUNs)

As previously stated, drive array technology distributes data across a series of individual hard drives to unite these physical drives into one or more higher-performance arrays. Distributing the data allows for concurrent access from multiple drives in the array, yielding faster I/O rates than non-arrayed drives.

While an array is a physical grouping of hard drives, a logical drive is the configuration of the arrays that is presented to the operating system.

When planning to allocate space on the NAS device, consider that the maximum number of LUNs in a virtual storage pool is eight. To have large pools from which to create virtual disks, make LUNs as large as possible. Pools and virtual disks are discussed in a later section of this chapter.

NOTE: LUNs should not be expanded after they are created because Virtual Replicator cannot recognize the new space associated with the expanded LUNs. To increase system capacity, new hard drives or unassigned hard drives can be configured into new arrays and new LUNs and can then be placed into new pools or added to existing pools.

After the physical drives are grouped into fault-tolerant arrays, they are ready to be converted into logical drives. Options include creating one large logical drive using the entire space of the array or dividing each array into multiple logical drives. Compaq recommends creating one logical drive from the array. Additional physical drives can be added to the array at a later time.

It is important to note that a LUN may extend over (span) all physical drives within a storage controller subsystem, but cannot span multiple storage controller subsystems.

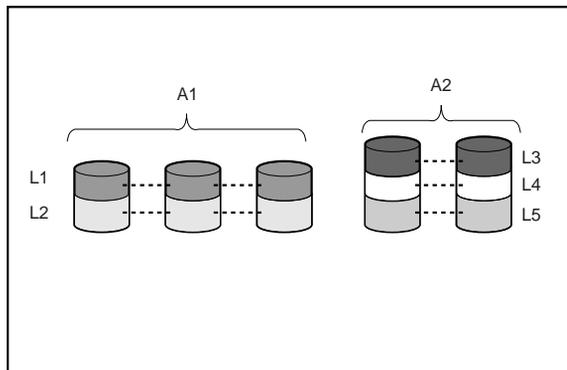


Figure 3-5: Two arrays (A1, A2) and five logical drives (L1 through L5) spread over five physical drives

Drive failure, although rare, is potentially catastrophic. For example, in the previous figure using simple striping, failure of *any* hard drive will lead to failure of *all* logical drives in the same array, and hence to data loss.

To protect against data loss from hard drive failure, arrays should be configured with **fault tolerance**. Several fault tolerance methods have been devised and are described in the following sections.

Fault-Tolerance Methods

Different RAID (redundant array of independent disks) types use different methods of striping the arrays and different ways of writing data and parity to the drives to offer a variety of fault-tolerance and capacity usage. The RAID methods supported by the NAS B2000 include:

- RAID 0—Data Striping only, no fault tolerance
- RAID 1 and RAID 1+0—Drive Mirroring
- RAID 5—Distributed Data Guarding
- RAID ADG—Advanced Data Guarding (ADG)

Further protection against data loss can be achieved by assigning an online spare to an array. This hard drive contains no data and is contained within the same storage sub-system as the other drives in the array. When a hard drive in the array fails, the controller can then automatically rebuild information that was originally on the failed drive onto the online spare. This quickly restores the system to full RAID-level fault tolerance protection.

IMPORTANT: To support ADG, you must install an additional hard drive. RAID 1 support requires an even number of drives. Refer to Table 3-1 for more information.

These fault-tolerance methods are discussed in the following paragraphs.

RAID 0—Data Striping

This configuration provides striping of the array to improve read and write performance, but offers no redundancy of data and therefore no protection against data loss when a drive fails. However, RAID 0 is useful for rapid storage of large amounts of noncritical data (for printing or image editing, for example) or when cost is the most important consideration.

NOTE: The ADG feature is available only with the optional Smart Array 5300 Controller installed.

When creating RAID 0 arrays, carefully consider how many drives to include in the array. While there is no theoretical limit to the number of drives that can be included in a RAID 0 array, there is a practical limit. Statistically, the chance of a drive failure increases with each additional drive that is included in an array. Based upon laboratory testing, Compaq recommends including no more than 7 drives in a RAID 0 array.

See Figure 3-4 for an illustration of the data striping technique.

Advantages

- Highest performance method for reads and writes
- Lowest cost per unit of data stored
- All drive capacity is used to store data—none is used for fault tolerance

Disadvantages

- All data on logical drive is lost if a hard drive fails
- Cannot use an online spare
- Data can only be preserved by being backed up to external drives

RAID 1—Drive Mirroring

In this configuration, information on one drive is duplicated onto a second drive, creating identical copies of the information as shown in Figure 3-6. Therefore, this method provides the best fault tolerance. RAID 1 requires an even number of drives and is the only method for fault-tolerance protection if only two drives are installed or selected for an array. If more than two drives are in an array, the data is striped across all of the drives in the array. This is referred to as RAID 1+0.

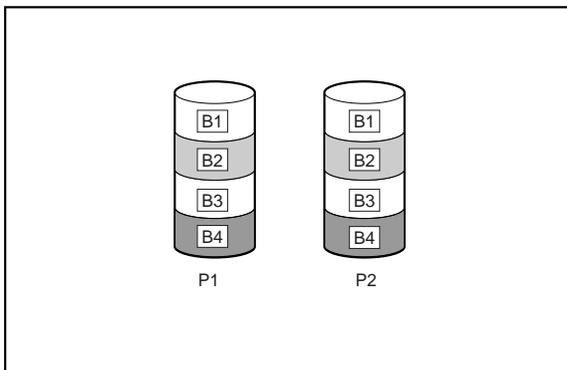


Figure 3-6: RAID 1 (drive mirroring) of P1 onto P2

This method is useful when high performance and data protection are more important than the cost of hard drives. The operating system drives are mirrored. If one drive fails, the mirror drive immediately takes over and normal system operations are not interrupted.

IMPORTANT: Compaq supports a configuration that uses RAID 1 on the system drives in a two-drive RAID array.

IMPORTANT: If two drives being mirrored to each other both fail, data loss occurs.

Advantages

Drive mirroring offers:

- The highest read and write performance of any fault-tolerant configuration.
- Protection against data loss if one drive fails.
- Data preservation in a RAID 1+0 system, when more than one drive fails, as long as none of the failed drives are mirrored to another failed drive.

Disadvantages

Some disadvantages of drive mirroring are:

- Increased expense—Since many drives must be used for fault tolerance and hard drives must be added in pairs.

- Decreased storage capacity—It is only 50% of the total drive capacity.
- Increase data loss—Since data will be lost if two failed drives happen to be mirrored to each other.

RAID 5—Distributed Data Guarding

By this method, a block of parity data (rather than redundant data) is calculated for each stripe from the data that is in all other blocks within that stripe. The blocks of parity data are distributed over every hard drive within the array, as shown in the figure below. When a hard drive fails, data on the failed drive can be rebuilt from the parity data and the user data on the remaining drives. This rebuilt data can be written to an online spare.

This configuration is useful when cost, performance, and data availability are equally important.

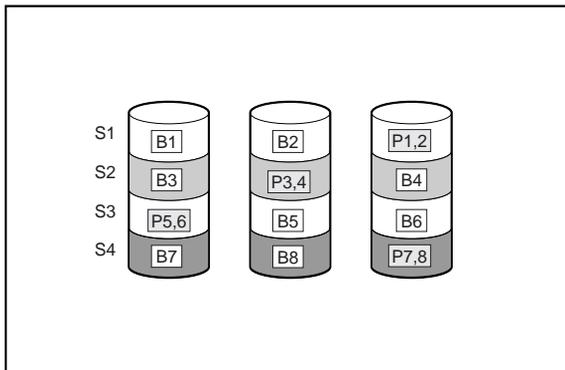


Figure 3-7: RAID 5 (distributed data guarding) showing parity information (P)

Spreading the parity across all the drives allows more simultaneous read operations and higher performance than data guarding (RAID 4). If one drive fails, the controller uses the parity data and the data on the remaining drives to reconstruct data from the failed drive. RAID 5 allows the system to continue operating with reduced performance until the failed drive is replaced. However, if more than one drive fails, RAID 5 also fails and all data in the array is lost.

Distributed data guarding uses the equivalent of one drive to store parity information and requires an array with a minimum of three physical drives. In an array containing three physical drives, distributed data guarding uses 33 percent of the total logical drive storage capacity for fault tolerance; a 14-drive configuration uses seven percent.

NOTE: Given the reliability of a particular generation of hard drive technology, the probability of an array experiencing a drive failure increases with the number of drives in an array. Compaq recommends the number of drives in an array not exceed 14.

Advantages

Distributed data guarding offers:

- High read and write performance.
- Protection against data loss if one drive fails.
- Increased usable storage capacity, since capacity equal to only one physical drive is used to store parity information.

Disadvantages

Some disadvantages of distributed data guarding are:

- Relatively low write performance.
- Increase data loss—Since data will be lost if a second drive fails before data from the first failed drive has been rebuilt.

RAID ADG—Advanced Data Guarding

RAID ADG is similar to RAID 5 in that parity information is generated (and stored) to protect against data loss caused by drive failure. With RAID ADG, however, two different sets of parity data are used. This allows data to still be preserved if two drives fail. As can be seen from Figure 3-8, each set of parity data uses up a capacity equivalent to that of one of the constituent drives, for a total parity usage of 2 drives of space.

This method is most useful when data loss is unacceptable, but cost must also be minimized. The probability that data loss will occur when configured with RAID ADG is less than when configured with RAID 5.

NOTE: The ADG feature is available only with the optional Smart Array 5300 Controller installed.

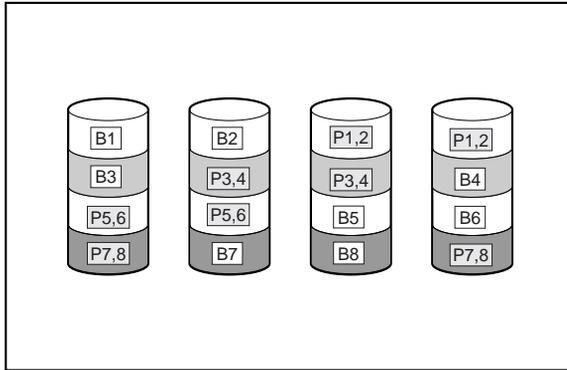


Figure 3-8: RAID ADG (advanced data guarding) with two sets of parity data

Advanced Data Guarding technology offers the best combination of fault tolerance and usable disk space among RAID levels.

This patented Compaq technology allows the safe deployment of large-capacity disk drives and the creation of very large storage volumes without expensive overhead to protect business critical data. This technology provides more flexibility in responding to drive failures without the fear of costly server downtime.

Advance Data Guarding protects against multiple disk failures, while requiring the capacity of 2 drives in an array of up to 56 disk drives to be set aside for dual sets of distributed parity data. It provides data protection greater than RAID 0+1 while having the capacity utilization efficiency similar to RAID 5.

Advantages

- High read performance
- High data availability—any two drives can fail without loss of critical data

Disadvantage

The only significant disadvantage of RAID ADG is a relatively low write performance (lower than RAID 5), due to the need for two sets of parity data.

The table below summarizes the important features of the different kinds of RAID supported by the MSA1000 controller. The decision chart in the following table may help determine which option is best for different situations.

Table 3-1: Summary of RAID methods

	RAID 0	RAID 1 / RAID 1+0	RAID 5	RAID ADG
	Striping (no fault tolerance)	Mirroring	Distributed Data Guarding	Advanced Data Guarding
Usable drive space*	100%	50%	67% to 93%	50% to 95%
Usable drive space formula	n	n/2	(n-1)/n	(n-2)/n
Minimum number of hard drives	1	2	3	4
Maximum number of hard drives	N/A	N/A	14	56
Tolerant of single hard drive failure?	No	Yes	Yes	Yes
Tolerant of multiple simultaneous hard drive failure?	No	For RAID 1+0, if the failed drives are not mirrored to each other	No	Yes

continued

Table 3-1: Summary of RAID methods *continued*

	RAID 0	RAID 1 / RAID 1+0	RAID 5	RAID ADG
Read performance	High	High	High	High
Write performance	High	Medium	Low	Lowest
RAID overhead	Low	High	Medium	Medium

***Note:** The value for usable drive space is calculated assuming a maximum of 14 hard drives of the same capacity (or a maximum of 42 for RAID ADG, which is the maximum number of drives in a single MSA1000 storage subsystem) with no online spares. Compaq recommends that these maximum figures are not exceeded when configuring a drive array, due to the increased likelihood of logical drive failure with more hard drives.

Physical Storage Best Practices

Minimally, choosing the best disk carving strategy includes the following policies:

- Analyze current corporate and departmental structure.
- Analyze the current file server structure and environment.
- Plan properly to ensure the best configuration and use of storage.
 - Determine the desired priority of fault tolerance, performance, and storage capacity
 - Use the determined priority of system characteristics to determine the optimal striping policy and RAID level.
- Include the appropriate number of physical drives in the arrays to create LUNs of desired sizes.
 - For RAID 0, include between 1 and 7 drives in each array.
 - For Raid 0+1, include between 2 and 56 drives in each array.
 - For RAID 5, include between 3 and 14 drives in each array.
 - For RAID ADG, do not exceed 56 drives in an array.

Complete and detailed planning information is included in Chapter 4.

RAID arrays and LUNs are created and managed using the Compaq Array Configuration Utility (ACU). See Chapter 5 for detailed information on creating and managing arrays and LUNs.

Virtual Storage Overview

After the LUNs are created, they are presented to the operating system for use. At this time, the virtual storage management software (Virtual Replicator) acquires the LUNs, groups them into pools, divides the pools into virtual disks, and re-presents the virtual disks to the operating system.

Virtual storage elements are managed using the Virtual Replicator (VR) software. VR provides advanced storage management capabilities in Windows-based computing environments. Innovative storage management features simplify storage configuration and management and enhance availability and scalability.

Virtual storage includes the following:

- Pools—formed from LUNs, can be up to 8 terabytes
- Virtual disks—created from available pool space, can be up to 2 terabytes
- Snapshots—are instant copies of the virtual disks

For specific information on pools, virtual disks, and snapshots, see the topic-specific sections of this chapter.

Pools

VR enables the grouping of hardware LUNs into a logically proportioned pool of drive space. Any number of pools can be created, using industry-standard storage components and controller-based, fault-tolerant drive RAID arrays and LUNs. The LUNs provide space for the pool in the same way that physical drives provide drive space for a RAID array.

Pool Facts:

- Each pool can contain up to eight logical disks (LUNs).
- If multiple LUNs are used to form a pool, they must all be from the same controller.
- All LUNs that make up a pool should have the same high-availability characteristics (RAID levels, striping methods, and drive capacity.)
- Each pool can be up to 8 terabytes (each LUN can be as large as 1 terabyte.)
- By default, 10% of the pool space is reserved for snapshots, leaving 90% of the pool available for the creation and use of the virtual disks.
- Virtual disks are created from pools.

Virtual Disks

Virtual Replicator controls how data is stored on a virtual disk. A virtual disk is composed of storage from a single pool and cannot be composed of storage from different pools. The virtual disks perform and behave in exactly the same way as physical drives do. They can be formatted and drive letters can be mapped to them, and the system can read from and write to them.

Disk virtualization allows drive space to be tailored to the size required by users and their applications. From a pool, the administrator may wish to create three separate virtual disks that exactly match the space needs of the users or groups requesting space, or they may create a single, large virtual disk from a pool. For example, if a user needs 650 MB of drive space, a 650-MB virtual disk can be created. Alternatively, if a system has a 2-terabyte database, LUNs can be combined into a single pool and a 2-terabyte virtual disk that spans the physical storage can be created. The size of the virtual disks can range from 10 MB to 2 terabytes, depending on free pool space and other limits set at the time of pool creation. A limitation is that virtual disks may not “span” pools—they must be created from a single pool.

Virtual disks are the storage entities that the operating system recognizes as drives and upon which data is placed.

Virtual Disk Facts:

- Must be between 10 MB and 2 terabytes.
- May only create eight virtual disks from a pool.
- Can initially be small, and then expanded through an online volume growth procedure.
- Can use up to 90% of the pool space, with 10% reserved for snapshots by default.

Snapshots

Virtual Replicator lets the administrator make instant replicas, called snapshots, of virtual disks in a matter of seconds. Snapshots enable the instant creation of multipurpose virtual replicas of production data without having to physically copy the data. They can be used to immediately recover a lost file or directory, to test a new application with realistic data without affecting the “real” data, and to serve as a source of data for backups. Snapshots are a temporary backup of the data and are not meant to be permanent.

When using snapshots, performance of the virtual disk may be affected, depending on the rate that data is changing and the number of snapshots kept for each virtual disk. Read performance of the virtual disk remains constant, regardless of the presence of snapshots. Read performance of the snapshot is identical to that of the virtual disk. Write performance, however, may vary. Each initial write to a virtual disk area causes a copy-out to the snapshot, and the initial write is slower than if a snapshot is not being used. Copy-out is not performed on subsequent writes to the same virtual disk block, so write performance is unaffected after the initial write to each block.

Predicting the exact effect of snapshots on any particular virtual disk is difficult, because several variables are involved. These variables include the type of applications accessing the data and the rate of change of the files on the virtual disk. When a high percentage of writes is made to the same area, as when a file is constantly rewritten, the effect is called write locality. Virtual disks with high write locality experience less performance degradation due to snapshots.

Snapshots can be read-write, and if they are shared, users can access a snapshot and edit the data. If snapshots are shared with write access enabled, a snapshot of the original snapshot should be created. There is no backup of the original snapshot unless a snapshot of it is taken. By controlling the share permissions, administrators can make snapshots read-only.

Snapshot Facts:

- Snapshots are created on a per-virtual disk basis.
- Snapshots of snapshots can be created.
- Snapshots can be read-only or read-write.
- Snapshots must be assigned to a drive letter to be accessible.
- Snapshots can be shared in the same manner as virtual disks.
- Snapshots are meant to be temporary in nature.
- Snapshots are automatically deleted if pool space becomes critical.

Snapshots and Pool Sizing

When initially implementing snapshots, the first priority is to gain an understanding of the amount of additional pool space required by the snapshots. As previously described, this assessment is a function of the write activity, its locality, the projected lifetime of the snapshot, and the number of snapshots per virtual disk. Because the space required for snapshots tends to grow over time, by default, 10 percent of the pool is reserved for snapshots

Virtual Replicator includes a Snapshot Planner utility that can be used before moving to a NAS device. The Snapshot Planner gathers information about your current server and its configuration to use when configuring the pools and logical drives on the new server. The Snapshot Planner utility must be installed and run only on existing servers and must not be run on the NAS device. This utility is included on the Compaq StorageWorks NAS Servers Configuration Utility CD. The Snapshot Planner provides complete pool and volume sizing information. The size of the virtual disks and the size of the pool are equally important.

Snapshots and Drive Defragmentation

A drive defragmenter attempts to consolidate files on a drive by reading various parts of the files and rewriting them to become contiguous on the drive. When virtual disks are created from the drive space pool, the software attempts to make them as contiguous as possible on the underlying storage units (RAID arrays and LUNs).



CAUTION: Defragmentation must not be performed if snapshots exist. To defragment a disk, first delete the snapshots.

IMPORTANT: Drive defragmentation only operates on virtual disks formatted with a 4-KB or smaller allocation size. By default, the NAS device uses a 16-KB allocation cluster size.

IMPORTANT: The Windows 2000 operating system native defragmenter works on VR virtual disks.

To perform drive defragmentation:

1. From the Group Policy Editor, select **Start, Run**. Enter: `gpedit.msc`
2. Follow the path **User Configuration \ Administrative Templates \ Windows Components \ Microsoft Management Console \ Restricted/Permitted Snap-ins**.
3. Double-click **Properties** to open the Disk Defragmenter properties.
4. Select the **not Configured** option and click **OK** to continue.
5. From the Disk Defragmenter Management Console, select **Start, Run**. Enter: `defrag.msc`.
6. Defragment the appropriate disks.

Snapshot Naming

Snapshots should be named to reflect the relationship with their parent virtual disk. In the case of a snapshot of a snapshot, the parent-child relationship of snapshots of snapshots is not maintained on the server. Appropriate snapshot naming and volume naming must be used to identify each snapshot individually. For example, if the virtual disk is named A:

- Name the parent “SnapshotA”
- Name the child “ChildofSnapshotA”

Snapshots and Backup

Because snapshots are quick to create, it is possible to capture a coherent view of the virtual disk data with little or no application downtime. Lack of application downtime removes the traditional backup window or the amount of time taken to back up to offline media. While many applications must be shut down to capture an accurate backup, snapshots capture a point-in-time view of the data that can be used as the source of backup data. Applications can continue processing against the virtual disk data. Therefore, applications may only have to be interrupted for a few seconds during the snapshot process.



CAUTION: Snapshots are not a replacement for reliable, periodic data backup. If free pool space becomes critical, snapshots are automatically deleted. See the “VR Lifeguard Service” section. In addition, snapshots are a short-term convenience and reside on the same physical drives as the data. If something happens to the data drives, the snapshots are also affected. Read Appendix A for suggestions on how to back up the NAS device.

Snapshots, Restoring Files, and Directories

Snapshots can be used as a highly efficient way to maintain online backups of a virtual disk, enabling immediate file and directory recovery. Snapshots can be assigned to a drive letter and used exactly like virtual disks. If a file or directory is lost or corrupted, it can be recovered easily. The administrator can switch to the snapshot drive and copy the file or directory back to its original location on the virtual disk or an alternate location.

IMPORTANT: To preserve the integrity of a point-in-time snapshot of a virtual disk, ensure that it is exported as read-only.

VR Lifeguard Service

VR Lifeguard is a utility that watches storage pools for pool-full conditions and deletes snapshots when a pool is nearly full. Pool-full conditions cause writes to the drive to be delayed. These conditions can occur when snapshots are being used and insufficient space is left in the pool to allow copying the data from the original virtual disk to the snapshot. If this situation occurs, the system delays the write to both the original virtual disk and the snapshot.



CAUTION: If the pool is improperly sized and there is an insufficient amount of free space in the pool, Lifeguard automatically deletes the oldest snapshots. This action prevents a critical space deficit that could delay or lose writes.

Because snapshots use free space in the pool, a pool-full condition can be a result of poor or improper pool space planning. Make sure there is sufficient space for snapshots. By default, 10 percent of the pool space is reserved for snapshots.

Before attempting to use a snapshot to restore a virtual disk, special care must be taken to make sure there is sufficient free space in the pool. If space becomes critical during the restore, the potential for a pool-full condition exists. In this case, Lifeguard may delete the snapshot from which it is restoring.

The Lifeguard service is installed and automatically started. The service checks each drive letter on a system to see if it is a virtual disk or a snapshot. If it is a snapshot, Lifeguard checks the storage pool that the drive belongs to for its amount of current free space. If the amount of free space goes below the threshold, Lifeguard deletes one or more snapshots to make more space available in the pool. This check is performed, by default, once every 10 seconds.

Lifeguard Facts:

- Checks for pool-full situations periodically based upon pool policy settings.
- Begins deleting snapshots when the free pool space drops below the default 1,024-MB threshold
- Deletes the oldest snapshot first
- Continues to delete up to 12 snapshots at the pre-set interval until free pool space is above the threshold
- Reserves 10 percent of the total pool for snapshot use and restricts online volume growth from causing a pool-full condition by not allowing a volume to grow into this reserved space
- Deletes a snapshot, even if it is the only snapshot in existence
- Posts events in the Application Event Log. An event is entered for each drive that exceeds the threshold, and additional events will be entered when a snapshot is deleted.

Virtual Storage Best Practices

Virtual Replicator is the software used to manage the storage on the NAS B2000. Some administrators may be familiar with other storage management software programs, such as Windows Disk Management.

Use the following guidelines when managing the storage on the NAS device:

- If a task can be performed using either VR or another tool, always use VR. Use VR to map drive letters to virtual disks and snapshots, not Disk Management.
- Use VR to monitor free space in a pool, rather than Windows Explorer. Windows Explorer does not recognize pools.
- Do not use VR in conjunction with Windows Disk Management/Logical Disk Manager. Doing so can result in unpredictable behavior.
- When using VR management tools, close Logical Disk Manager.
- Do not perform Disk Management tasks on disks that belong to VR pools.
- Do not perform online volume growth of either basic or virtual disks at the same time as other major disk management operations, such as defragmenting and disk checking.
- As an extra precaution, before performing any task with Disk Management, use Device Manager to scan for hardware changes to update disk information.

Refer to the *Compaq SANworks Virtual Replicator System Administration Guide* for additional information about VR.

Storage Management Planning

This chapter details issues surrounding storage configuration, storage sizing, and performance planning for the *StorageWorks* NAS B2000. Proper planning for system storage and performance is critical to the successful deployment of the NAS B2000. Improper planning or implementation can result in wasted storage space, degraded performance, or inability to expand the system to meet growing storage needs.

There are several key points to keep in mind when planning system storage needs. These range from desired priorities of different system characteristics, to software restrictions, to hardware behaviors, to the ramifications of specific configurations.

This chapter documents the main issues and rules to follow when planning and configuring the storage of the NAS B2000, including:

- Fundamental Storage Configuration Planning Issues
 - System Priorities
 - Array Configuration (Striping) Methods
 - Recommended System Configurations
- Physical Storage Planning Issues
 - Hard Drive Sizes and Types
 - Use and Number of Spare Disks
 - LUN Sizing

- Virtual Storage Planning Issues
 - Pool Considerations
 - Virtual Disk Considerations
 - Shares Considerations
- Storage Management Planning Scenarios
 - A Complete and Detailed Storage Planning Example
 - A Simple Sizing Comparison
 - An Example of a Storage Subsystem Using Different Array Configurations
 - Planning Worksheet
- Migration Issues

All of the factors mentioned in this chapter must be taken into consideration when planning, laying out, and implementing the storage architecture. The decisions and implementations made during the planning and configuration stage affect the performance, availability, and expandability of the configuration. Any oversights or mistakes made during this phase will be difficult to correct later.

Fundamental Storage Configuration Planning Issues

Prior to actually configuring the drives in the storage enclosures into arrays, LUNs, pools, virtual disks, and shares, extensive analysis of desired system performance must be completed. Therefore, preliminary storage planning topics include:

- System Priorities
- Array Configuration (Striping) Methods
- Recommended System Configurations

System Priorities

The first and most important part of storage management planning is the ranking of basic desired system characteristics. Based on the type of data that will be stored on and accessed from the system, the importance of the following three system characteristics must be determined:

- Fault Tolerance
- Capacity Utilization
- I/O Performance

The optimal configuration method of the system storage depends on the ranking of these characteristics.

As shown in Figure 4-1, these system traits are independent and unrelated. One trait must be chosen as the most important. The ranking of one trait as most important automatically ranks the other characteristics as second and third. After the primary characteristic has been determined, one of the remaining two traits must be declared as second in importance. With the desired system characteristics now ranked, the optimal configuration method can be selected from the following paragraphs and figures in this section of this document.

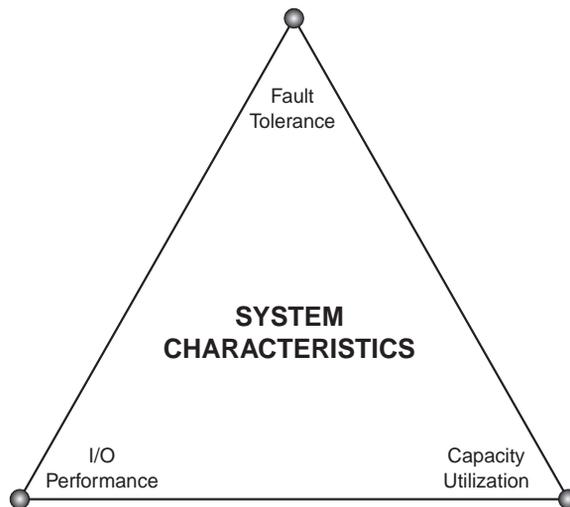


Figure 4-1: System characteristics

Because the physical configuration of the arrays determines whether the system is optimized for fault tolerance, performance, or capacity utilization, a preliminary discussion on the configuration striping methods is warranted.

Array Configuration (Striping) Methods

LUNs are composed of the physical storage arrays (RAID arrays) and are presented to the operating system as disk devices. The physical configuration of the arrays affects both the performance and the high-availability characteristics of the units.

The arrays must be configured in a manner that strikes the desired balance between fault tolerance, storage efficiency, and performance.

There are two methods for configuring the physical layout of the disk arrays:

- Vertical striping
- Horizontal striping

In a vertical configuration, a single RAID array uses one physical drive from each storage enclosure. In a horizontal configuration, the RAID array uses multiple drives contained within one or more storage enclosures. Figure 4-2 and Figure 4-3 illustrate examples of possible array striping configurations.

As a point of planning, horizontal and vertical arrays have their advantages and disadvantages. Each should be used in the appropriate environment to optimize the desired system characteristics.

Vertical Array Configurations

For most RAID configurations, vertical striping of arrays is the only configuration method that ensures no single point-of-failure (NSPOF). Because vertical striping uses drives from different storage enclosures in the same array, an MSA1000 storage enclosure with no additional disk storage enclosures does not permit vertically striped arrays. To most effectively implement vertical striping of RAID arrays, Compaq recommends using a fully populated MSA1000 storage enclosure with two additional disk storage enclosures. This configuration allows for greater choices of RAID configurations, as well as for the creation of larger arrays than one additional disk storage enclosure would allow.

Using an example of RAID 5 vertically striped arrays with one drive from each enclosure included in an array and parity information distributed throughout the array, if a single storage enclosure fails, only one drive in each array will fail. All arrays are still online and the data can be rebuilt from the distributed parity information. For some RAID configurations, vertical striping must be used for ultimate fault tolerance.

Depending on the RAID configuration chosen, a vertical array will consist of at least one drive from each storage enclosure. See Figure 4-2 and Table 4-1 for RAID-specific disk use information of vertical carving configurations.

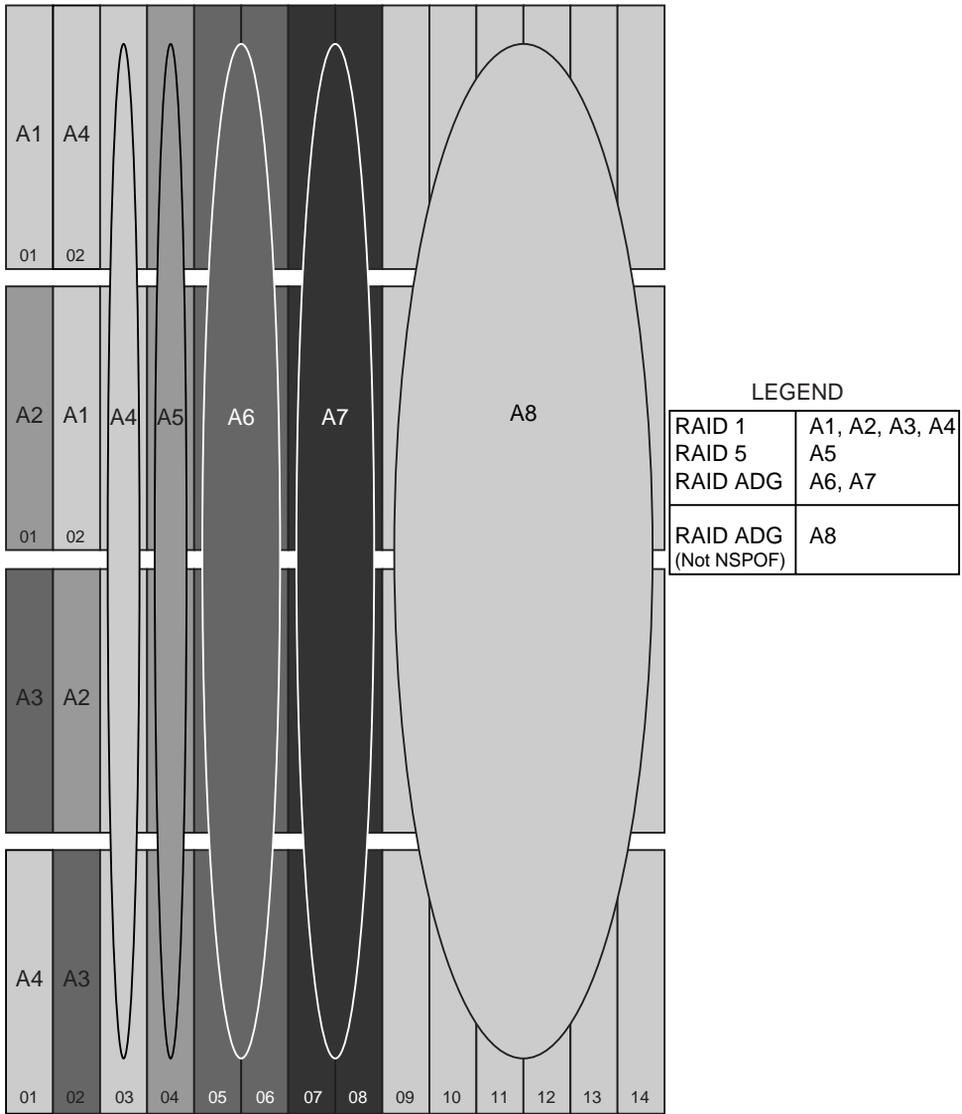


Figure 4-2: Vertical Array configurations

Table 4-1: Vertical Carving Disk Use per RAID Level

RAID Method	Number of Drives	Space Used for Fault Tolerance
RAID 0	Not recommended	
RAID 1 (NSPOF)	2 (one per enclosure)	1 drive (50%)
RAID 1+0	Not recommended	
RAID 5 (NSPOF)	4 (one per enclosure)	1 drive (33%)
RAID ADG (NSPOF)	8 (two per enclosure)	2 drives (50%)
RAID ADG	User defined	2 drives

NOTE: The ADG feature is available only with the optional Smart Array 5300 Controller installed.

IMPORTANT: RAID 0 and RAID 1+0 are not recommended when using vertical striping.

One disadvantage of vertical configurations is that the arrays are relatively small, which may result in the rapid exhaustion of available drive letters. After all letters have been used, no additional virtual disks can be presented to or used by the system. Creating small arrays divides the storage capacity among more LUNs, pools, and virtual disks than large arrays. With each virtual disk and snapshot using a letter, deployments using vertical striping can quickly run out of available letters.

One potentially major disadvantage of vertical arrays is that they are not very flexible when it comes to growing the storage at a later date. Using small arrays means that the same amount of storage would occupy more LUNs. Because a pool can contain no more than eight LUNs, the pool configuration resulting from adding many small LUNs into a single pool would limit or eliminate the ability of the pool to be expanded.

An additional disadvantage of vertical configurations is that the creation of many small arrays, LUNs, pools, virtual disks, snapshots, and shares increases administrative and management overhead.

A final disadvantage of vertical array configurations and their small arrays and LUNs is that a greater proportion of the total capacity must be used for fault tolerance.

Horizontal Array Configurations

Horizontal striping allows for the creation of large arrays and LUNs, and offers the best combination of capacity utilization and performance. See Figure 4-3 and Table 4-2 for information about horizontal array configurations and disk use.

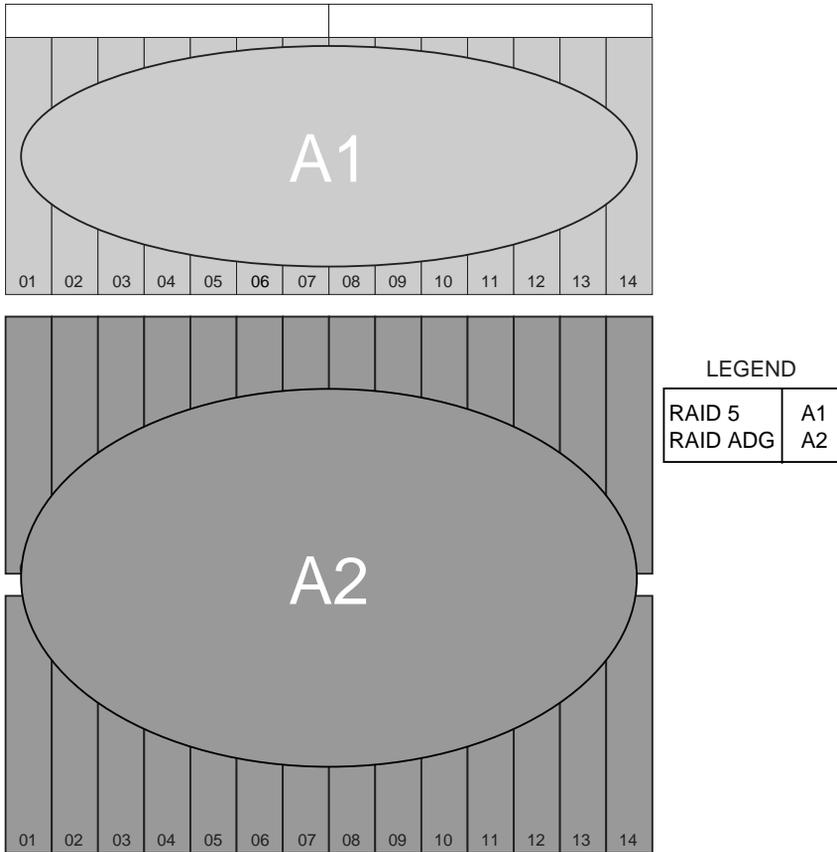


Figure 4-3: Horizontal Array configurations

Table 4-2: Horizontal Configuration Disk Use per RAID Level

RAID Method	Number of Drives	Space Used for Fault Tolerance
RAID 0	Between 2 and 56	None
RAID 1	2	1 (50%)
RAID 1+0	Between 2 and 56	(50%)
RAID 5	Between 3 and 14	1 drive (33% to 7%)
RAID ADG	Between 4 and 56	2 drives (50% to 4%)

In addition to creating larger LUNs than vertical configurations, the biggest advantage of horizontal arrays is the ability to dynamically grow the storage on an as-needed basis. It is very easy to purchase an additional storage enclosure full of drives, connect the storage enclosure, set up the arrays and LUNs, and create additional pools and virtual disks. Alternatively, the new capacity could be used to expand existing arrays, pools and virtual disks.

The major disadvantage of horizontal arrays is loss of access to the data if a storage enclosure fails. The array and its data may not survive the failure, depending on the circumstances in which the failure occurred. In the worst case, there is a total loss of data and the array must be reconfigured and the data must be restored from backup. In the best case, the array and its data will reappear after the failure is repaired.

Keep in mind that regardless of the number of drives in an array, RAID 5 reserves one disk worth of space to contain the parity data, meaning that one disk worth of space is not available for data storage. Similarly, RAID ADG uses the equivalent of two drives of space for the two sets of parity information maintained. Using horizontal striping rather than vertical striping allows more drives to be incorporated into an array, thus reducing the proportion of capacity reserved for parity information.

NSPOF Horizontal Array Configurations

A hybrid RAID striping configuration combines horizontal striping with vertical mirroring, allowing NSPOF fault protection in horizontally striped arrays. This configuration offers the best NSPOF combination of fault tolerance, I/O performance, and capacity utilization.

Using horizontal RAID 1+0 arrays with the mirrored drives in a separate disk storage enclosure permits the creation of larger arrays than a vertically striped RAID 1 NSPOF configuration.

See Figure 4-4 for an illustration of this technique.

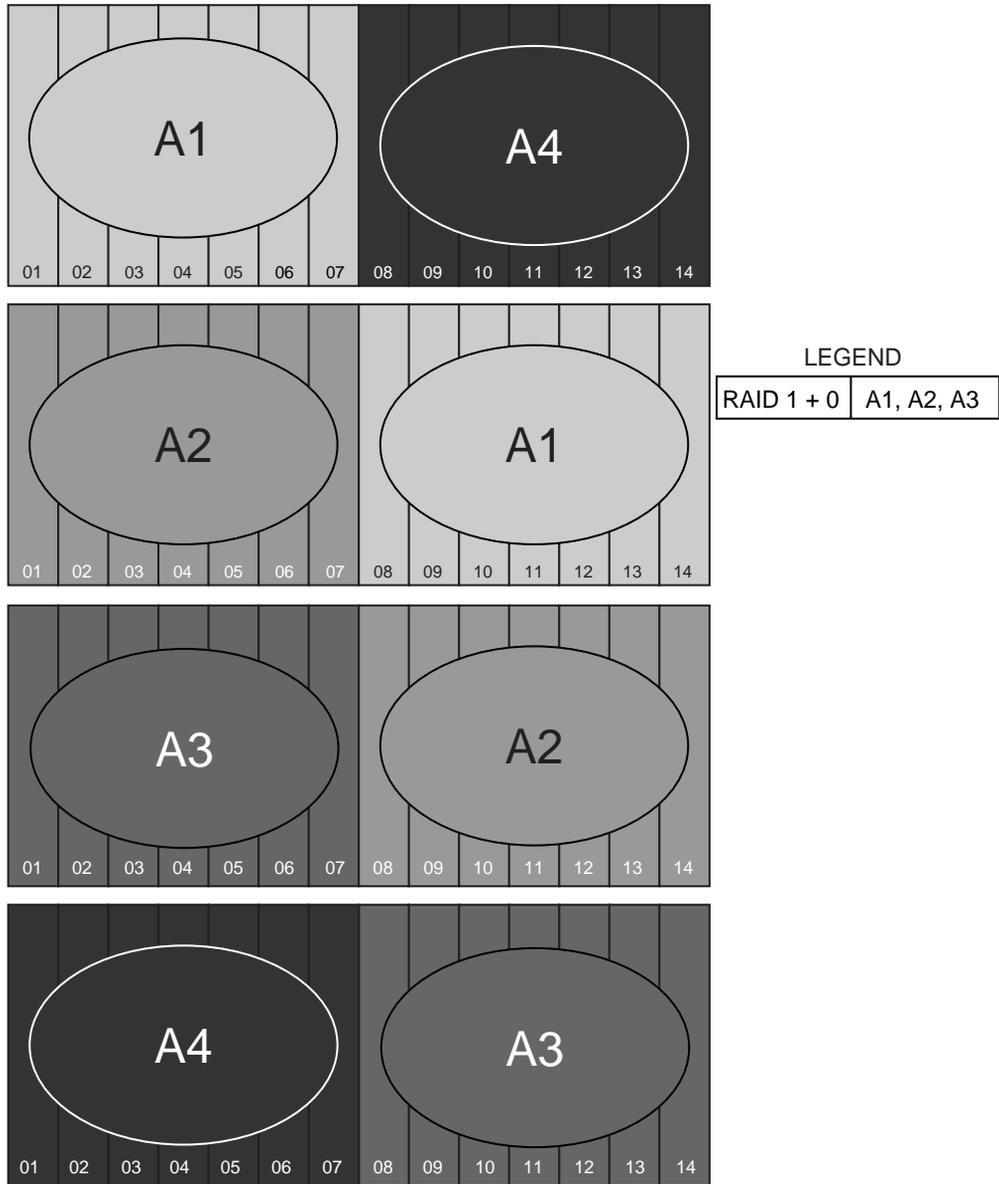


Figure 4-4: NSPOF Horizontal Array configuration

The only requirement of this configuration is that the array must have an equal number of drives in each storage enclosure, with the drives in one storage enclosure containing the data and the drives in the other storage enclosure mirroring the data.

The example in Figure 4-4 illustrates a configuration that maximizes storage capacity and I/O performance, while maintaining an NSPOF configuration and reducing administrative overhead. This configuration divides each storage enclosure in half, including seven drives from each storage enclosure in an array, resulting in a total of 14 drives per array (of which seven drives are available for data storage.) Three arrays are created using this hybrid configuration, compared to 14 arrays when using vertically striped RAID 1 or RAID 5 arrays.

Recommended System Configurations

- Storage Enclosure Configuration Options
- Recommended Configuration Methods
- When Fault Tolerance is Most Important
- When Capacity Utilization is Most Important
- When I/O Performance is Most Important

Storage Enclosure Configuration Options

The number of disk storage enclosures attached to the Smart Array 5300 or Smart Array 5i controller determines the possible configuration methods. Table 4-3 provides a list of the striping and array configuration methods available for the different hardware configurations.

Table 4-3: Suggested Storage Enclosure Configurations

Striping Method	One 4300 Family Storage Enclosure	Two 4300 Family Storage Enclosures	Four 4300 Family Storage Enclosures
Vertical (NSPOF)	None	RAID 1	RAID 1 RAID 5 RAID ADG
Horizontal	RAID 0 RAID 1 RAID 5 RAID ADG	RAID 0 RAID 5 RAID ADG	RAID 0 RAID 5 RAID ADG
NSPOF Horizontal	None	RAID 1+0	RAID 1+0

Recommended Configuration Methods

As discussed in the previous sections of this chapter, a variety of configuration methods are available to choose from when configuring the system storage of the NAS B2000. Different RAID levels and array striping methods can be combined to create many possible configurations for each NAS B2000 deployment. Administrators must choose between RAID levels and striping methods that offer different levels of fault tolerance, capacity utilization, and I/O performance.

Figure 4-5 illustrates the same System Characteristics triangle that was presented earlier. Depending on the system environment and the type of data that will be stored on the NAS device, different administrators will prioritize different desired system characteristics.

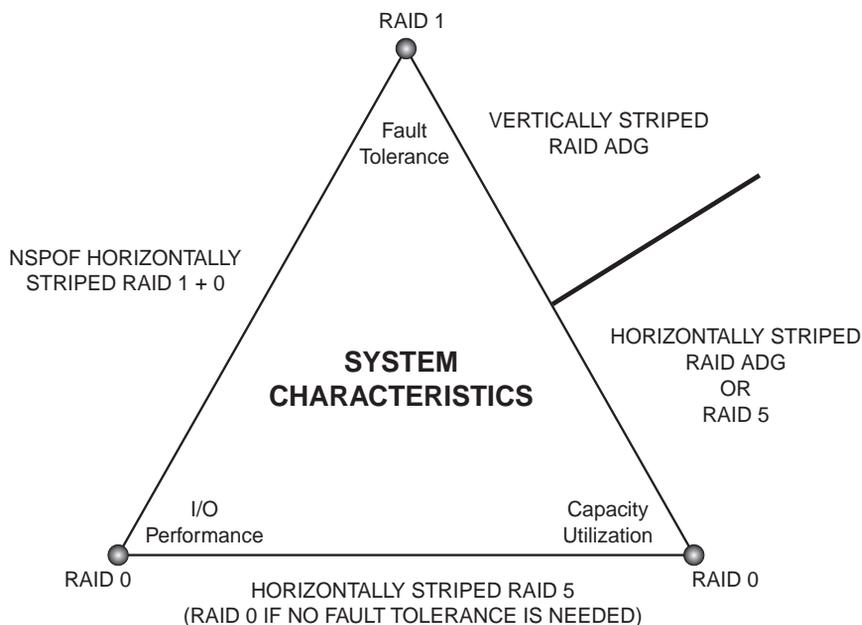


Figure 4-5: Recommended Configuration methods

The following paragraphs discuss the best configurations for the different priority rankings of these desired system characteristics.

When Fault-Tolerance is Most Important

When high availability is of paramount importance, arrays must be configured using an NSPOF configuration. In this configuration, if a storage enclosure fails or if one of the SCSI channels fails, all the arrays are still available.

As shown in Figure 4-5, if I/O performance is considered important as well as fault tolerance, the arrays should be striped using an NSPOF horizontal RAID 1+0 configuration. These mirrored drives offer the ultimate protection against data loss. Although RAID 1 and RAID 1+0 are mirrored configurations with 50 percent of raw capacity reserved for fault tolerance, RAID 1+0 incorporates more drives into an array, allowing the creation of larger LUNs than RAID 1.

In large deployments, only the most sensitive data should be stored on mirrored drives, with the remainder of the data stored on larger, more capacity-efficient arrays.

If capacity utilization is considered second to fault tolerance, the arrays should be striped vertically using RAID 5 or RAID ADG. One set of parity information is maintained for RAID 5, so one disk from each storage enclosure can be included in an array. Because RAID ADG maintains two sets of parity information, two drives from each storage enclosure can be included in an array. These fault-tolerance methods offer a high level of protection against data loss and reserve much less space for fault tolerance than RAID 1 or RAID 1+0.

NSPOF configurations offer ultimate data protection, trading off some capacity utilization and I/O performance for high availability.

When Capacity Utilization Is Most Important

When the maximum utilization of storage capacity is considered more important than fault tolerance or performance, large horizontally striped arrays should be created.

If no fault tolerance is necessary, drives in a storage enclosure can be incorporated into one or several horizontally striped RAID 0 arrays.

More commonly, some level of fault tolerance protection is needed. Therefore, the arrays should be horizontally striped RAID ADG or RAID 5 arrays. Figure 4-5 illustrates that if fault tolerance is more important than I/O performance, the arrays should be striped horizontally using RAID ADG. If I/O performance is more important, the arrays should be striped horizontally using RAID 5. These fault-tolerance methods offer efficient use of capacity, while offering appropriate levels of security and performance.

When I/O Performance Is Most Important

If I/O performance is determined to be more important than fault tolerance or capacity utilization, horizontal striping should be used.

If capacity utilization is the next most important system characteristic, then the storage can be configured into horizontally striped RAID 5 arrays. (If no level of fault tolerance is necessary, use RAID 0 striping.)

Alternatively, if fault tolerance is ranked second to I/O performance, the most effective configuration uses NSPOF horizontal RAID 1+0 arrays.

Physical Storage Planning Issues

- Hard Drive Sizes and Types
- Use and Number of Spares
- LUN Sizing

Hard Drive Sizes and Types

RAID arrays should be composed of hard drives of the same size and type. When drive types are mixed within a storage enclosure, the usable capacity and the processing ability of the entire storage subsystem is affected. Figure 4-6 provides several examples of arrays in a mixed-drive storage subsystem.

Mixed Drive Sizes

When a RAID array is composed of different-sized drives, the RAID array defaults to the smallest individual drive size, and capacity in the larger drives goes unused.

For example, array A9 in Figure 4-6 includes seven drives, with three 18.2-GB drives and four 36.4-GB drives. The controller creates the array as if it were composed of seven 18.2-GB drives. This process results in a waste of 72.8 GB (18.2 GB wasted per 36.4-GB disk x 4 disks = 72.8 GB.) The better solution is to create two different RAID arrays, one with the three 18.2-GB drives and one with the four 36.4-GB drives.

Although it is an extreme example, array A8 in Figure 4-6 illustrates another example of inefficient drive placement. In this vertical array, one drive is 18.2 GB and two drives are 72.8 GB. This configuration results in a waste of 109.2 GB (72.8 - 18.2 = 54.6 GB per disk x 2 disks = 109.2 GB.)

Note: Except for the indicated drives, all drives are 72 GB Ultra 3 drives.

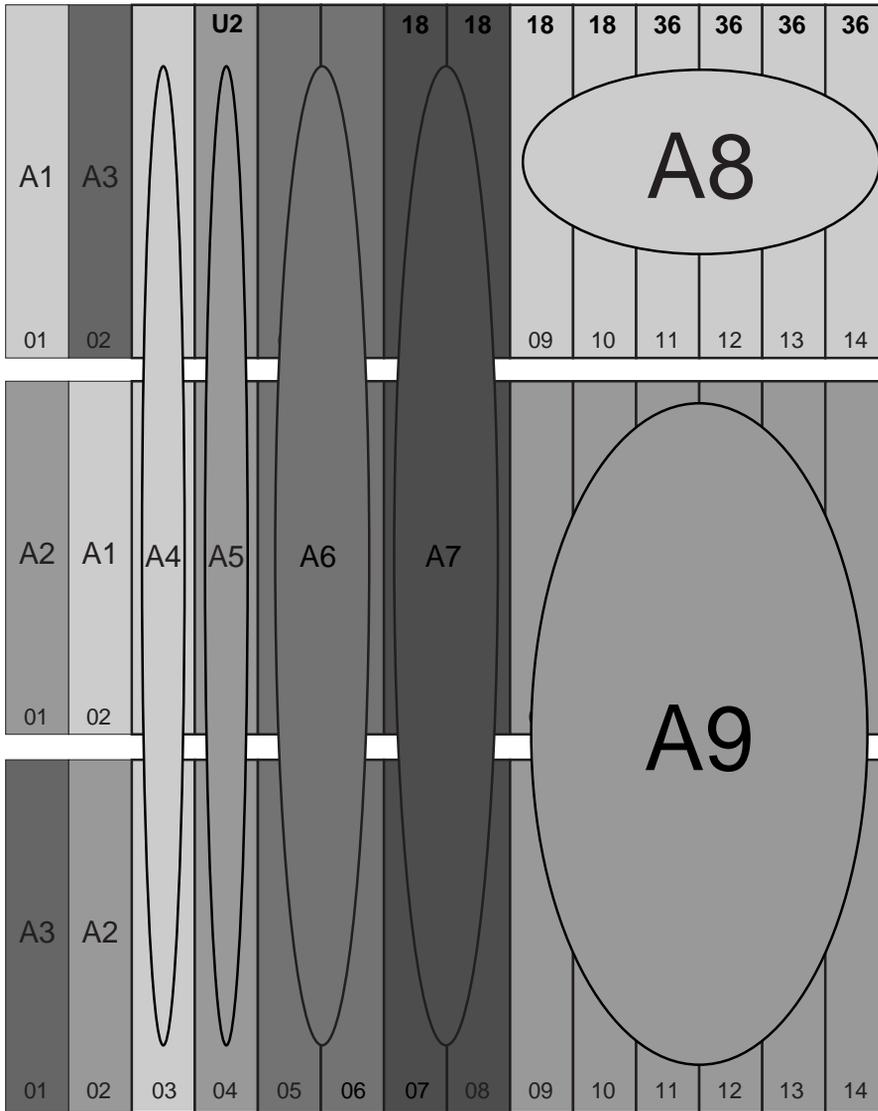


Figure 4-6: How drive characteristics affect array performance

Table 4-4: Table 4-6 Legend

Array Number	Array Configuration	Processing Characteristics	Effective Drive Capacity
A1, A3	Vertically striped RAID 1	U2	72 GB
A2	Vertically striped RAID 1	U3	72 GB
A4, A5, A6, A7	Vertically striped RAID 5	U2	72 GB
A8	Vertically striped RAID 5	U2	18 GB
A9	Horizontally striped RAID 5	U2	18 GB
A10	Horizontally striped RAID ADG	U3	72 GB

Mixed Drive Types

When different drive types are included in the same enclosure, the processing characteristics of the entire enclosure are reduced to that of the slowest drive.

In addition to the size of the drive affecting the usable capacity in an array, the processing characteristics of the drives must be considered. Do not mix different generations of hard drives (such as Ultra 2 and Ultra 3) in the same storage enclosure.

The processing characteristic of any array that includes a drive from that storage enclosure is reduced to that of the slowest drive. Using the arrays in Figure 4-6 as an example, any array containing a drive in the top storage enclosure is reduced to the processing ability of the Ultra 2 drive in drive bay 4. This includes array A1, A3, A4, A5, A6, A7, A8, and A9.

Spare Drive Sizes

The sizes of spare drives in relation to the active drives must be taken into consideration.

A RAID array composed of 36.4-GB drives cannot use an 18.2-GB spare to replace a failed drive, but a RAID array composed of 18.2-GB drives can use a 36.4-GB spare. Compaq recommends that the spare set consist of the largest drives in the entire storage subsystem. This configuration ensures that any array in the storage subsystem can use any of the spares.

The following section defines and discusses using spares.

Use and Number of Spare Disks

Compaq recommends that spare disks be designated for use on the NAS B2000. Spares are disks that are not active members of any particular array but have been configured to be used in the event that a disk in one of the arrays should fail. If a spare is present, it will immediately be used to begin rebuilding the information that was on the failed disk, using the parity information from the other member disks. During the rebuilding process, the array is operating in a reduced state and, unless it is a RAID ADG or RAID 1+0 array, it cannot tolerate another disk failure in the same array. If another disk should fail at this time, the array would become inaccessible and the information stored there would have to be restored from backup.

After the rebuild of the data onto the spare is completed, when a replacement drive is inserted to replace the failed drive, the system will automatically transfer the data from the spare onto the replacement drive and return the spare to an available-spare state. It is important to note that the process of rebuilding the spare or the replacement drive must not be interrupted, or the process will be aborted.

Some administrators deem it necessary or desirable to have multiple spare disks, so that multiple arrays can experience failure and successfully recover, before administrative intervention would be required to replace the spare or failed disk. When assigning a spare to an array, the administrator chooses which arrays and how many arrays are protected by that spare.

LUN Sizing

When planning for optimal file serving performance, you must determine the number of hard drives necessary to maintain an optimum performance level. As a general rule, the greater the number of drives that are included in an array, the greater the performance level that can be achieved. However, the performance considerations are offset by fault tolerance considerations. The greater the number of drives in an array, the higher the probability of one or more disk failures in that array. The administrator must strike a balance between performance and fault tolerance.

In addition to other reasons mentioned throughout this chapter, planning the size of the arrays and LUNs is important for the following reasons:

- LUNs cannot be extended.
- LUNs are concatenated when they are added into a pool.

These limitations lead to the recommendation that system administrators create large LUNs and place a single LUN into a pool. Each of these points is discussed in further detail in the following paragraphs.

LUNs Can Not be Extended

The array controller does not support the growing of LUNs (called logical drives in the Array Configuration Utility) after they have been created. In order to grow an array and its corresponding LUN, the LUN and the array would have to be deleted and then recreated including more drives in the new array and LUN. Doing this results in the loss of the data in the array. After the decision to recreate the arrays has been made, but prior to actually deleting and recreating the array and LUN, all data on the underlying folders and virtual disks must be backed up. This backup will be used to restore the data onto the system after the new configuration is established. The data need not be erased before deleting or recreating the arrays and LUNs, because the re-creation process deletes the data.

An alternate method of incorporating new drives in the array is to expand the array itself. This does not alter the size of the existing LUN, but it does allow the array to incorporate the new drives and gain the additional performance benefit of having more drives in the array. Additionally, the LUN and all of its data will remain intact. However, this means that the original LUN is not taking advantage of the new storage space.

To use the additional space that was added by incorporating these new drives into the array, a second LUN will have to be carved out of the new free space inside the array. This LUN will then need to be incorporated into an existing pool or used to create a new pool. This method is the most flexible with respect to adding storage on an as-needed basis. However, it comes with the same planning issues as using vertical array striping (creating a large number of LUNs out of smaller storage chunks.) Compaq recommends that no fewer than seven drives be added each time this procedure is used. This will help reduce the overhead required by the new storage, while also helping to preserve the flexibility and ability of the pools to dynamically grow as storage needs increase.

Although both methods allow for modification of existing array configurations, each has its own challenges and issues. Thus, planning for future growth and creating the arrays wisely is extremely important.

LUNs are Concatenated When They Are Added into a Pool

The VR software included with the NAS B2000 device does not stripe data across the LUNs that are incorporated into a pool. Rather, LUNs are concatenated inside the pool. This means that data is not written to subsequent LUNs until the capacity of the first LUN is fully used.

When multiple LUNs are members of a pool, performance will typically be less than if one larger LUN containing all of the hard drives was used to form the pool. All writes will go to a single LUN, until there is no more space available on the LUN. After that LUN is full, then the writes will move to the second LUN, and so on. This performance will typically be less than if a single LUN comprises the pool and all drives are incorporated initially. In configurations where performance is the paramount concern, Compaq recommends importing only one large LUN into a pool, rather than composing the pool of multiple smaller LUNs.

Virtual Storage Planning Issues

The rules associated with the Virtual Replicator (VR) software must also be taken into consideration when planning system storage size and capabilities. These issues directly affect the planning and configuration of the storage. Knowing and understanding these rules is crucial to properly configuring the storage.

Virtual storage planning issues include:

- Pool Considerations
- Virtual Disk Considerations
- Share Considerations

Pool Considerations

The primary consideration that must be addressed when planning VR pools is that there is a limit to the amount of growth a single pool can undergo. Pools are limited in size and are limited to no more than eight separate LUNs. VR pools cannot exceed 8 terabytes of disk space. When using vertically striped RAID 1 arrays, 21 LUNs will be created. When using vertically striped RAID 5 arrays, 14 LUNs will be created. A single pool cannot span all of these arrays. Because a pool can contain no more than eight LUNs, these vertical configurations require the administrator to divide the storage into at least two or three separate pools.

The RAID level of the LUNs included in a pool must be considered. All the units that make up a pool should have the same high-availability characteristics. In other words, the units should all be of the same RAID level. For example, it would be a bad practice to include both a RAID 1 and a RAID 5 array in the same pool. By keeping all the units the same, the entire pool retains the same performance and high-availability characteristics, making managing and maintaining the pool much easier.

Finally, the administrator must carefully consider how the pools will be carved up and what groups or applications will be using them. For example, putting several storage-intensive applications or groups into the same pool would not be efficient. These applications or groups would be better served by being divided up into separate pools, which could then grow as their space requirements increased, following the allowable growth limits.

Virtual Disk Considerations

The primary considerations when planning VR virtual disk carving are that virtual disks are limited in size and that there is a limit to the number of virtual disks that can be carved out of a VR pool. Currently, a virtual disk cannot be larger than 2 terabytes and no more than eight virtual disks can be carved out of a single pool. This means that in larger configurations, storage must be divided up and organized into these 2-terabyte chunks. For example, if a group has 20 different departments, it is not feasible to create an 8-terabyte pool and then carve out a separate virtual disk for each department. Rather, the storage would have to be broken down into several separate pools, and then the virtual disks for the departments would be carved out of those pools. (As mentioned later in this chapter, an alternative method of dividing out storage space among departments is to create file shares on the virtual disk for each department.)

Unless mount-points are being used, in order for users and the operating system to be able to access the virtual disk, each virtual disk must be presented as a drive letter on the device. Although it is possible to create a virtual disk and not map it to a drive letter, neither the users nor the administrator will have access to the storage space contained within that virtual disk. After a drive letter is mapped, the space becomes available for use. This is also true of snapshots. A snapshot can be made but is not accessible until it is assigned a drive letter.

A finite number of drive letters can be assigned. There is theoretically a maximum of 26 drive letters available, of which several are already reserved. For example, the diskette drive is A:, and the boot disk is C:. Because snapshots are competing with virtual disks for available drive letters, the administrator must carefully plan the number of virtual disks to create and the number of snapshots to keep. After all the available drive letters are used up, there is no mechanism to make additional virtual disks or snapshots available.

Share Considerations

Planning the content, size, and distribution of shares on the NAS B2000 can improve performance, manageability, and ease of use.

The content of shares should be carefully chosen to avoid two common pitfalls: either having too many shares of a very specific nature or of having very few shares of a generic nature. For example, shares for general usage are easier to set up in the beginning, but can cause problems later. Frequently, a better approach is to create separate shares with a specific purpose or group of users in mind. However, creating too many shares also has its drawbacks. Take care to avoid creating shares unnecessarily. For example, if it is sufficient to create a single share for user home directories, create a “homes” share rather than creating separate shares for each user.

By keeping the number of shares and other resources low, the performance of the NAS B2000 is optimized. For example, instead of sharing out each individual user’s home directory as its own share, share out the top-level directory and let the users map personal drives to their own subdirectory.

Storage Sizing Considerations

Estimating the number and configuration of storage enclosures, disks, arrays, pools, and virtual disks that are needed in a particular environment can be a complicated endeavor. A variety of factors affect the actual usable amount of space on a particular NAS B2000, including:

- RAID Issues
- Spare Disk Issues
- Snapshot Issues
- Growth Issues
- Allocation Unit Size Issues
- Consolidation Concerns

RAID Issues

The RAID level defines the high-availability characteristics of the RAID array. The variety of RAID types offers different combinations of fault tolerance, performance, and efficiency.

NOTE: For a complete description of RAID types, see Chapter 3. Additional information on configuring the RAID types is included in the “Fundamental Storage Configuration Planning Issues” section earlier in this chapter.

Unless drives are being mirrored using RAID 1, no fewer than three disks should be included in an array, so that more efficient RAID types can be used. As more disks are used in a RAID array, the percentage of total space devoted or “lost” to parity for fault tolerance goes down. Thus, in general, larger RAID arrays are more space efficient. For example, in a three-disk RAID 5 array, the equivalent of one disk of the three is devoted to parity, accounting for about 33 percent of the total raw storage capacity. In a 14-disk RAID 5 array, the equivalent of a single disk is still needed for parity, but the percentage drops to one in 14 disks, or about 7 percent. Larger arrays offer better performance and use of storage capacity than smaller arrays.

Adequate planning must increase the amount of raw storage capacity needed by the estimated amount of storage space used for fault tolerance.

Spare Disk Issues

Compaq recommends that spare disks be designated for use on the NAS B2000. Spares are disks that are not active members of any particular array, but have been configured to be used in the unlikely event that a disk in one of the arrays fails.

Some administrators deem it necessary or desirable to have multiple spare disks, so that multiple arrays can experience failure and successfully recover, before administrative intervention would be required to replace the spare or failed disk. When assigning a spare to an array, the administrator chooses which arrays and how many arrays are protected by that spare.

The administrator must include in the storage planning process the increase of raw storage capacity needed to account for the hard drives that will be used as spares.

Snapshot Issues

Snapshots are a feature of VR and are a point-in-time view of a VR virtual disk. Even though snapshots do not use up any space initially, maintaining the point-in-time view of the virtual disk requires that snapshots must allocate space and create a copy of original data before allowing any changes to be made to the parent virtual disk. Thus, snapshot space usage tends to climb over time, and the rate of usage is related to the rate of change of the data on a virtual disk, the total lifetime of the snapshot, and the number of snapshots maintained for a virtual disk.

By default, the NAS B2000 reserves 10 percent of the pool space exclusively for the use of snapshots. This percentage may be altered for those expecting to generate more or less snapshot data.

The administrator must include in the storage planning the increase in raw storage capacity needed by either 10 percent or the estimated amount of space that will be reserved for snapshot use.

Growth Issues

It is also important to allow for significant changes in the total size of all files on a particular virtual disk. Some applications generate large temporary files during their execution, or create extremely detailed log files that can cause a significant, but temporary, expansion in needed disk space. Knowing the type of data that will occupy a particular share or disk is valuable in planning for this type of variability.

Planning for growth can seriously alter sizing requirements. Storage needs commonly grow from 50 percent to over 500 percent in a single year, depending on the system deployment. In general, it is better to accommodate an expected rate of growth from the beginning rather than attempting to address it in several smaller changes throughout the year.

When necessary, an additional storage enclosure or groups of storage enclosures can be added to the original configuration. After new arrays and LUNs are created, they can be formed into additional pools and virtual disks, mounted on drive letters, and put into use.

Allocation Unit Size Issues

Depending on the allocation unit size, a given amount of data from one server may take up either more or less disk space when moved to the NAS B2000.

Allocation units define the smallest segment of a disk that is read or written at a time. If a file size is smaller than the allocation unit, the extra space is “wasted” because it is not used. If there is a large number of files that are smaller than the allocation unit size, a large amount of unused space may exist in those files. In contrast to the storage space considerations, larger allocation unit sizes offer better disk I/O and file-serving performance. If the allocation unit size is too small relative to the size of the files, more disk I/Os are required to read and write the data, and consequently, both the disk I/O and the file-serving performance degrades. To achieve a good balance between performance and space efficiency, the default allocation unit size on the NAS B2000 is 16 KB.

When migrating data from an existing server to the NAS B2000, the administrator must make sure that the allocation unit size on the NAS B2000 virtual disk exactly matches that of the original virtual disk, unless this modification is intentionally being performed and the impact on required space has been planned. If the administrator does not carefully compare and plan for the allocation unit size, the performance and storage space can be negatively impacted. Depending on the allocation unit size of the destination virtual disk, the allocation unit size of the source virtual disk, and the actual file sizes, the destination disk may need to be smaller than or larger than the source disk. If the raw capacity of the destination virtual disk is not larger than the capacity of the source disk, and during the migration process the difference in allocation unit size causes the storage used to grow in relation to the source disk, the destination disk may fill up before all of the data is migrated. The allocation unit size of the virtual disks is easily selectable when creating virtual disks on the NAS B2000.

Consolidation Issues

Consolidating the data from several file servers into a single NAS B2000 can improve availability and decrease the administrative burden of fileserver management. However, care should be taken when estimating the amount of space that will be required to replace existing servers. In most cases, it is not merely a matter of adding all the space used by the replaced servers. All of the previously mentioned factors must be taken into account.

Storage Management Planning Scenarios

This section provides examples of system configurations. The first scenario details the analysis and planning process of determining the configuration of an example system deployment. Other examples compare different possible configurations of the same amount of system storage.

Additionally, a planning worksheet is included at the end of this section. System administrators can use this worksheet to guide them through the storage planning process.

In summary, this section includes:

- A complete and detailed storage planning example
- A simple sizing comparison
- An example of a storage subsystem using different array configurations
- A planning worksheet

A Complete and Detailed Storage Planning Example

This example progresses through the analysis and planning process of setting up and configuring the storage for a NAS B2000 device. The assumptions and numbers used in this example are also illustrated in a completed Planning Worksheet located at the end of this section.

Assume that a NAS B2000 device is obtained to consolidate file storage. Before any migration of data can occur and before any configuration of the storage on the NAS device can take place, complete analysis of the current file-serving environment must be completed. Primary analysis of the current file-serving environment involves the determination of amount of space that is currently used to store the data that will be migrated.

Initial Storage Needs

Assume that the following information was obtained during the analysis process:

- 1000 GB of space is needed to accommodate data that will be migrated.
- 500 GB of additional space for data is needed to accommodate a projected 50 percent future growth in storage needs.
- Therefore, the initial usable storage need is 1500 GB.

See Table 4-5 for a sample of the initial entries to the Planning Worksheet.

Snapshot Storage Needs

Snapshots may be a new feature to many environments that are moving to the NAS B2000. As mentioned previously in Chapter 3, a snapshot is an instantaneous copy of pointers to data on the disk. When the data on the disk changes, the original blocks are copied to a reserved snapshot area of the same pool in which the parent virtual disk is located. In environments where a large volume of data changes rapidly, it is possible for snapshots to use a significant amount of storage capacity.

- Determine the percentage of space to reserve for snapshots.

When snapshots are going to be used, proper planning includes an allowance for an adequate amount of space to be reserved for snapshots. For most environments, Compaq recommends that an additional 10 percent of disk space beyond the storage requirements be allocated for snapshots.

- Determine the percentage of pool space that will be reserved for snapshots, such as 10 percent.
- Convert the percentage to a decimal equivalent, such as 0.10.

- Manipulate this percentage to derive a factor that can be applied to the Initial Usable Storage Need, resulting in a Total Storage Need.

The following formula ensures that the required space will be available after the snapshot reserve is deducted:

- The Snapshot Factor is 1.00 minus the reserve percentage, such as $1.00 - 0.10 = 0.90$.

This scenario will use snapshots and will use the default snapshot reserve of 10 percent.

Total Storage Need

By applying the Snapshot Factor to the Initial Usable Storage Need, the Total Storage Need can be calculated. The result is the amount of required storage space. This figure can then be used as a constant during other steps of the planning process, such as determining the best array size and configuration.

If the Initial Usable Storage Need is divided by the Snapshot Factor, the needed space is increased by a sufficient amount to allow for the snapshot reserve.

For example, this Total Storage Need is $1500 \text{ GB} / .90 = 2142.86 \text{ GB}$, or 2143 GB.

Table 4-5: Example Storage Need Worksheet

	Formula	Value
Initial Storage Need		
1. Data Space Needed		1000 GB
2. Future Growth Space		500 GB
3. Total Initial Usable Storage Need	Data Space + Growth Space (Step 1 + Step 2)	$1000 + 500 = 1500 \text{ GB}$
	Initial Storage Need	1500 GB
Snapshot Storage Need		
4. Reserve Percentage		$10\% = 0.10$
5. Snapshot Factor	$1.00 - \text{Reserve Percentage}$ $1.00 - \text{Step 4}$	$1.00 - 0.10 = 0.90$
Revised Storage Need		
6. Total Storage Need	Total Initial Usable Storage Need / Snapshot Factor (Step 3 / Step 5)	$1500 / .90 = 1667 \text{ GB}$
	Total Storage Need	1667 GB

Array Configuration Requirements

Before the physical hard drives can be configured into arrays and LUNs, the preliminary planning steps as outlined in the beginning of this chapter must be completed. These primary planning decisions include:

NOTE: See Table 4-6 for an example of a completed Planning Worksheet for this scenario.

- Determine the sizes and types of physical hard drives that will be used.

The NAS B2000 supports the use of 1-inch Compaq Ultra 2 and Ultra 3 drives in the following capacities: 18 GB, 36.4 GB, and 72.8 GB.

In this scenario, Compaq Ultra 3 72.8-GB drives were chosen instead of the 36.4 GB drives to maximize the use of the storage enclosure and its available drive bays.

- Determine which array configuration strategy will be used, how many drives will be reserved for parity, and how many drives will be included in each array.
 - The NAS B2000 supports the use of several different RAID configurations of the arrays and LUNs. In conjunction with striping strategies, these RAID configurations offer varying combinations of fault tolerance, I/O performance, and capacity utilization.

This scenario assumes that while fault tolerance is important, capacity utilization and I/O performance are important as well. For this reason, horizontally striped RAID 5 arrays will be used.
 - After the RAID level has been selected, the amount of space required for fault tolerance per array can be determined.

In this scenario using horizontally striped RAID 5 arrays, one drive from each array must be reserved for parity information.

- As in previous discussions in this chapter, having more drives translates into having more space, and more drives per array means that the percentage of space required for parity is reduced. However, the effective reliability of the underlying disk system is lessened, as more and more drives are included in any individual array. Using moderately sized RAID 5 arrays of 14 or fewer drives, it is unlikely that two drives within the same array would fail at the same time. While fourteen drives is still a reasonable RAID 5 array size, increasing the size to twenty-eight or forty-two drives steadily increases the potential for multiple-drive failure within the same array. It is for this reason that Compaq recommends placing 14 or fewer drives in a RAID 5 array.

An additional point to consider when deciding on the array size is the potential restore window for data recovery from offline media in case of a disk subsystem failure. Fourteen drives'-worth of data may be restorable in a few hours, and users of data resident on other arrays can continue using that data while the restore takes place. However, the time required to restore arrays containing twenty-eight or forty-two drives is much greater. Since the failure of a larger array takes with it a greater chunk of the potential storage capacity of the NAS B2000, fewer arrays are available for use while the restore takes place.

Based on these considerations, this scenario will create seven-drive RAID 5 arrays.

- Determine how many drives in the array will be usable for storage.

This figure is easily obtained using the decision from the previous steps. With the RAID configuration, array size, and fault tolerance requirement known, the following equation can be formed:

Usable Drives per Array = Number of Drives per Array – Number of Drives for Fault Tolerance

In this scenario, the numbers are:

7 drives per array – 1 drive for fault tolerance = 6 usable drives per array

- Determine the amount of usable space in the array.

As these calculations progress, it can be seen that these figures are building upon one another to ultimately determine the total number of arrays, the number of drives that must be purchased, and the number of storage enclosures needed to hold all of the drives. Therefore:

Usable Space per Array = Usable Drives per Array x Individual Drive Size

In this scenario, the numbers are:

6 Usable Drives x 72.8-GB drives = 436.8 GB usable space per array

- Determine the total number of arrays required.

One of the most important steps in the storage planning process is the determination of how many arrays must be created to house the needed storage. Knowing the required number of arrays leads to a conclusion of how many drives and storage enclosures must be purchased to house the storage.

Additionally, because the arrays and LUNs are used to build the pools, the number of arrays in a system directly affects the size of the storage pools and virtual disks that are created from the arrays.

The formula for determining the required number of arrays is straightforward: divide the total amount of needed usable storage by the amount of usable storage in each array. The Total Storage Need can be determined and obtained from the Usable Storage Needs Worksheet.

Total Number of Arrays Required = Total Storage Need / Usable Space per Array

In this scenario, the numbers are:

1667 GB Total Storage Need / 436 GB Usable Space per Array = 3.82
Total Number of Arrays Required = 4 arrays required

Table 4-6: Example Array Configuration Requirements Worksheet

	Formula	Value
Array Configuration Requirements		
1. Individual Drive size		72.8 GB
2. Number of Drives per Array		7 drives
3. Number of Drives for Fault Tolerance per Array		1 drive
4. Usable Drives per Array	Number of Drives per Array – Number of Drives for Fault Tolerance (Step 2 - Step 3)	6 drives
5. Usable Storage Space per Array	Usable Drives per Array x Individual Drive Size (Step 4 x Step 1)	6 x 72.8 GB = 436.8 GB
6. Total Number of Arrays Required	Total Storage Need / Usable Storage Space per Array (Step 6 from previous worksheet / Step 5 from this worksheet)	1667 GB / 436 = 3.82
Total Number of Arrays Required		4 arrays

Drives Required

At this point in the planning process, one potential configuration strategy has been presented, and the number of arrays required to support the needed amount of usable storage has been calculated.

To continue with the planning process, the required number of drives must be determined, as well as the number of storage enclosures needed to house them.

See Table 4-7 for an example of a completed worksheet.

- Determine the number of hard drives required for all of the arrays.

This figure is an interim figure, one that when increased by the number of spare drives to be used results in the total number of drives needed for the entire storage subsystem of the NAS B2000.

Number of Drives Required for all Arrays = Total Number of Arrays Required x
Number of Drives per Array

For this scenario:

4 Arrays x 7 Drives per Array = 28 Drives Required for all Arrays

- Determine the number of spare drives that will be used.

Compaq recommends that at least one online spare be allocated for each storage subsystem. One spare can be dedicated to each array or a single spare can be shared with multiple arrays. These extra drives must be added to the Number of Drives Needed for an Array to show the Total Number of Required Drives.

This example will use a single spare, shared between the five seven-drive RAID 5 arrays.

- Determine the total number of drives required

Total Number of Drives Required = Drives Required for all Arrays + Spare
Drives

Therefore, at this time, 36 drives must be purchased, seven each for the five arrays, plus an additional drive to be used as an online spare.

28 Drives Required for the Arrays + 1 Spare Drive = 29 Total Drives Required

Table 4-7: Example Drives Required Worksheet

Drives Required for the Arrays			
1.	Number of Drives Required for the Arrays	Total Number of Required Arrays x Number of Drives per Array (Step 6 x Step 2)	4 x 7 = 28 drives
Spare Drives Need			
2.	Number of Spare Drives		1 drive
Drives Required for the Storage Subsystem			
3.	Total Number of Drives Required	Number of Drives for the Arrays + Number of Spare Drives	28 + 1 = 29 drives
Total Drives Required			29 drives

Storage Enclosures Required

After the number of required drives is calculated, the number of storage enclosures that are required to house them can be determined. Depending on the results, the array configuration may need to be adjusted. For example, the number of drives may expand the configuration into an additional storage enclosure by only one or two drives. In this case, different array configurations should be studied that may meet the storage needs, while not requiring an additional enclosure.

- Determine how many disk storage enclosures are required.

Because each disk storage enclosure accommodates fourteen drives, calculate this figure by dividing the Total Number of Drives Required by a constant of 14.

$$\text{Number of Storage Enclosures} = \text{Total Number of Drives Required} / 14$$

For this scenario, 3 storage enclosures are required:

29 Total Number of Drives Required / 14 = 2.07 = 3 storage enclosures

Table 4-8: Example Enclosures Required Worksheet

Storage Enclosures		
1. Total Number of Storage Enclosures Required	Total Number of Drives / 14	36 / 14 = 2.57 enclosures
Total Number of Storage Enclosures Needed		3 enclosures
MSA1000s		
2. Total Number of MSA1000s Required	Total Number of Storage Enclosures / 3	3 / 3 = 1 MSA1000

Conclusion

A final look at the figures in this scenario shows the following:

- 36 hard drives must be obtained and placed into the storage enclosures.
Because the storage subsystem can house up to 42 drives, six open drive bays are available for future drive purchases.
- Five seven-drive horizontal RAID 5 arrays will be created.
- One spare will be assigned to serve all of the arrays.

At this time, one possible storage configuration has been presented. This scenario was planned to provide a conclusion that needs no adjustments.

For most deployments, this analysis and configuration process must be repeated several times, each time using different variables.

If all data is considered equal and will be placed in identical arrays, the Total Usable Storage Need will remain constant. However, when differing RAID configurations and array sizes are suggested, the resulting storage configurations will be quite different.

If some data needs special protection, the Total Usable Storage Need may need to be divided into different groups. A combination of configurations can be created in order to protect the most sensitive data in an NSPOF configuration, while other data may be stored in more capacity-efficient configurations.

A Simple Sizing Comparison

In terms of raw storage capacity, four fully populated storage enclosures using fifty-six 72.8 GB drives provides about 4076 GB or 4 TB of raw storage space. This compares to approximately 2038 GB or 2 TB of raw storage space if 36.4-GB drives are used. Twice the amount of space per drive translates to twice the amount of total capacity in the storage subsystem available for carving into arrays and virtual disks.

An Example of a Storage Subsystem Using Different Array Configurations

Assume that four populated storage enclosures using fifty-six 72.8-GB drives are available. The examples in the following paragraphs illustrate how different configuration strategies provide different levels of fault tolerance and deliver slightly different total capacities.

If fault tolerance is of paramount importance when configuring the system, the arrays should be configured using an NSPOF horizontal RAID 1+0 array configuration. Four arrays will be created, each using seven drives from one storage enclosure and seven drives from a different enclosure to serve as the mirrored drives. Because RAID 1 and RAID 1+0 reserve 50 percent of raw storage capacity for the mirrored fault tolerance, the usable storage space is 2038 GB.

(14 drives – 7 drives for mirroring = 7 drives; 7 drives x 72.8 GB = 509.6 GB per array)
(4 arrays x 509.6 GB per array = 2038.4 GB)

If fault tolerance is still important, but better capacity utilization is desired, eight horizontal seven-drive RAID 5 arrays can be created. One drive per array will be used for parity information, resulting in 3494 GB of total usable storage space.

(7 drives – 1 drive for parity = 6 drives; 6 drives x 72.8 GB = 436.8 GB)
(8 arrays x 436.8 GB = 3494.4 GB)

If even better capacity utilization is needed, 4 fourteen-drive RAID 5 arrays can be created. One drive per array is still used for parity information, but because the arrays incorporate more drives than the previous example, only 4 drives in total will be used for parity. The resulting usable space is 3857 GB of total usable space. Simply creating larger arrays has resulted in 363 GB of additional usable storage space.

(14 drives – 1 drive for parity = 13 drives; 13 drives x 72.8 GB = 946.4 GB)
(4 arrays x 964.4 GB = 3857.6 GB)

A final example offers the ultimate capacity utilization while still offering an acceptable level of fault tolerance. All 56 drives of the storage subsystem are placed into one large RAID ADG array. Regardless of the number of drives in the array, RAID ADG arrays reserve 2 drives of capacity for parity information. So this example configuration presents 3931 GB of usable storage space.

(42 drives – 2 drives for parity = 40 drives; 40 drives x 72.8 GB = 3931 GB)

Table 4-9 consolidates these calculations.

Table 4-9: Example Usable Space Using Different Configurations

Array Configuration	Usable Space	Fault Tolerance
RAID 1 arrays	2038 GB	28 drives - 50%
7-drive RAID 5 arrays	3494 GB	8 drives - 14%
14-drive RAID 5 arrays	3857 GB	4 drives - 7%
42-drive RAID ADG arrays	3931 GB	2 drives - 5%

NOTE: This example highlights usable capacity.

Planning Worksheet

In general, the following steps should be performed to effectively plan the needed storage capacity and its configuration:

- Analyze current data storage demands and determine the amount of data that will be migrated to the NAS B2000.
- Determine the amount of space to allow for future growth.
- Determine whether snapshots will be used, and if so, if a 10 percent reserve is adequate.
- Determine the total amount of space needed.
- Determine the sizes and types of hard drives that will be used.
- Determine the ranking of desired system characteristics and decide upon the RAID striping and configuration method, space for fault tolerance, and the size of the arrays.
- Determine the number of arrays that will be required.
- Determine the number of drives required to build these arrays.
- Determine the number of spare drives to use to support these arrays.
- Determine the number of drives required to build this system configuration.
- Determine the number of storage enclosures necessary to hold all of the drives.
- Adjust, rework, and finalize and the sizing and configuration plan.

Use the following worksheet as a guide when evaluating different storage configurations.

Table 4-10: Usable Storage Need Worksheet

	Formula	Value
Initial Storage Need		
1. Data Space Needed		
2. Future Growth Space		
3. Total Initial Usable Storage Need	Data Space + Growth Space (Step 1 + Step 2)	
Snapshot Storage Need		
4. Reserve Percentage		
5. Usable Storage Percentage	1.00 – Reserve Percentage 1.00-Step 4	
Revised Storage Need		
6. Total Storage Need	Total Initial Usable Storage Need / Usable Storage Percentage (Step 3 / Step 5)	
Total Storage Need		

Table 4-11: Array Configuration Storage Needs Worksheet

	Formula	Value
Array Configuration Requirements		
1.	Individual Drive size	
2.	Number of Drives per Array	
3.	Number of Drives for Fault Tolerance per Array	
4.	Usable Drives per Array	Number of Drives per Array – Number of Drives for Fault Tolerance (Step 2 - Step 3)
5.	Usable Storage Space per Array	Usable Drives per Array x Individual Drive Size (Step 4 x Step 1)
6.	Total Number of Arrays Required	Total Storage Need / Usable Storage Space per Array (Step 6 from previous worksheet / Step 5 from this worksheet)
Total Number of Arrays Required		

Table 4-12: Drive and Enclosure Requirements

Drives Required for the Arrays	
1. Number of Drives Required for the Arrays	Total Number of Required Arrays x Number of Drives per Array (Step 6 x Step 2)

Spare Drives Need	
2. Number of Spare Drives	

Drives Required for the Storage Subsystem	
3. Total Number of Drives Required	Number of Drives for the Arrays + Number of Spare Drives

Total Number of Drives Required	

Storage Enclosures	
4. Total Number of Storage Enclosures Required	Total Number of Drives / 14

Migration Issues

The process of moving data from the old file servers over to the NAS B2000 includes the following steps:

- Developing a migration plan
- Performing the migration

The following sections discuss the conceptual aspects and procedural highlights of different migration methods.

Developing a Migration Plan

There are only two ways to transition from one file server to another. Each method requires some downtime, but the amount of downtime varies, as does the impact on the client users. These migration methods are:

- System-wide migration
- Departmental migration

System-Wide Migration

During a system-wide migration, the entire system is taken offline. During this downtime, the entire content of the old system is migrated over to the NAS B2000. All migration procedures are performed once, perhaps over a weekend.

In large deployments, this method may not be practical. If the amount of data to transfer is great, a system-wide migration may not be possible, due to the increased amount of time necessary to move all data.

Departmental Migration

During a departmentalized migration, sections of the operation are individually taken offline, migrated over, and brought online. This method is also referred to as a “rolling” migration. The migration steps are performed several times, once for each department. During a departmental migration, one department is brought over at a time according to a set schedule, such as each Sunday night or each weekend.

In addition to requiring a relatively small window of time, there are additional advantages of a departmental migration. A departmental migration allows for a “trial run” of a portion of the business to go live on the new system. During this trial run, the administrator can address any questions or concerns about the migration process.

Departmental migration is accomplished by first identifying the level at which data will be transitioned as a unit. Transitioning data to the NAS server in more granular logical units has the advantage of allowing the administrator to carefully analyze and plan an appropriate storage management scheme, rather than simply moving the data. Regardless of the name, the departmental migration method may be applied at many distinct levels, as outlined below:

- **Server level:** Much like a system-wide migration, data from each particular file server is transitioned “en masse,” when one is performing the consolidation of many individual file servers into a single NAS B2000.
- **Volume level:** In many cases, the volume is the unit of transition, since a volume is typically the unit at which a particular group of users is granted storage. In the case of the NAS B2000, a volume is called a virtual disk.
- **Project Directory level:** Some companies choose to partition their storage into very large volumes for convenience, and then subdivide the large volumes into separate project directories, which may then be allocated to project groups as needed.
- **Share level:** Frequently the most effective level of movement is at the share level, since this provides very granular control over the users and groups who may be affected, and the share point of the old server is available to the new NAS server as a source for copying the data.

Performing the Migration

As with the migration plans, there are two methods of actually moving the data from the old file server to the NAS B2000:

- Backup and restore
- Ethernet copy

IMPORTANT: Regardless of the method chosen, backing up and verifying the entire content of the server whose data is being moved is very important. There is always the possibility that mishaps can occur during the data transition. Even if not all data on the old server is being migrated to the NAS device, having a complete, verified backup of the old server's data is insurance against accidental deletion of directories and missed content.

Although the actual migration procedures of these two methods are different, some of the preparatory and completion procedures of each method are the same. These procedures for each of these migration methods are outlined below.

Backup and Restore

The following steps provide an outline of the necessary procedures when migrating to the NAS B2000 using backups. Procedural details are located within the appropriate topic-specific chapters of this administration guide.

1. Complete the initial system configuration.
Complete the Rapid Startup procedures.
2. Configure the NAS B2000 storage, *excluding* file shares.
Configure the arrays, LUNs, pools, and virtual disks.

IMPORTANT: When creating the virtual disks, verify that the Allocation Unit Size is set to the desired size. If the allocation unit size is not the same on the source volumes and the target virtual disks on the NAS device, there may be unintended growth in the file sizes. See the "Allocation Unit Size Issues" section earlier in this chapter for a discussion on the issues surrounding selecting an appropriate allocation unit size.

Do not create file shares at this time. The file structure is contained in the backup of the data and will be restored along with the data during the restoration process.

3. Create any local users or groups.

4. Install the backup software onto the NAS device.
5. Create the backups of the data being migrated.
6. Restore the data from the backup onto the target virtual disks of the NAS device.
7. Finalize the configuration of the NAS B2000.
 - a. Re-create the file shares.
 - b. Create any desired snapshot schedules.
 - c. Build the permissions lists for the shares and files.
 - d. Add any trust relationships. These trust relationships allow users from one domain to access resources in another domain.
 - e. Install additional software, such as antivirus programs.

Ethernet Copy

If both the old file servers and the NAS B2000 can be attached to the same Ethernet network, data can be copied directly from the old servers on to the NAS device, without using backups. However, while this may be a convenient option, the actual copy process is slower than the restoration process from tape.

The procedural details of this type of migration are the same as when using tapes, with the following exception: when the storage of the NAS B2000 is configured, the file shares must also be created.

1. Complete the initial system configuration.
2. Configure the NAS B2000 storage—including file shares.
3. Create any local users or groups.
4. Create a complete system backup of the old file server.
5. Copy the data from the old file servers to the target folders on the NAS B2000.
6. Finalize the configuration of the NAS B2000.

Storage Capacity Expansion Issues

The NAS device is designed for expansion. Up to two additional external disk storage enclosures can be added to the MSA1000, for a total of 42 drives. Up to eight additional storage enclosures can be added, for a total of 56 drives. Depending on the needs of each system deployment, many configuration options are available, including:

- Add new drives to an existing array and create a new logical drive or drives
- Add new drives, create a new array, and create a new logical drive or drives
- Add the new logical drive or drives to existing pool or pools
- Create a new pool or pools using the new logical drive or drives
- Use new pool space to expand existing virtual disks
- Use new pool space or new pools to create new virtual disks

These options give the administrator great flexibility in adding new storage space or effectively expanding existing space. When adding space to an existing pool, the virtual disks created in that pool can be expanded, or new virtual disks can be created to take advantage of the space. Use the procedures described in the previous sections for creating arrays, logical drives, pools, and virtual disks.

When adding space to the system, consider the following:

- Add drives in groups of three or four. Because new drives must be configured as new logical drives, for best performance add drives in sets of at least three, and preferably seven or more.
- The maximum number of storage units (logical drives) in a pool is eight. To have large pools from which to create virtual disks, make logical drives as large as possible.

To expand system capacity:

1. Physically install the new hard drives.
2. Use the ACU to add the new hard drives to an existing array or create a new array using the new drives.

Existing logical drives automatically expand across the physical drives, including newly added ones.

3. Create a new logical drive to use the extra space in the expanded array.
4. Add the new LUN to an existing pool or create a new pool with the new LUN.
5. Create new virtual disks, as needed.
6. Create the necessary folders and file shares.

Physical Storage Management

For discussion purposes, storage in the *StorageWorks* NAS B2000 is divided into two categories: physical and virtual. Physical storage includes the physical hard drives, as well as the initial configuration of the hard drives into arrays and logical units (LUNs). Virtual storage includes the management and grouping of the LUNs into storage pools and virtual disks.

Complete conceptual information on managing physical and virtual storage is discussed in the “Storage Management Overview” and “Storage Management Planning” chapters.

This chapter discusses the procedural aspects of managing the physical storage. Procedural information on virtual storage is included in the “Virtual Storage Management” chapter.

Therefore, the storage topics discussed in this chapter include:

- Hard Drive Management
 - Adding New Hard Drives
 - Defining Hard Drive LED Indicators
 - Replacing Failed Hard Drives
 - Using Spare Drives
 - Moving Hard Drives
 - Moving Arrays

- Array and LUN Management
 - Compaq ACU Overview
 - Accessing the ACU
 - Entering Controller Settings
 - Creating a New Array
 - Creating a New Logical Drive
 - Expanding the Capacity of an Array
 - Migrating an Existing LUN to a New RAID Level

Hard Drive Management

Some of the administrative tasks for managing the physical hard drives include:

- Adding New Hard Drives
- Defining Hard Drive LED Indicators
- Replacing Failed Hard Drives
- Using Spare Drives
- Moving Hard Drives

Defining Hard Drive LED Indicators

The hard drive LEDs, located on each physical drive, are visible on the front of the server or on the front of the storage enclosure. They provide activity, online, and fault status for each corresponding drive when configured as a part of an array and connected to a powered-up controller. LED behavior can vary depending on the status of other drives in the array.

This section provides the following information about hard drive LEDs:

- An illustration detailing the location of each LED
- A table of the possible LED configurations and what each combination means



CAUTION: Read “Hot-Plug Drive Replacement Guidelines” in the *Compaq Servers Troubleshooting Guide* before removing a hard drive.

For additional information on troubleshooting hard drive problems, refer to “Hard Drive Problems” and “SCSI Device Problems” in the *Compaq Servers Troubleshooting Guide*.

Use the following illustration in conjunction with Table 5-1 to analyze the status of hot-plug hard drives.

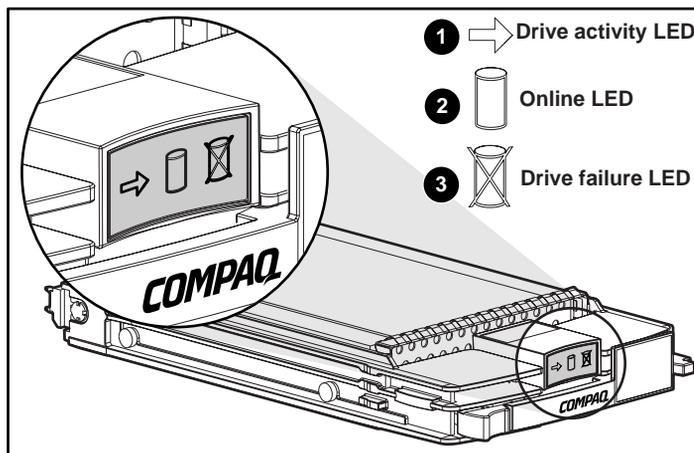


Figure 5-1: Hot-plug hard drive LED indicators

Table 5-1: Hard Drive LED Combinations

Activity	Online	Drive Failure	Indication
On	Off	Off	<p>Do not remove the drive. Removing a drive during this process causes data loss.</p> <p>The drive is being accessed and is not configured as part of an array.</p>
On	Flashing	Off	<p>Do not remove the drive. Removing a drive during this process causes data loss.</p> <p>The drive is rebuilding or undergoing capacity expansion.</p>
Flashing	Flashing	Flashing	<p>Do not remove the drive. Removing a drive during this process causes data loss.</p> <p>The drive is part of an array being selected by the ACU.</p> <p>-Or-</p> <p>The Options <i>ROMPaq</i>[™] is upgrading the drive.</p>
Off	Off	Off	<p>OK to replace the drive online if a predictive failure alert is received (see the "Predictive Failure Alert" section in <i>Compaq Servers Troubleshooting Guide</i> for details) and the drive is connected to an array controller.</p> <p>The drive is not configured as part of an array.</p> <p>-Or-</p> <p>If this drive is part of an array, then a powered-up controller is not accessing the drive.</p> <p>-Or-</p> <p>The drive is configured as an online spare.</p>

continued

Table 5-1: Hard Drive LED Combinations *continued*

Activity	Online	Drive Failure	Indication
Off	Off	On	OK to replace the drive online. The drive has failed and has been placed offline.
Off	On	Off	OK to replace the drive online if a predictive failure alert is received (see the "Predictive Failure Alert" section in <i>Compaq Servers Troubleshooting Guide</i> for details), provided that the array is configured for fault tolerance and all other drives in the array are online. The drive is online and configured as part of an array.
On or flashing	On	Off	OK to replace the drive online if a predictive failure alert is received (see the "Predictive Failure Alert" section in <i>Compaq Servers Troubleshooting Guide</i> for details), provided that the array is configured for fault tolerance and all other drives in the array are online. The drive is online and being accessed.

Replacing Failed Hard Drives

The NAS B2000 is designed with many fault-tolerant features to prevent common problems. However, if a drive fails, it must be replaced. Because the operating system drives are set up in a RAID 1 array, the loss of one of the two drives does not result in any loss of system function. If the external drives are set up in RAID 0 arrays with no fault tolerance, the loss of a single drive will cause data loss. However, it is more likely that the external drives are set up in RAID arrays that offer fault tolerance, so the loss of a single drive in an array does not result in any loss of data. Regardless, failed drives must be replaced as soon as possible. Until a failed drive is replaced, the NAS device is operating in a non-fault-tolerant mode. If the other drive is lost before the failed one is replaced and rebuilt, the system will fail.

NOTE: Refer to the maintenance and service guide for detailed procedures about hardware failures on the NAS device.

Failed drives can be replaced without the server being powered down. The drives in your server are hot pluggable. To identify failed drives, look for one or more of the following:

- An amber LED is illuminated on a failed drive. This LED indicates that the storage enclosure is powered up and that the SCSI cable connecting the storage enclosure to the server is working.

NOTE: The amber light may briefly illuminate when the drives are inserted. This illumination is normal and does not indicate a failure condition unless the LED remains illuminated.

- An amber LED is illuminated on the storage enclosure. This LED indicates that one or more drives in the storage enclosure has failed, a fan failure has occurred, or a high-temperature condition exists.
- A (POST) message lists failed drives when the NAS device is restarted. This message is dependent on the array controller detecting one or more “good” drives.
- The Array Diagnostics Utility (ADU), found on the SmartStart and Support Software CD, reports failed drives.
- Compaq Insight Manager software reports failed drives or prefailure conditions on one or more drives.

Follow these guidelines when replacing a failed drive:

- **Never remove more than one drive at a time (two drives if ADG is being used).** When a drive is being replaced, the controller uses data from the other drives in the array to reconstruct data on the replacement drive. If more than one drive is removed (or two with RAID ADG), a complete data set is not available to reconstruct data on the replacement drive, and permanent data loss could occur.
- **Never remove a working drive.** The amber Drive Failure indicator on the drive carrier indicates a drive that has been failed by the controller. Unless the drive is a member of a RAID ADG array, permanent data loss will occur if a working drive is removed while replacing a failed drive.

- **Never remove a drive while another drive is being rebuilt.** A drive's online indicator flashes green (once per second) while it is being rebuilt. A replaced drive is rebuilt from data stored on the other drives.
- **If the system has an online spare drive, wait for it to complete rebuilding before replacing the failed drive.** When a drive fails, the online spare becomes active and begins rebuilding as a replacement drive. After the online spare has completed Automatic Data Recovery (the Online indicators will be continuously lit), replace the failed drive with a new replacement drive. **Do not** replace the failed drive with the online spare. The system will automatically rebuild the replacement drive and reset the spare drive to an available state.

To replace a failed drive:

1. Determine the position of the failed drive.

A lighted LED indicator means that the drive has failed. When identifying drives by drive number, refer to the following table showing how bay numbers correspond to SCSI ID numbers. It is imperative to identify the correct drive before starting the replacement procedure.

NOTE: All Compaq utilities report drive information in terms of SCSI ID. Only the physical drive enclosure is labeled by bay number. Use Table 6-3 to properly identify the appropriate drive bay.

IMPORTANT: Compaq utilities report drive state through the SCSI ID. Table 5-2 assists you in associating the appropriate drive bay to the corresponding SCSI ID.

Table 5-2: Storage Enclosure Drive Bay Configuration

Description														
Bay Number	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	0	1	2	3	4	5	8	9	10	11	12	13	14	15

2. Remove the failed drive by unlatching and sliding it out of the drive bay about one inch.

3. Allow a few seconds for the drive to spin down before removing it completely.



CAUTION: Be sure to use a drive bay blank to cover open bays if you do not plan to insert a new drive within a few seconds. The drive bay blanks are necessary to maintain proper airflow for cooling. Failure to cover open bays can lead to thermal failure of your drives and loss of data.

4. Insert a new drive of the same type into the bay where the failed drive was located.

Compromised Fault Tolerance

If the fault tolerance of an array is compromised due to multiple concurrent drive failures, the condition of the logical drive or drives on that array is “failed,” and unrecoverable errors are returned to the operating system. Data loss is probable. Inserting replacement drives at this time does not improve the condition of the logical drive or drives.

If this situation occurs, first try powering down the entire system and then powering up. Then remove power from both the server and storage enclosure. Reapply power to the storage enclosure, then the server, and then restart the system. In some cases, a drive with intermittent problems can work long enough for you to make copies of important files. If a 1779 POST message is displayed, press the **F2** key to re-enable the logical drive or drives. Remember, it is likely that data loss has occurred, and any data on a failed logical drive is suspect.

Fault tolerance can also be compromised due to non-drive issues. These issues include faulty SCSI and power cables, power supplies, facility power, or accidental unplugging of storage enclosure power. In such cases, the physical drives in the storage enclosures need not be replaced. However, data loss is possible.

In cases of actual drive failure, replace any drives that have failed, to prevent further problems. Afterwards, the fault tolerance may again be compromised, power may need to be recycled, and the 1779 POST message may again be displayed. Press the **F2** key to re-enable the logical drive or drives. Then recreate your pools and virtual disks, and restore data from backup media.

Moving Hard Drives



CAUTION: All data must be backed up before removing drives or changing configurations. Failure to do so could result in permanent loss of data. The system must be powered down.

Drives in the external storage enclosures can be moved, but this should only be done after a complete and successful backup and when the server has been powered down and turned off. To move drives, the following conditions must be met:

- System is powered down (includes all system components).
- No drive failures are identified. The array must be in its original configuration.
- Capacity expansion or drive rebuild is not in progress.
- Controller firmware is the latest version (recommended).

When the above conditions are met, move the drives using the following procedure:

1. Remove one drive at a time by unlatching and sliding it out of the drive bay about one inch.

IMPORTANT: Before removing a drive from the drive bay, be sure to allow a few seconds for the drive to completely spin down before handling it. If the external storage enclosure has already been powered down, the drives should already have spun down.

2. Move the drive to the desired location in an external storage enclosure.
3. Repeat steps 1 through 2 for all of the drives that need to be moved.

IMPORTANT: All drives in an array must be moved at the same time or data loss will occur.

4. Restore power to the external storage enclosure or enclosures, and then power up the server. As the system restarts, a 1724 POST message is displayed, indicating that the drive positions have changed and have been updated. If a 1785 POST message is displayed, immediately power down the system components to prevent data loss and return each drive to its original location. Then power up the server as usual.
5. If desired, run the ACU to view and confirm the new drive positions.

When moving drives, refer to Table 5-2, “Storage Enclosure Drive Bay Configuration.”



CAUTION: All data must be backed up before removing drives or changing configurations. Failure to do so could result in permanent loss of data. The system must be powered down before removing drives.

Moving Arrays

When a company has multiple NAS devices, situations may arise in which an entire array needs to be moved from one server to another. Because the data on the servers is stored in virtual disks created from drive space pools that may contain space from multiple arrays, care must be taken when moving arrays. The following guidelines are provided:

- All drives with arrays and logical drives used in a pool must be moved at the same time.
- Use the ACU to identify which drives belong to the array being moved. In the ACU interface, highlighting the array causes the drive lights to flash on the member disks of the array.
- Move all drives in the array at the same time.
- Make sure that no failed drives are present and that no rebuild processes are in progress.
- Make sure that the positions of other drives on the system to which the array is being moved are not being changed at the same time.
- Make a full backup of all data on the server from which the array is being moved before starting the move process.
- To move the array, power down the system, and unplug power from the server. Then, remove power from the storage enclosure or enclosures and carefully remove the drives.



CAUTION: To avoid data loss, replace failed drives according to configuration guidelines in Table 5-2.

Array and LUN Management

NOTE: For consolidation purposes, detailed overview information of all storage issues is included in a separate chapter. The “Storage Management Overview” chapter introduces and discusses in detail arrays, LUNs, RAID levels, pools, virtual disks, and snapshots.

In addition, the “Storage Management Planning” chapter includes detailed planning information that guides the administrator through the decision-making process of determining the best configuration of their storage.

The physical disks, arrays, and their corresponding logical units (LUNs) are managed using the Compaq Array Configuration Utility (ACU). Virtual Replicator then uses the LUNs to create storage pools, virtual disks, and snapshots.

This section provides instructions for using the ACU. VR procedures are discussed in the following chapter. The following topics are included in this section:

- Compaq ACU Overview
- Accessing the ACU
- Creating a New Array
- Creating a New Logical Drive
- Expanding the Capacity of an Existing Array
- Migrating an Existing LUN to a New RAID Level or Stripe Size

Compaq ACU Overview

The ACU is graphical tool, incorporating wizard-style interfaces to create RAID arrays and logical drives from the physical drives installed in the storage subsystems. The drive arrays should be configured using the RAID level that meets the fault tolerance, cost effectiveness, and I/O performance needs of the environment using those arrays. For detailed planning discussions of recommended configuration striping methods and RAID levels, see Chapter 4.

The ACU can be used to grow an existing array by incorporating new drives into the array. The associated LUN and its data are automatically rewritten and re-striped over all of the drives now in the array. The new extra capacity in the array can be used for a variety of purposes, including reconfiguring associated LUNs or creating new LUNs and using them to grow the size of the array's pool.

The ACU can also completely reconfigure an existing LUN, including changing the RAID type and the stripe size. While these procedures are time consuming, they are not data destructive and can be performed online.



CAUTION: Existing logical drives must be deleted only after deleting the virtual disks and pools associated with these logical drives. See the "Virtual Storage Management" chapter for information on deleting virtual disks and pools.

Features of the ACU

- Graphical representation of drive array configurations with wizards that help optimize array configuration
- Online spare (hot spare) configuration
- Separate fault-tolerant configurations on a logical drive (LUN) basis
- Easy capacity expansion of arrays
- Online RAID level and stripe size migration

Accessing the ACU

To access the ACU:

1. Log on to the NAS device as an administrator and go to the WebUI.
2. From the WebUI, navigate to **Disks, Array Configuration Utility**.

A Terminal Services session is automatically opened, and prompts you for user identification. Enter an administrator-level name and password to access the ACU.

If the hard drives and the arrays are unconfigured or if the configuration is less than optimal, a configuration wizard guides you through the configuration process of adding the drives to arrays or creating LUNs.

NOTE: The automatic wizard can be bypassed to manually configure arrays by clicking the "Cancel All" selection.

3. After the Configuration wizard is finished, or if it is bypassed, the main configuration screen of the ACU is displayed. All array and LUN configuration tasks can be performed in the ACU main configuration screen.

In the **Controller Selection** drop-down box near the top of the screen, select the controller that needs to be configured.



CAUTION: Do not modify the settings for the Integrated Smart Array Controller in Array A. Changes to the configuration of Array A can corrupt the operating system on the server.

After a controller is selected, the screen display is refreshed to include information about the arrays and logical drives associated with the chosen controller.

If the selected controller is unconfigured, the screen will show only unassigned drives. When the ACU utility detects an unconfigured controller, a configuration wizard leads you through the controller configuration process.

If the selected controller has already been configured, arrays, logical drives, and unused space are displayed.

4. The ACU presents either a logical or a physical view of the drives, arrays, and LUNs. Control the presentation of the view by using the radio buttons at the bottom left of the screen.

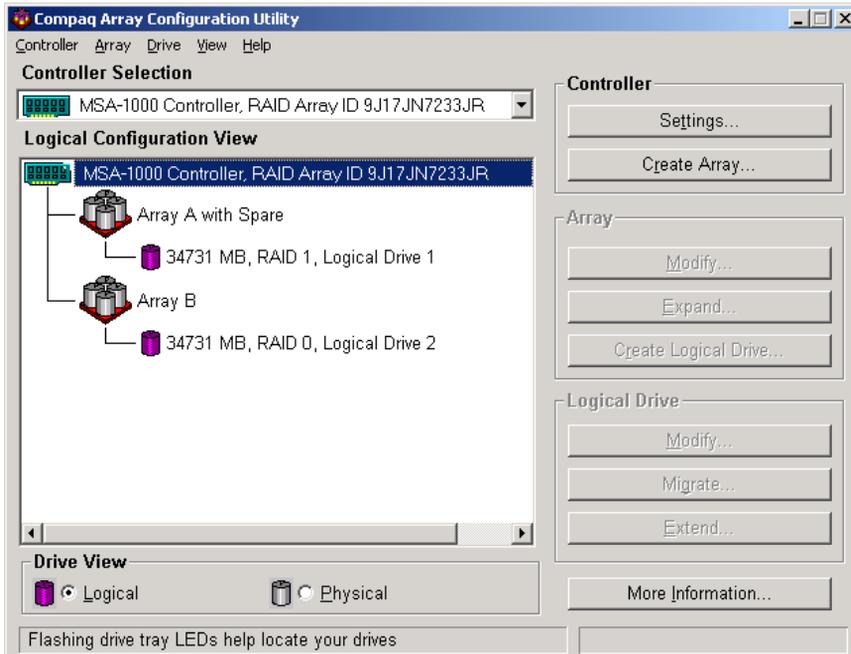


Figure 5-2: ACU Logical Drive view

Select either **Logical Drive View** or **Physical Drive View**. Figure 5-2 and Figure 5-3 are examples of the two presentation methods.

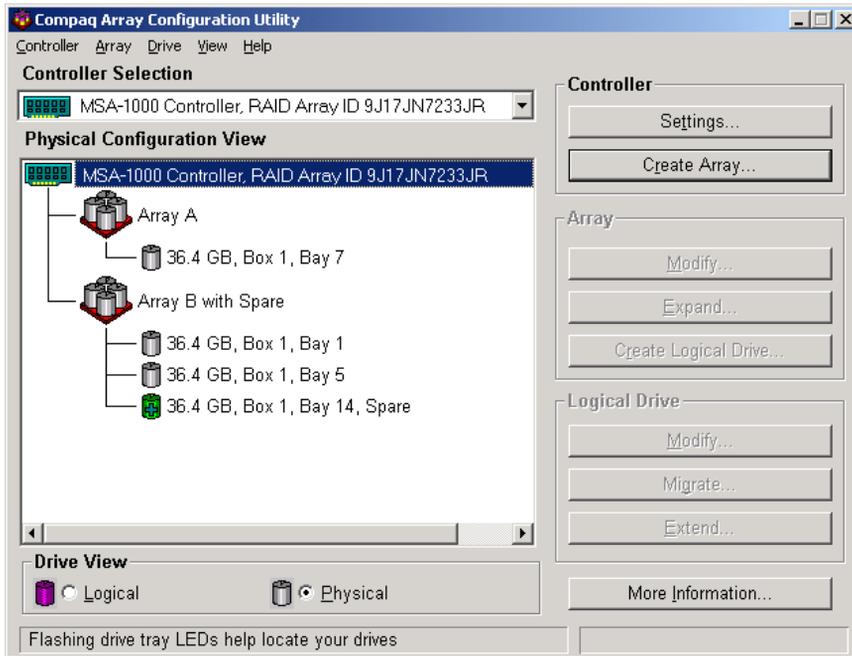


Figure 5-3: ACU Physical view

NOTE: Selecting an item—a controller, array, logical drive, or physical drive—in the Configuration View box will cause the hard drive tray LEDs to blink. Use this feature to identify a specific physical drive or to identify the drives in the storage enclosures that are attached to the controller.

5. Action buttons are displayed in the right-hand portion of the dialog box. Some action buttons on the screen are highlighted and some appear gray. Appropriate buttons are available depending on the items selected to configure.

Actions include:

- **Controller Settings**—used to enter parameter settings for the arrays and LUNs associated with this controller.
- **Create Array**—used to create new arrays from unassigned physical hard drives.

- **Modify Array**—used to make changes to an existing array. This option can be used to add or delete physical drives from an array and is data destructive.
 - **Expand Array**—used to grow an array by adding unassigned physical hard drives to an existing array.
 - **Create Logical Drive**—used to convert an array into a logical drive (LUN). It is during this process that the RAID level is assigned.
 - **Modify Logical Drive**—used to change the Array Accelerator setting for an existing logical drive.
 - **Migrate Logical Drive**—used to convert a LUN from one RAID configuration to a new RAID configuration.
 - **Extend Logical Drive**—this option is not available for this operating system.
6. To see detailed information about a specific controller, array, or LUN, click **More Information** at the bottom right of the screen.

Entering Controller Settings

The controller settings determine how much importance to place on array expansion or rebuilding relative to normal I/O operations. Additionally, if the controller has a battery-backed cache, the ratio of read cache to write cache can be changed.

The default controller settings provided by the ACU will be adequate for most environments. However, to change the controller settings:

1. Access the ACU and select the desired controller in the Controller Selection drop-down box.



CAUTION: Do not modify the settings for the Integrated Smart Array Controller in Array A. Changes to the configuration of the embedded controller in Array A can corrupt the operating system on the server.

2. Click **Controller Settings**. The Controller Settings dialog box is displayed.

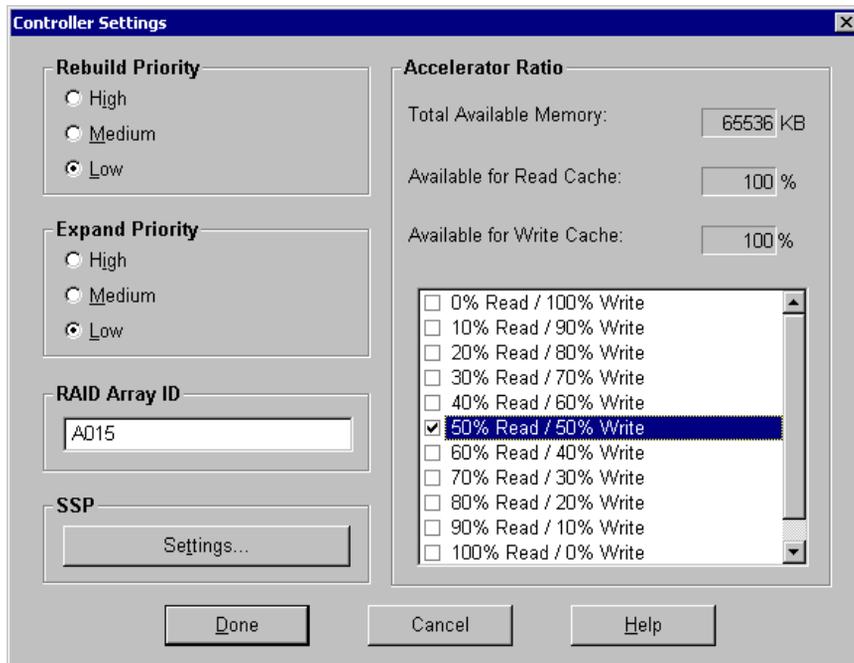


Figure 5-4: Controller Settings dialog box

3. Select **Rebuild Priority**.

The Rebuild Priority affects the amount of time the controller spends rebuilding data after a failed drive has been replaced.

If the array rebuild is set to **low** priority, the rebuild will take place only when the array controller is not busy handling normal I/O requests. This setting has minimal effect on normal I/O operations. With a low rebuild priority, however, there is an increased risk that data will be lost if a physical drive fails while the rebuild is in progress.

If the rebuild has **high** priority, the rebuild occurs at the expense of normal I/O operations. Although system performance is affected, this setting provides better data protection because the array is vulnerable to additional drive failures for a shorter time.

When the priority is set to **medium**, rebuild still has greater priority than I/O requests, but less than with the high setting.

4. Select **Expand Priority**.

The Expand Priority affects the amount of time the controller spends rewriting data and restriping the LUNs during an array expansion.

If the expansion is set to **low** priority, the expansion will take place only when the array controller is not busy handling normal I/O requests. This setting has minimal effect on normal I/O operations. With a low expansion priority, however, there is an increased risk that data will be lost if a physical drive fails while the expansion is in progress.

If the expansion has **high** priority, the expansion occurs at the expense of normal I/O operations. Although system performance is affected, this setting provides better data protection because the array is vulnerable to additional drive failures for a shorter time.

When the priority is set to **medium**, expansion still has greater priority than I/O requests, but less than with the high setting.

5. Select the **Accelerator Read/Write Ratio**.

The Accelerator Read/Write Ratio determines the amount of memory allocated to the read and write caches on the array accelerator. Some applications may perform better with a larger write cache while others may perform better with a larger read cache. This setting can only be changed if the controller has a battery-backed cache.

6. After all settings are entered, click **Done**. The ACU main configuration screen is displayed again.

Creating a New Array

Before creating any arrays, become familiar with the information in the “Storage Management Overview” chapter and use the information and recommendations in the “Storage Management Planning” chapter to develop a corporate storage management plan.

To create a new array:

1. Access the ACU as described previously. In the ACU main configuration screen, select the desired controller from the Controller Selection drop-down menu.



CAUTION: Do not modify the settings for the embedded Integrated Smart Array Controller in Array A. Changes to the configuration of the embedded controller can corrupt the operating system on the server.

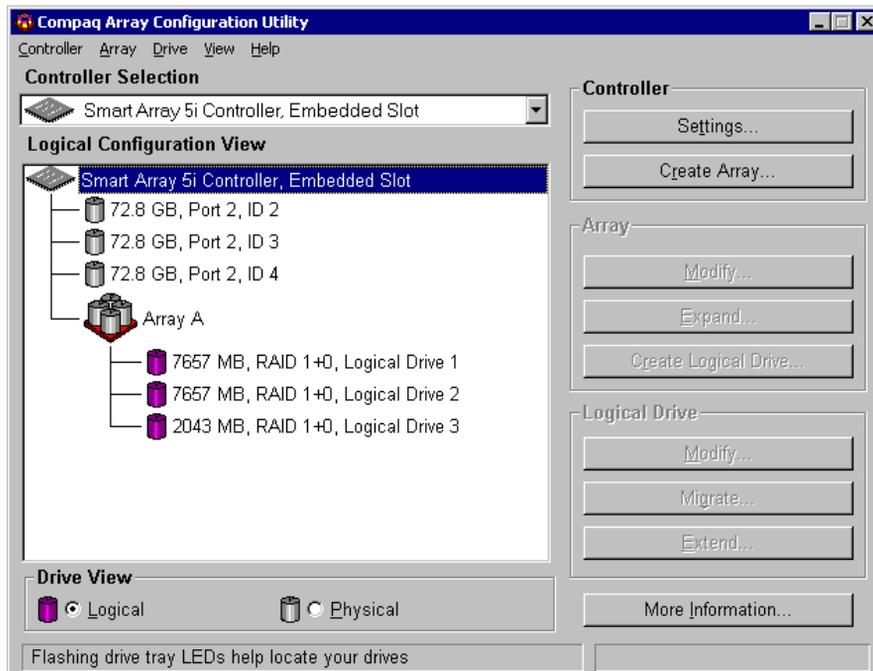


Figure 5-5: ACU main configuration screen

2. Click **Create Array**. The Create Drive Array dialog box is displayed.

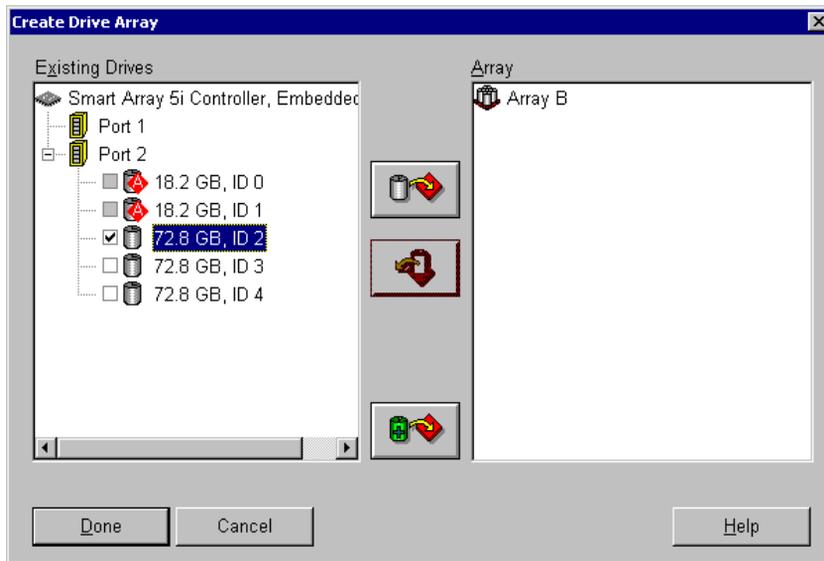


Figure 5-6: Create Drive Array screen

IMPORTANT: Always group physical drives of the same size and type. If drive sizes are mixed, some capacity of the larger drives is wasted. See the “Storage Management Planning” chapter for detailed information, examples, and consequences of mixing drive types.

3. Select all of the physical drives to include in the array by clicking on them in the **Existing Drives** box.

4. Add the drives to the array by clicking the icon showing the **right-pointing arrow (Assign Drives to Array)** option in the center of the screen. (When the cursor passes over this icon, the title for the button is displayed.)

All selected drives are moved into the right pane on the interface, designating their inclusion in the array.

5. If desired, indicate a drive to use as a spare for this array.

This online spare is used automatically and immediately when one of the other member drives in the array fails. Spares can be shared among arrays.

To assign a spare to this array, select the desired drive in the **Existing Drives** box and then click the **right-pointing arrow (Assign Spare to Array)** icon at the bottom center of the screen.

At this point, the Create Drive Array screen may look similar to the following figure.

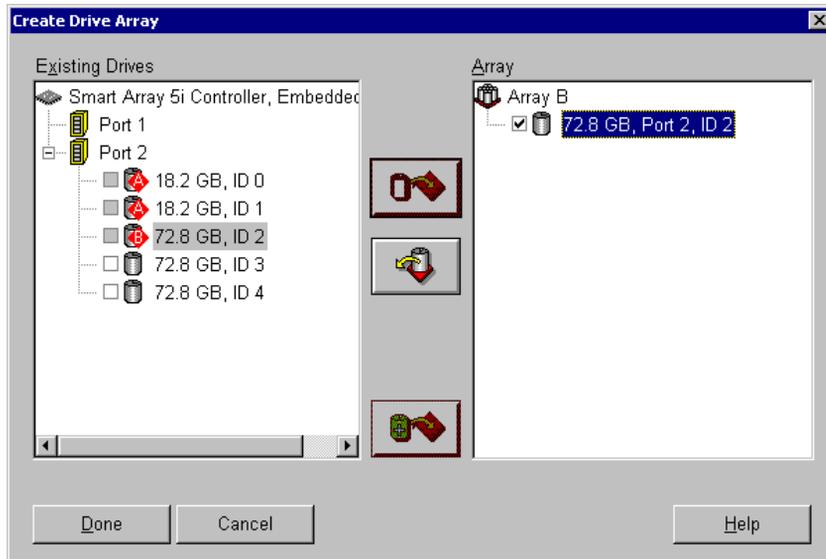


Figure 5-7: Example Array B

NOTE: The same spare drive can be assigned to multiple arrays. However, spare drives should have the same or greater capacity as the drives in the array.

6. Click **Done** to return to the ACU main configuration screen. The Configuration View is updated to show the new configuration. Array information is now displayed instead of the individual hard drive information.

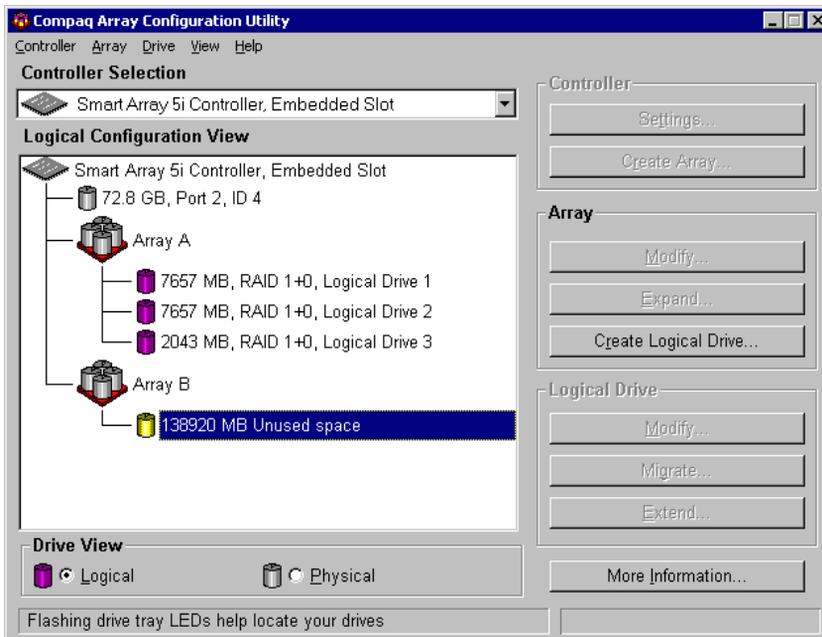


Figure 5-8: Example Array Logical Configuration view with two arrays

IMPORTANT: A logical drive (LUN) must be created for the arrays before exiting the ACU or the array setup will not be saved. Use the procedure in the following section to create a logical drive from the array.

7. Create additional arrays for this controller, using any remaining unused physical hard drives attached to this controller.
8. Create the logical drives (LUNs) for these arrays.

Creating Logical Drives (LUNs)

After the physical drives are gathered into arrays, they need to be converted into fault-tolerant LUNs. When creating a logical drive, the fault-tolerance (RAID level) is selected along with settings regarding the LUN size, array accelerator use, and the stripe size.

Although multiple LUNs can be created from one array, Compaq recommends creating one LUN from the array.

To create a new LUN:

1. Access the ACU, and in the ACU main configuration screen, select the desired controller and array.
2. Click **Create Logical Drive**. A screen similar to the following figure is displayed.

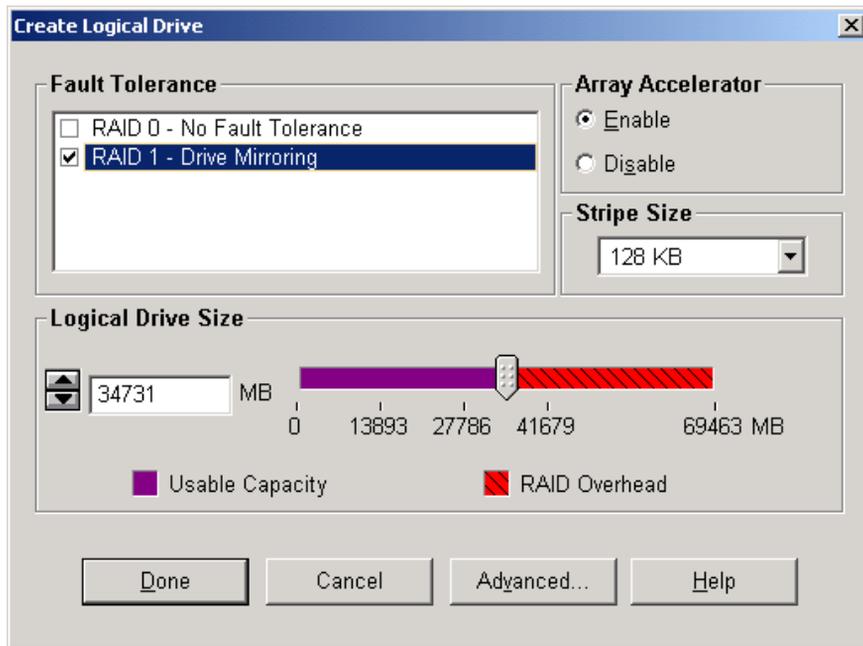


Figure 5-9: Create Logical Drive dialog box

3. In the Fault Tolerance box, indicate the desired level of fault tolerance for the LUN. Select the RAID type from the displayed list.

NOTE: Only the applicable RAID types are displayed for the array. For example, in a two-drive array, only RAID 0 and RAID 1 are available

4. Select **Enable Array Accelerator**.
5. Set the **Stripe Size** to the desired value or accept the default.

Stripe size refers to the amount of data stored on each physical drive in one stripe of a logical drive. Each RAID level has a default value plus a range of supported sizes. The default values provide optimum performance for that RAID level in most applications.

To select a stripe size other than the default, click the down arrow next to the displayed default stripe size and select from those available.

Table 5-3: Optimum Stripe Sizes for Different Environments

Server Application Environment	Suggested Strip Size Change
Mixed read/write	Accept the default value.
Mainly sequential read (such as audio/video applications)	Larger stripe sizes work best.
Mainly write (such as image manipulation applications)	Smaller stripe sizes for RAID 5 and RAID ADG Larger stripe sizes for RAID 0, RAID 1, and RAID 1

6. Set the **Logical Drive Size**.

To create a logical drive that uses the entire array, use the default values. The utility does not allow you to create a logical drive larger than the maximum size supported by the operating system.

The Logical Drive Size area of the screen includes a scale marked with the amount of raw storage capacity in the array. The left side of the Logical Drive Size scale indicates the amount of space available for data. The right side of the scale indicates the amount of space required for storing parity or mirrored information, depending on the fault-tolerance method.

7. Click **Done**.

The Configuration View screen may now look like Figure 5-10.

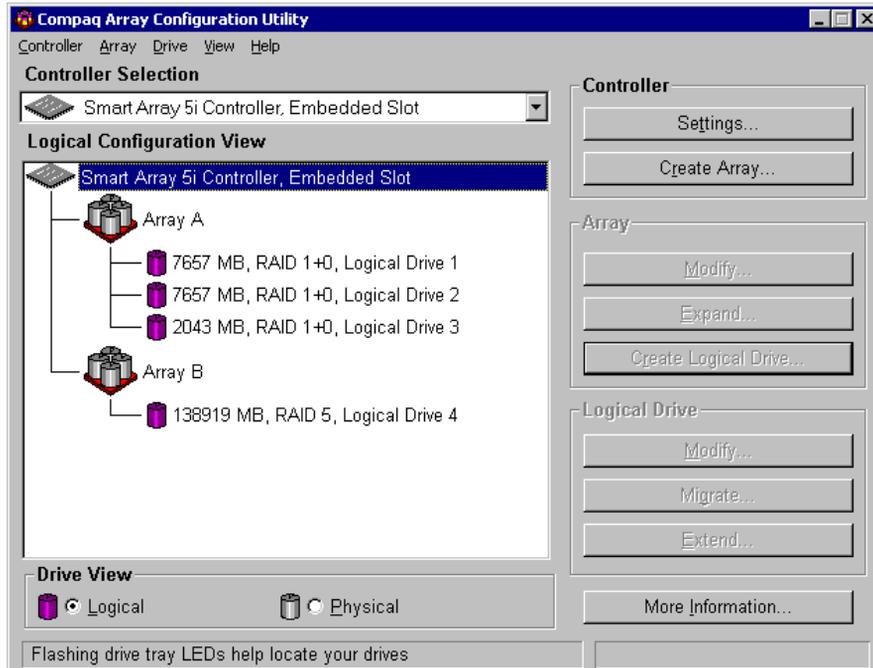


Figure 5-10: Example array - Configuration View screen with two arrays

8. If there are other arrays that need to be converted into LUNs, repeat steps 1 through 8 to create the LUNs for those arrays as well.

IMPORTANT: A logical drive (LUN) must be created for arrays before exiting the ACU or the array setup will not be saved.

9. After the arrays and logical drives are set up, the LUNs can be placed in to virtual storage pools from which virtual disks are created. Pools and virtual disks are detailed in the “Virtual Storage Management” chapter.

Expanding the Capacity of an Existing Array

Capacity expansion increases the storage capacity of an existing array. The Smart 5i and Smart 5300 array expansion feature included in the NAS B2000 allows unused physical drives to be added to an array. During the expansion process, the controller will rearrange (re-stripe) the existing logical drives and their data so that they span all of the physical drives in the expanded array. The size of the existing LUNs remains constant and the data is preserved.

The new capacity in the array can then be used to migrate the RAID level of the LUNs of this array, to change the stripe size of LUNs of this array, or to create new LUNs. If the purpose of expanding the array was to grow a pool, this new LUN can now be added to that pool. Adding units to a pool is discussed in the “Virtual Storage Management” chapter.

IMPORTANT: The expansion process takes about 15 minutes per GB or considerably longer if the controller does not have a battery-backed cache. While array expansion is taking place, no other expansion or migration can occur on the same controller.

NOTE: During a hard drive expansion, migration, or extension process, the redundancy feature of the controllers will be temporarily disabled. At the end of the expansion, migration, or extension process, the redundancy feature of the controllers will be automatically reinstated; no action is required by the user. This scenario is for expansion and migration only; the controllers remain fully redundant during a drive rebuild process.

To grow an array:

1. If necessary, install new physical drives.
2. Back up all data on the array. Although array expansion is unlikely to cause data loss, this precaution will provide additional data protection.
3. From the ACU main configuration dialog box, click **Controller Settings** and verify that the **Expand Priority** setting is acceptable.
4. From the ACU main configuration dialog box, select the array to grow and then click **Expand Array**.

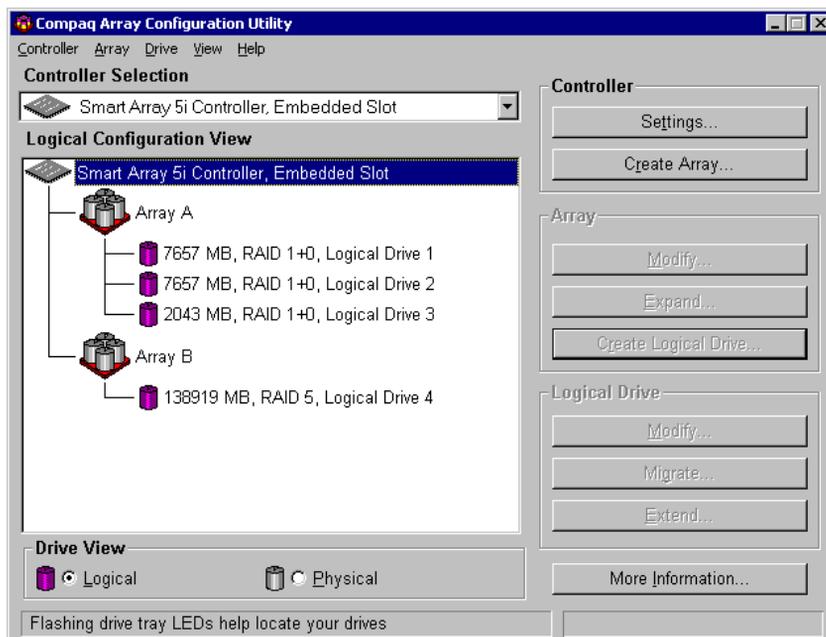


Figure 5-11: Array expansion example – Logical Configuration View screen

5. A subscreen is displayed. Select the intended unassigned drives to add to the array.
6. Click the **right-pointing arrow (Assign Drives to Array)** in the center of the screen to add the drives to the array.

- Click **Done** at the bottom of the screen. A screen similar to the following figure is displayed.

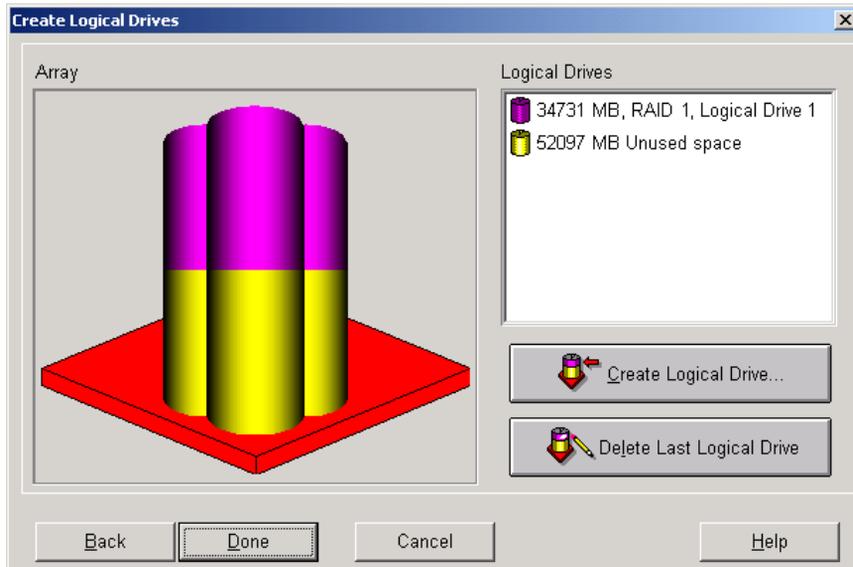


Figure 5-12: Expansion wizards - Logical Drive screen

- To create a new LUN with the available capacity, click **Create Logical Drive**.
- In the Create Logical Drive subscreen, indicate the settings for the Fault Tolerance, Array Accelerator, Stripe Size, and Logical Drive size.
To enable the Maximum Boot Size, click the **Advanced** button.
- Click **Done**.
- In the ACU Main Controller Configuration dialog box, save these settings.

On the menu bar at the top of the screen, select **Controller**. Then, select the **Save Configuration** option. The settings for the new LUN are saved and the capacity expansion process starts.



CAUTION: In case of power loss, capacity expansion process information is temporarily stored in the Array Accelerator memory. To prevent the loss of data in the expanding logical drive, do not interchange Smart 5i or Smart 5300 Controllers or Array Accelerator boards during a capacity expansion process.

NOTE: If several arrays are to be expanded, the system will queue the requests, expanding one array at a time.

NOTE: The new LUN will not be accessible until the capacity expansion process has completed on the array.

Migrating an Existing LUN to a New RAID Level or Stripe Size

Use the Online RAID Level and Stripe Size Migration screen to reconfigure a currently configured logical drive to a new fault-tolerance (RAID) level or to change an existing logical drive's stripe size to a new stripe size (data block size).

Depending on the initial settings and the new RAID type and Stripe Size settings for the LUN, unused capacity may need to be available for the migration. Additional drives may need to be included in the array.

Both of these procedures can be done online without causing any data loss.

IMPORTANT: The migration process takes about 15 minutes per GB or considerably longer if the controller does not have a battery-backed cache. While migration is taking place, no other expansion or migration can occur on the same controller.

To migrate a logical drive to a different RAID level or stripe size:

1. Back up all data on the LUN. Although migration is unlikely to cause data loss, this precaution will provide additional data protection.
2. From the ACU Main Controller Configuration dialog box, select the appropriate controller from the drop-down box, select the target logical drive, and click **Migrate Logical Drive**. A screen similar to the following is displayed.

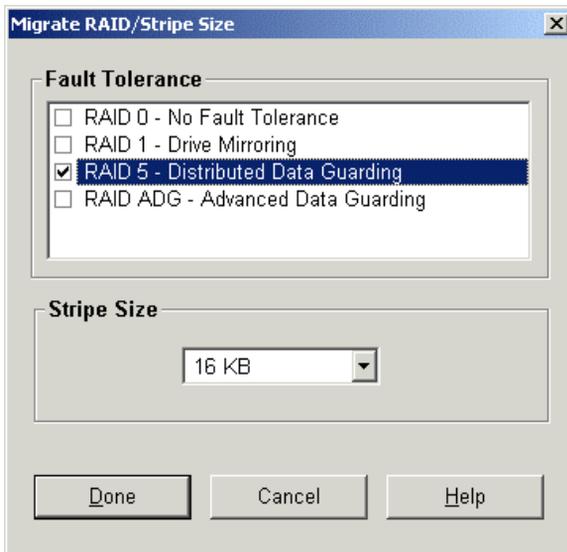


Figure 5-13: Migrate RAID/Stripe Size screen

3. To set the new level of fault tolerance, select the desired RAID type in the Fault Tolerance portion of the screen.
4. To set a new stripe size, either accept the default size for the selected RAID level, or set to another value.
5. Click **Done**.

Virtual Storage Management

Virtual storage elements of the *StorageWorks* NAS B2000 include storage pools, virtual disks, and snapshots. These virtual elements are created from the physical drive arrays and LUNs. For complete conceptual information on physical and virtual storage, see the “Storage Management Overview” chapter.

Because all storage overview and planning information is discussed in a separate chapter of this guide, this chapter is dedicated to the following virtual storage procedural tasks:

- Storage Management Wizard
- Pool Management
 - Creating a New Pool
 - Deleting a Pool
 - Viewing Pool Properties
 - Setting Pool Policies for a Specific Pool
 - Adding Storage Units
 - Bringing a Pool Online and Offline
- Virtual Disk Management
 - Creating a New Virtual disk
 - Setting the Drive Letter of a Virtual disk
 - Formatting a Virtual Disk
 - Deleting a Virtual Disk

- Viewing Virtual Disk Properties
- Growing a Virtual Disk
- Backing Up (Creating a Scheduled Snapshot of) a Virtual Disk
- Snapshot Management
 - Creating a New Snapshot
 - Deleting a Snapshot
 - Viewing Snapshot Properties
 - Setting the Drive Letter for a Snapshot
 - Creating a Snapshot Schedule
 - Displaying and Deleting a Snapshot Schedule
 - Scheduling Snapshot Deletions
 - Restoring a Virtual Disk from a Snapshot
 - Enabling Incremental Backup Support
- Global Pool Policy Settings
- Namespace Recovery
- Drive Quotas
 - Enabling and Disabling Quota Management on a Virtual Disk
 - Creating New Quota Entries for a User or Group
 - Deleting Quota Entries for a User or Group
 - Modifying Quota Entries for a User or Group

Storage Management Wizard

The Storage Management wizard is one of three major wizards included in the Web-based user interface (WebUI) of the NAS device. This wizard contains some of the functionality of the detailed storage management screens and is designed to be intuitive and easy-to-use.

The functionality included in this wizard can also be performed through the **Virtual Replicator** menu option of the WebUI. Detailed information on storage management topics is included in later sections of this chapter.

To use the Storage Management wizard:

From the WebUI, select **Wizard Tasks**, and **Storage Management**. The Welcome screen of the wizard is displayed. Click **Next** to continue.

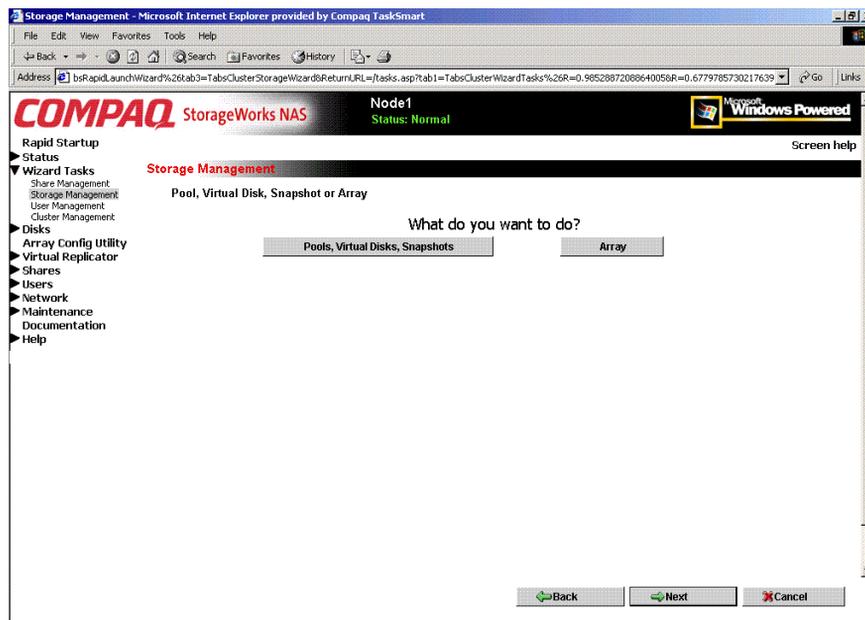


Figure 6-1: Storage Management wizard

The primary screen of the Storage Management wizard is an expandable selection box, listing all existing storage components. Depending on the components that have been set up, the following components may be listed:

- Pools
- Virtual Disks
- Snapshots

Depending on the component that is selected, different options become available. This discussion is organized according to these storage components.

Wizard Tasks

Selecting a Server

Within the Storage Management wizard, if a server is selected to manage, the Create New Pool option is displayed:

To create a new pool:

1. Click **Create New Pool**.

The next screen of the wizard is displayed, listing the available storage units, their capacity, and type.

2. Select the Storage Unit (LUN) to use, indicate a pool name to assign, and specify the Segment Size. Then, click **Create**.

Selecting a Pool

Within the Storage Management wizard, if a pool is selected to manage, two options are displayed:

- **Create virtual disk**
- **Delete this pool**

To create a new virtual disk:

1. Open a new Web browser window and log in to the NAS operating system console as *administrator* through the Remote Insight Lights-Out Edition board.

2. Select **Create virtual disk**.

The next screen of the wizard is displayed, listing the pool and the amount of free space that is available.

3. Enter the name and size of the virtual disk, indicate whether to assign a drive letter, and specify the allocation unit size. Then, click **Create**.

NOTE: The console session must be established through the Remote Insight Lights-Out Edition board or a directly connected keyboard, monitor, and mouse. The login session must remain active while creating a virtual disk. The console session may be locked for security purposes.

To delete the pool:

1. Select **Delete this pool**.

A verification box is displayed.

NOTE: There is no additional warning prompt.

2. Confirm that the intended pool has been selected and then click **OK**.

Virtual Disk Wizard Tasks

Within the Storage Management wizard, if a virtual disk is selected to manage, two options are displayed:

- **Create snapshot**
- **Delete this virtual disk**

To create a snapshot:

1. Select **Create snapshot**.
2. Enter a name to assign the snapshot, a drive letter, and a label. Click **Create**.

To delete the virtual disk:

1. Select **Delete this virtual disk**.

A verification box is displayed.

NOTE: There is no additional warning prompt.

Confirm that the intended virtual disk has been selected and then click **OK**.

Snapshot Wizard Tasks

Within the **Storage Management** wizard, if a snapshot is selected to manage, two options are displayed:

- **Delete this snapshot**
- **Restore virtual disk from this snapshot**

To delete this snapshot:

1. Select **Delete this snapshot**.

A verification box is displayed.

NOTE: There is no additional warning prompt.

Confirm that the intended snapshot has been selected and then click **OK**.

To restore a virtual disk from a snapshot:

1. Select **Restore virtual disk from this snapshot**.

Pool Management (Details)

Virtual Replicator enables the grouping of LUNs into a single logical pool of drive space. The LUNs provide drive space for the pool in the same way that physical drives provide drive space for a RAID array.

Virtual Replicator is a utility for managing disk space logically and efficiently. VR combines logical drives into a pool of space, from which the administrator can create virtual disks of exactly the correct size needed by a particular department or application.

Pools are managed through the WebUI, with the **Virtual Replicator** menu option from the main menu. To create and manage storage pools, select the **Pools** option from the **Virtual Replicator** menu. The Manage Pools dialog box is displayed.

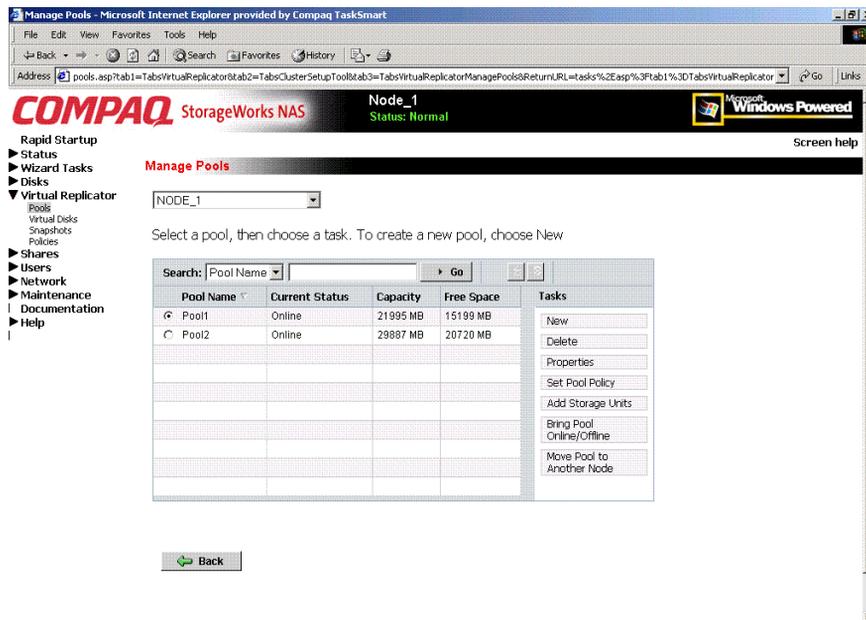


Figure 6-2: Manage Pools dialog box

All existing pools are displayed. Status and capacity information is displayed for each pool.

The **Manage Pools** dialog box allows the following tasks:

- Creating a new pool
- Deleting a pool
- Modifying pool properties
- Setting pool policies
- Adding storage units to a pool

Each of these tasks is discussed in the following sections.

Creating a New Pool

The process of creating a pool of disk space and dividing it into virtual disks is straightforward:

1. From the **Virtual Replicator**, **Pools** option, click **New**. The **Create New Pool** dialog box is displayed.

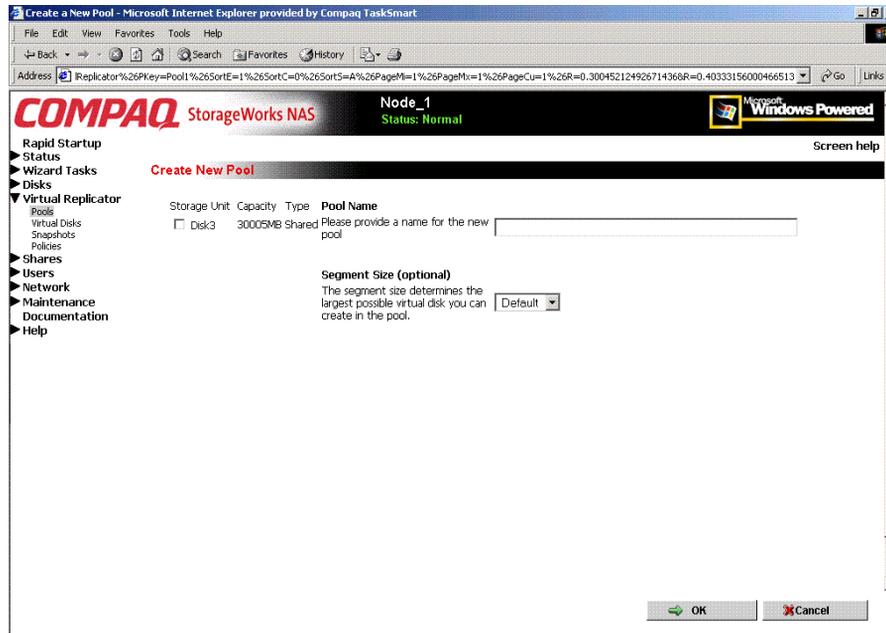


Figure 6-3: Create New Pool dialog box

2. Available LUNs are displayed. Select the units to include in the new pool, and enter a descriptive name for the new pool.

NOTE: Default system settings allow up to eight logical drives as members of a pool. All LUNs must come from the same controller pair.

3. Select the desired segment size from the **Segment Size** drop-down box. Compaq recommends accepting the default size of 256 KB.
4. Click **OK**.

Deleting a Pool

To delete a pool:

1. From the **Manage Pools** dialog box, select the pool to delete from the displayed list and then click **Delete**.
2. A verification screen is displayed. To proceed with the deletion, click **OK**. The pool is deleted and the Manage Pools dialog box is displayed again.

Viewing Pool Properties

To view the properties of a pool, from the Manage Pools dialog box, select a pool from the displayed list and then click **Properties**. The Pool Properties screen is displayed. This interface is a summary display. To change any of these properties, use the appropriate action in the Manage Pools dialog box.

The screenshot shows the 'Pool Properties' summary display in a web browser. The interface includes a navigation menu on the left with options like 'Rapid Startup', 'Status', 'Wizard Tasks', 'Disks', 'Virtual Replicator', 'Shares', 'Users', 'Network', 'Maintenance', 'Documentation', and 'Help'. The main content area is titled 'COMPAQ StorageWorks NAS Node_1 Status: Normal' and 'Pool Properties'. It displays a table of pool information and sections for Storage Units, Read Operations, and Write Operations.

Pool	Pool1
Node	NODE_1
Status	Pool Online
Capacity	21995 MB
Created	10/2/2001 2:23:26 PM
Modified	10/2/2001 3:17:13 PM
Free Space	21797 MB
Segment Size	32 KB

Unit	Capacity	Type
Disk:4	22082 MB	Shared

Read Operations	
Reads received	959
Reads issued to disk	959
Split Reads	0

Write Operations	
Writes received	140509
Writes issued to disk	142637
Split Writes	2128
Copy-outs	2128

At the bottom of the window, there are 'OK' and 'Cancel' buttons.

Figure 6-4: Pool Properties summary display

Setting Pool Policies for a Specific Pool

Virtual Replicator parameters are preset on the NAS device, with one primary set of parameters controlling all pools. Through this **Set Pool Policy** action, these parameters can be customized for a specific pool. The global set of parameters will continue to manage all pools that do not have customized policies.

Experienced users can modify these global and pool-specific policy settings. This section discusses setting pool-specific policy parameters. To change the global set of policy settings, see the “Global Pool Policy Settings” section later in this chapter.



CAUTION: Use caution when changing these parameters. Some choices can result in system inefficiencies, difficulties, and data loss.

To change the policy settings for a specific pool:

1. From the **Manage Pools** dialog box, select the target pool.
2. Click **Set Pool Policy**. An introductory **Manage policies** screen is displayed. Click **Next** to continue. An additional Manage policies dialog box is displayed.

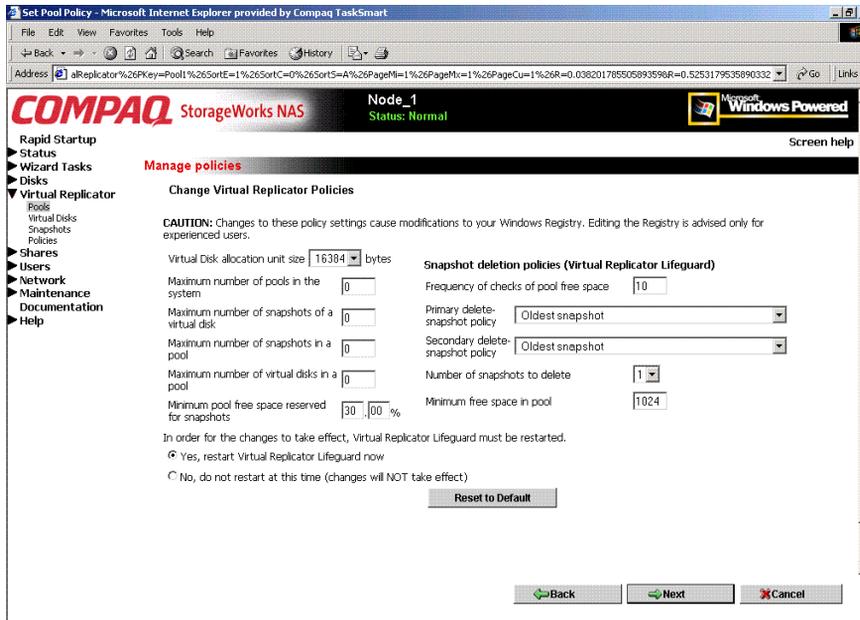


Figure 6-5: Manage Policies dialog box

3. Within this dialog box, enter the new parameter settings.

Default settings include:

Table 6-1: Pool Policy Default Settings

Policy	Default Setting
Virtual disk allocation unit size	16 KB
Maximum number of pools in system	
Maximum number of snapshots of a virtual disk	
Maximum number of snapshots in a pool	
Maximum number of virtual disks in a pool	8
Maximum pool free space reserved for snapshots	10%
Restart Lifeguard now or later	Now
Frequency of pool free space checks	10 seconds
Primary delete snapshot policy	Oldest snapshot first
Secondary delete snapshot policy	Oldest snapshot first
Number of snapshots to delete	1
Minimum free space in pool	1024 KB

4. After all settings are entered, click **Next** and then click **Finish**.

Adding Storage Units to a Pool

Two scenarios may occur when the size of a pool needs to be expanded:

- There are additional, unused LUNs available.
- There are not any additional, unused LUNs available.

If no space is available, physical drives must be added. Then the drives must be configured into a new array, or perhaps an existing array could be expanded to include them. Then, a new LUN can be created and added to the existing pool.

For information on creating and managing arrays and LUNs, see the following sections in the “Physical Storage Management” chapter:

- Creating a New Array
- Expanding the Capacity of an Array
- Creating a New Logical Drive

When the additional LUNs are available, to add them to an existing pool:

1. Select the desired pool in the **Manage Pools** dialog box. Then click **Add Storage Units**. The Add Storage Units dialog box is displayed.
2. All available LUNs are listed. Select the desired LUNs and then click **OK**.

The Manage Pools dialog box is displayed again.

NOTE: Additional LUNs cannot be added if doing so violates any of the Pool Policy parameters.

IMPORTANT: All LUNs in a pool must be managed by the same controller pair.

Bringing a Pool Online and Offline

For maintenance or other reasons, a pool may need to be taken offline.

Within the Manage Pools dialog box, the current status of each pool is displayed. Pool status will be either Online or Offline.

To toggle the status of a pool from Online to Offline or from Offline to Online:

1. Select the desired pool in the **Manage Pools** dialog box.
2. Click **Bring Pool Online/Offline**. A verification screen is displayed.
3. After confirming this is the correct pool, click **OK**. The Manage Pools dialog box is displayed again.

Virtual Disk Management (Details)

Virtual disks are managed through the WebUI, with the **Virtual Replicator** option. From the Virtual Replicator menu, select **Virtual Disks**. The Manage Virtual Disks dialog box is displayed. All existing virtual disks are displayed, including capacity information for each disk.

Figure 6-6 is an illustration of the Manage Virtual Disks dialog box.

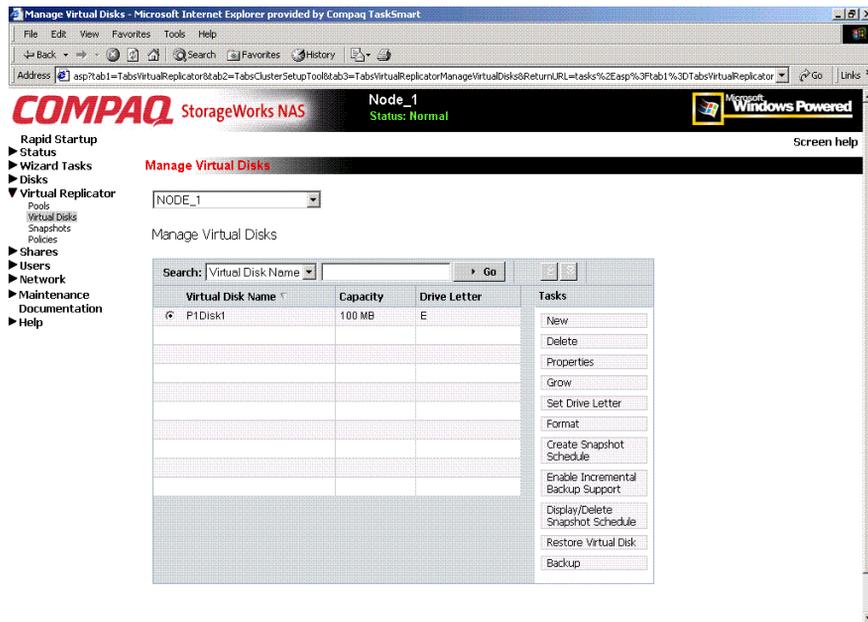


Figure 6-6: Manage Virtual Disks dialog box



CAUTION: Virtual disks must be managed exclusively through the Virtual Replicator interface. Do not manage virtual disks or pools through any other interface. Use of other methods, such as Windows Disk Manager could result in data loss.

The **Manage Virtual Disks** dialog box allows the following tasks:

- Creating a new virtual disk
- Setting the drive letter of a virtual disk
- Formatting a virtual disk
- Deleting a virtual disk
- Viewing virtual disk properties
- Growing a virtual disk
- Creating a snapshot schedule (see “Snapshot Management”)
- Enabling incremental backup support (see “Snapshot Management”)
- Displaying and deleting a snapshot schedule (“see “Snapshot Management”)
- Restoring a virtual disk (see “Snapshot Management”)
- Backing up (creating a snapshot of) a virtual disk

Each of these tasks is discussed in the following sections.

Creating a New Virtual Disk

To create a new virtual disk from a storage pool:

1. Return to the window containing the NAS WebUI.
2. From the **Manage Virtual Disks** dialog box, click **New**. The Welcome screen of the New Virtual Disk wizard is displayed. Click **Next** to continue.

NOTE: The console session must be established through the Remote Insight Lights-Out Edition board or a directly connected keyboard, monitor, and mouse. The login session must remain active while creating a virtual disk. The console session may be locked for security purposes.

The **Create a New Virtual Disk** dialog box is displayed. Figure 6-7 illustrates this dialog box.

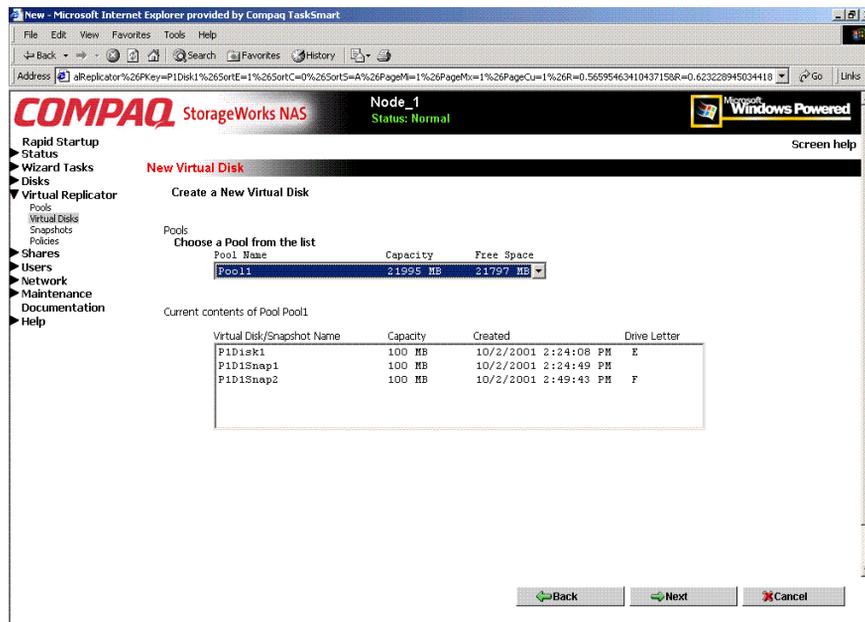


Figure 6-7: Create a New Virtual Disk dialog box

In the **Create a New Virtual Disk** dialog box, the following information is displayed:

- All existing pools, including total capacity and amount of free space
 - For the selected pool, all existing virtual disks and snapshots, including total capacity of each, when they were created, and their drive letters
3. Select a pool from the list from which to create a virtual disk and then click **Next**. The **New Virtual Disk Step 1** dialog box is displayed.

Figure 6-8 is an illustration of the New Virtual Disk Step 1 dialog box.

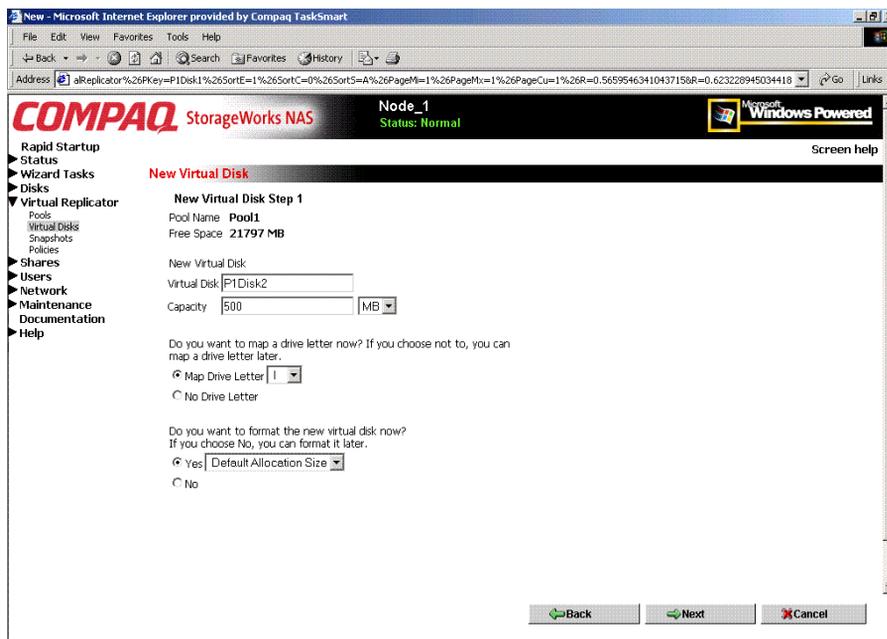


Figure 6-8: New Virtual Disk Step 1 dialog box

4. Enter a name for the new virtual disk and indicate its size.

5. Assign a drive letter (optional).

The operating system does not recognize a virtual disk if it does not have a drive letter assigned. If a drive letter is not assigned at this time, it can be set later, using the **Set Drive Letter** option. See “Setting the Drive Letter of a Virtual Disk.”

6. Format the new virtual disk and indicate the allocation size (optional). Compaq recommends using the default allocation size of 16 KB.

The operating system does not recognize a virtual disk if it is not formatted. If a drive is not formatted at this time, it can be formatted later, using the **Format** option. See “Formatting a Virtual Disk.”

7. Click **Next**. A completion screen showing summary information about the new virtual disk is displayed. Click **Finish** to continue. The Manage Virtual Disks dialog box is displayed again.

Setting the Drive Letter of a Virtual Disk

To set or change the drive letter of a virtual disk:

1. From the **Manage Virtual Disks** dialog box, select the virtual disk to work with and then click **Set Drive Letter**. The **Set Drive Letter** dialog box is displayed.

Current information about the selected virtual disk is displayed, including its current drive letter mapping.

2. In the **Drive Letter** drop-down box, expand the display and select a letter. Only available letters are displayed in the Drive Letter drop-down box.
3. Click **OK**.

The selected letter is assigned to the virtual disk and the Manage Virtual Disks dialog box is displayed again.

Formatting a Virtual Disk

To format a virtual disk:

1. From the **Manage Virtual Disks** dialog box, select the virtual disk to format, and then click **Format**. The **Format Virtual Disk** dialog box is displayed.

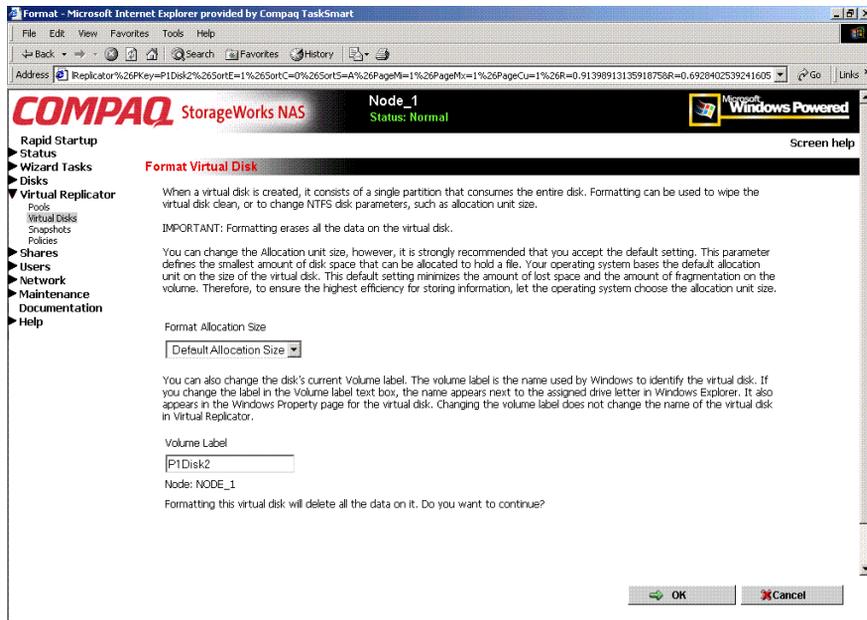


Figure 6-9: Format Virtual Disk dialog box

Summary and descriptive information about formatting is displayed.

2. In the two data entry boxes, enter the allocation size and a label for the disk.

IMPORTANT: No additional warning screen is displayed.

3. View the screen to verify that this is the correct virtual disk and the appropriate allocation size and labels are indicated. Click **OK**.

The virtual disk is formatted and the Manage Virtual Disks dialog box is displayed again.

Deleting a Virtual Disk

When deleting virtual disks, it is important to remember the following:

- Before the underlying disk structure is deleted, virtual disks must be deleted using the Virtual Replicator.
- Do not delete the LUNs before deleting the virtual disks, pools, and snapshots from VR.
- Delete all snapshots associated with a virtual disk before deleting the virtual disk.
- If the virtual disk contains data, back up the data. Deleting a virtual disk destroys all the data stored on the disk.
- Make sure no one is accessing the virtual disk. The disk cannot be deleted if any files on the disk are open.

To delete the virtual disk:

1. From the WebUI, navigate to **Virtual Replicator, Virtual Disks**.
2. In the **Manage Virtual Disks** dialog box, select the virtual disk to delete, and then click **Delete**.
3. A verification screen is displayed. After viewing the screen display to confirm that this is the correct virtual disk, click **OK**. All information on that virtual disk is erased, and the virtual disk is deleted. The Manage Virtual Disks dialog box is displayed again.

NOTE: When deleting a virtual disk, make sure to delete any scheduled tasks for that disk. Otherwise, if the old virtual disk name is reused for another virtual disk, the old tasks will run on the new disk. Use the Windows Scheduled Tasks applet to delete scheduled tasks.

Viewing Virtual Disk Properties

To view the properties of a virtual disk, from the Manage Virtual Disks dialog box, select a virtual disk from the displayed list and then click **Properties**. The Virtual Disk Properties screen is displayed.

This interface is a summary display. To change any of these properties, use the appropriate action in the Manage Virtual Disks or the Manage Snapshots dialog boxes.

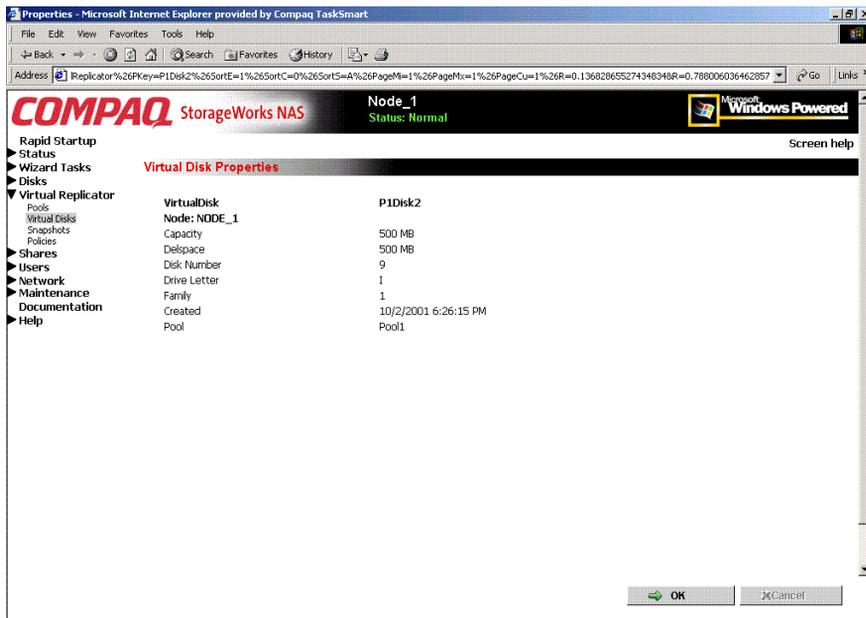


Figure 6-10: Virtual Disk Properties screen

Growing a Virtual Disk

Virtual Replicator can increase storage capacity without disrupting operations on the NAS device. The online volume growth feature directs the operating system to update the size of a virtual disk without requiring a system restart. With online volume growth, storage capacity can be easily increased as required by users and applications, with zero downtime.

The VR online volume growth feature operates only on basic and virtual disks formatted with the New Technology File System (NTFS). It does not work with dynamic drives, disks that are not virtual disks, or virtual disks that are not formatted with NTFS.

Preparing for Online Volume Growth of Basic or Virtual Disks

Before growing a volume:

- Make a reliable backup copy of the data.
- Plan for new storage. For example, adding capacity can increase the time required to perform backups.
- Make sure the pools have adequate free space. When the size of a virtual disk is increased, it consumes more space in the pool and reduces the amount of free space that can be used for snapshots.

NOTE: Online volume growth operates only on virtual disks formatted with New Technology File System (NTFS). If the file system is not NTFS, an error message is displayed instructing the administrator to cancel.

NOTE: The amount of free space shown is the amount of space available for use by virtual disks. The pool space reserved for snapshots is not included. Online volume growth does not grow the volume beyond the amount of free pool space available. If an attempt is made to grow the volume beyond the available space, an error occurs.

Adding Additional Storage Units to the Pool

Before growing a volume, additional storage units may need to be added to the pool.



CAUTION: Virtual Replicator does not support changing the size of a logical unit number (LUN) after it has been incorporated into a pool. An existing RAID array must be modified or a new RAID array must be added. Changing the size of the LUN results in data loss.

If additional LUNs were added, but they are not displayed in the user interface, use Windows Device Manager to rescan for the new drive.

The following procedure explains how to grow a virtual disk volume using the WebUI.

NOTE: Drive volume growth must not be performed at the same time as other major drive management operations.

Growing the Disk

To increase the capacity of a virtual disk:

1. Add the additional LUNs to the desired pool, using the procedures described in the “Pool Management” section.
2. From the **Manage Virtual Disks** dialog box, select the virtual disk to work with and click **Grow**. The Grow Virtual Disk dialog box is displayed. Information about the selected virtual disk is included in the screen display.
3. In the data entry box, enter the amount of additional capacity (MB) to add to the virtual disk.
4. Click **OK**. The disk is updated and the Manage Virtual Disks dialog box is displayed again.

Backing Up (Creating a Scheduled Snapshot of) a Virtual Disk

To create a scheduled snapshot of a virtual disk:

1. From the **Manage Virtual Disks** dialog box, select the virtual disk to back up and then click **Backup**. The Welcome screen of the **Virtual Disk Backup** wizard is displayed.
2. Click **Next** to continue. The **Snapshot Information** screen is displayed, including the name (and node, if applicable) of the selected virtual disk.

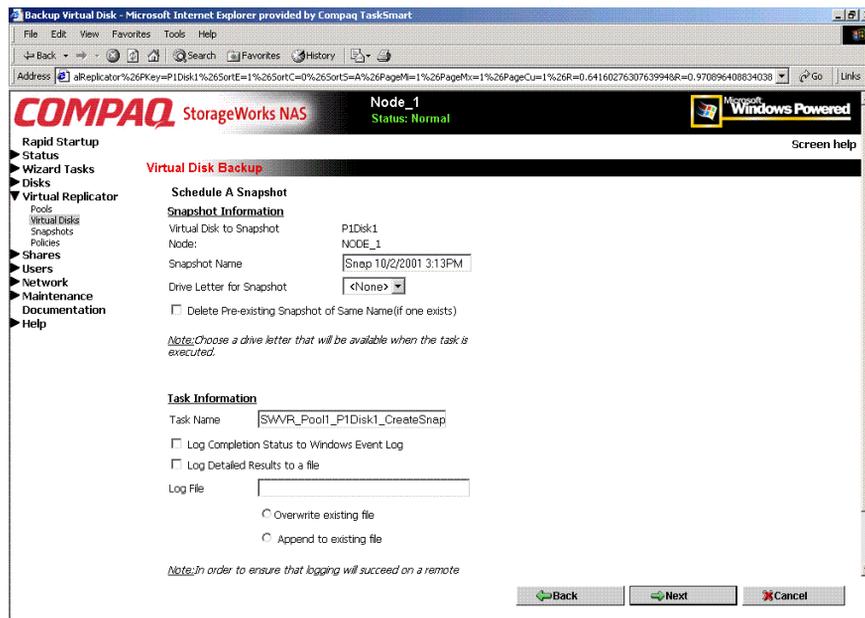


Figure 6-11: Schedule a Snapshot screen

3. Enter snapshot information, including:
 - A name for the snapshot
 - The drive letter to assign to the snapshot (optional)

- Indicate whether the system should delete any existing snapshots with the same name. If this option is not checked and a snapshot with the same name already exists, the new snapshot will not be created.
4. Enter task information.

The screen displays a default name for the task. Indicate whether the system will log the task into an audit file:

- Completion status can be posted to the Windows Event Log. Compaq recommends posting to the log to confirm that the task was begun and was completed.
 - Task details can be posted to a custom file. To post the details to a file, enter the filename and also indicate whether the system should append this information onto the end of the file or to overwrite information in the existing file.
5. Click **Next**. The **Schedule Information** screen is displayed.

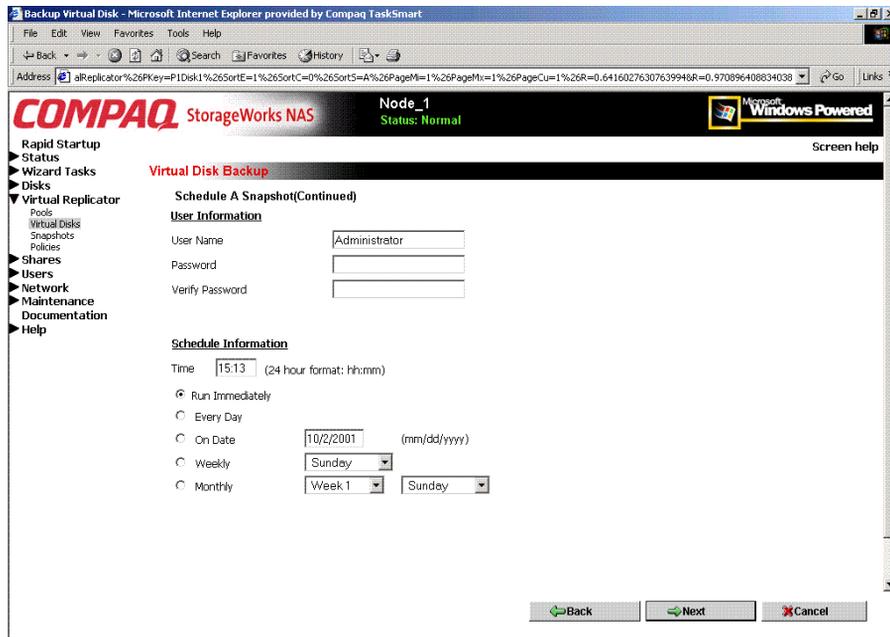


Figure 6-12: Schedule Information screen

6. Enter schedule information, including:
 - Time to create the snapshot (using a 24-hour clock)
 - When to create the snapshot. Options include:
 - Immediately
 - Daily
 - On a specific day
 - Weekly
 - Monthly.
7. Click **Next**. A summary screen is displayed. Figure 6-13 illustrates the Virtual Disk Backup summary screen.

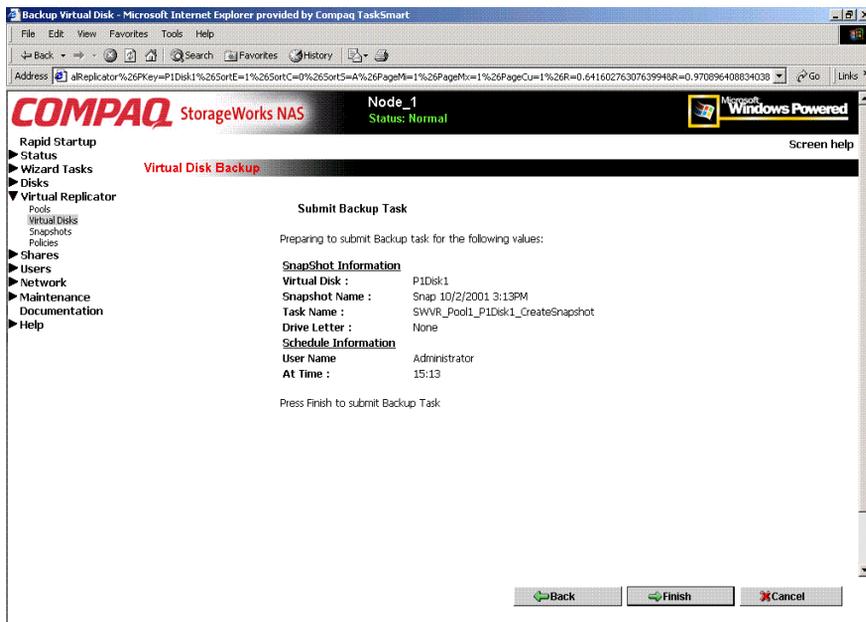


Figure 6-13: Virtual Disk Backup Summary screen

8. After verifying the accuracy of the schedule information, click **Finish**. The Manage Virtual Disks dialog box is displayed again.

Snapshot Management (Details)

Snapshots are managed through the WebUI, with the **Virtual Replicator** menu option. To create and manage snapshots, in the Virtual Replicator menu, select **Snapshots**.

The **Manage Snapshots** dialog box is displayed. All existing snapshots are displayed, including capacity and drive letter information for each snapshot.

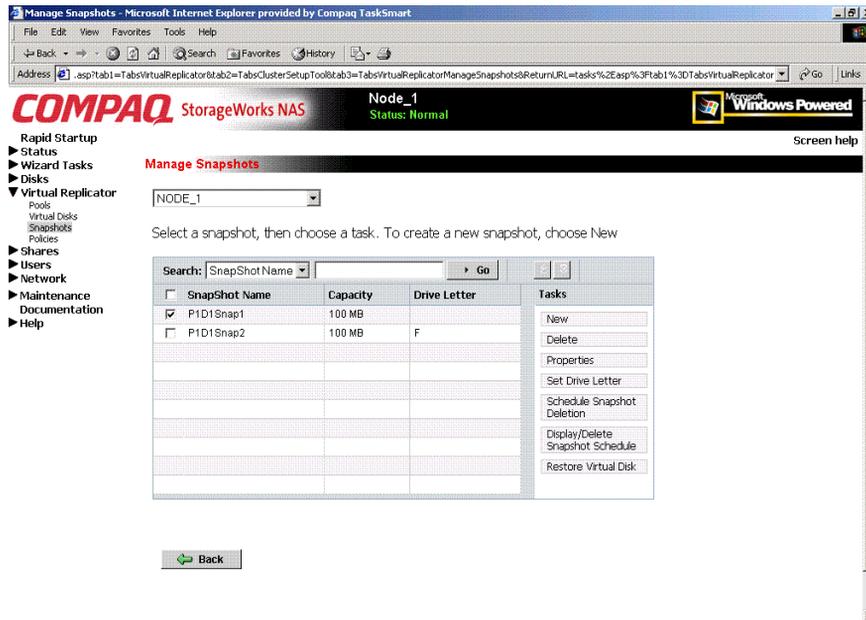


Figure 6-14: Manage Snapshots dialog box

Snapshot tasks include:

- Creating a new snapshot
- Deleting a snapshot
- Viewing and Modifying snapshot properties
- Setting the drive letter for a snapshot
- Creating a Snapshot Schedule (managed through the Virtual Disks option)
- Displaying and deleting snapshot schedules (managed through the virtual Disks option and the Snapshots option)
- Scheduling snapshot deletions (managed through the Virtual Disks option)
- Restoring a virtual disk from a snapshot
- Enabling incremental backup support (managed through the Virtual Disks option)

Each of these tasks is discussed in the following sections.

Creating a New Snapshot

Snapshots are a routine maintenance tool for the administrator of the NAS device. Before creating a snapshot, the NAS device must have pools of space and one or more virtual disks established. Verify that there is adequate free pool space to support snapshots.

IMPORTANT: Snapshots are temporary in nature and are automatically deleted by the system if free pool space becomes critical. Snapshot data loss can occur. Snapshots are not a replacement for regular tape backups. See “Lifeguard Service” for more information.

After the pool or pools of space and one or more virtual disks are established, the process to create snapshots can begin.

1. From the **Manage Snapshots** dialog box, click **New**. The Welcome screen of the Create New Snapshot wizard is displayed. To continue, click **Next**. The next screen of the wizard is displayed.

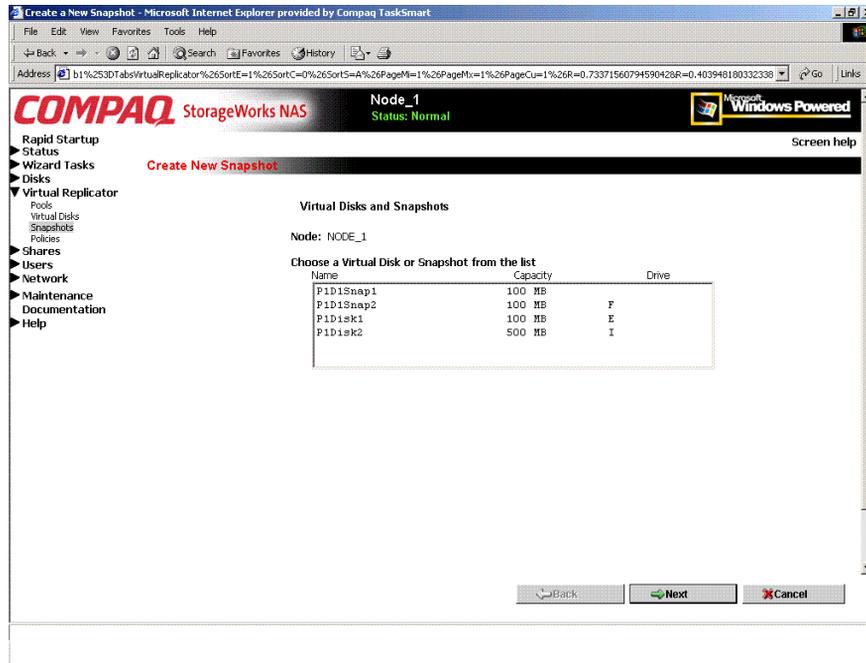


Figure 6-15: Virtual Disks and Snapshots screen

2. All existing virtual disks and snapshots are listed in the display box. Select the virtual disk or the snapshot to take a snapshot of, and then click **Next**. The New Snapshot information screen is displayed.

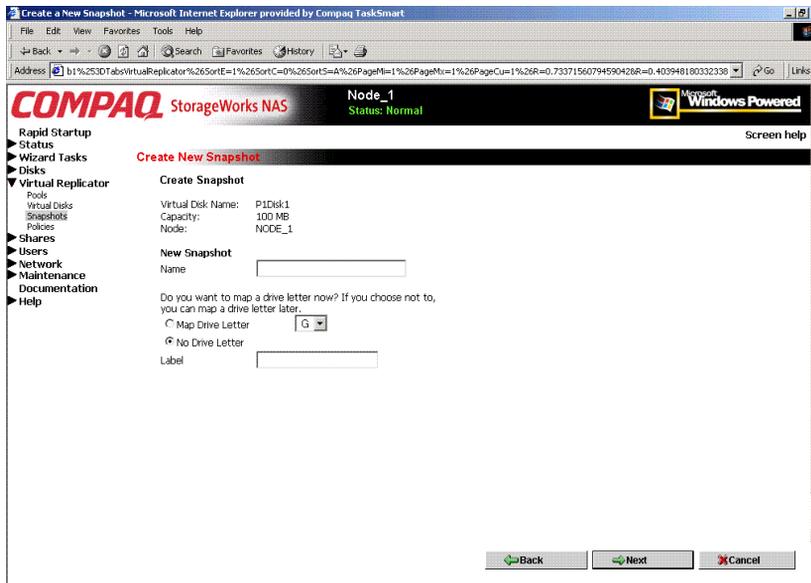


Figure 6-16: New Snapshot Information screen

3. The name and size of the selected item is displayed. Enter:

- Snapshot name
- Drive letter for the snapshot
- Label for the drive letter of the snapshot.

NOTE: Use appropriate snapshot and volume naming to identify each snapshot individually.

NOTE: It is not necessary to assign a drive letter to the snapshot until access to the snapshot is needed. If one snapshot per virtual disk will be created, it is advisable to assign a drive letter at this time. If a drive letter is not set now, it can be set at a later time.

4. Click **Next**. A completion screen is displayed. Click **Finish**.

Deleting a Snapshot

To delete a snapshot:

1. From the **Manage Snapshots** dialog box, select the snapshot to delete from the displayed list. Then click **Delete**.
2. A confirmation screen is displayed. After verifying that the displayed snapshot is the one that needs to be deleted, click **OK**.

Viewing Snapshot Properties

To view snapshot properties, from the **Manage Snapshots** dialog box, select a snapshot from the displayed list and then click **Properties**. The Snapshot Properties screen is displayed.

This interface is a summary display. To change any of these properties, use the appropriate action in the Manage Snapshots or Manage Virtual Disks dialog boxes.

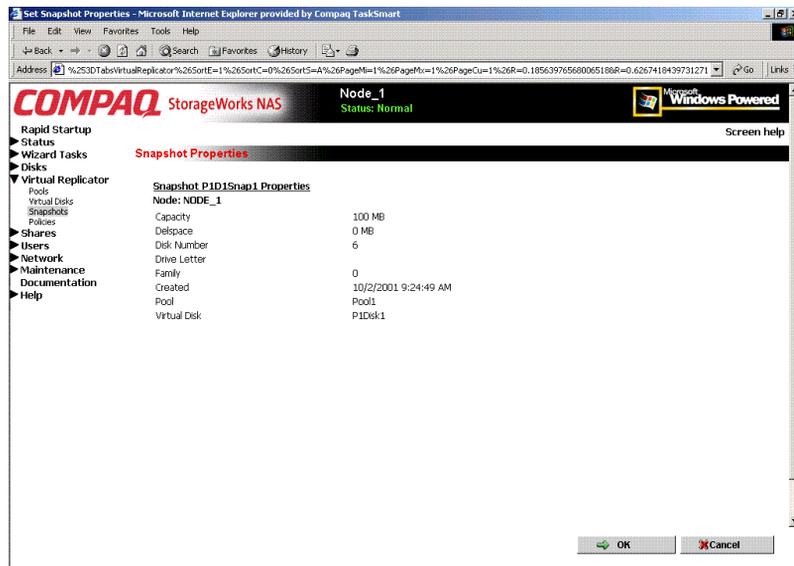


Figure 6-17: Snapshot Properties summary display

Setting the Drive Letter for a Snapshot

To set or change the drive letter for a snapshot:

1. From the **Manage Snapshots** dialog box, select the desired snapshot. Then, click **Set Drive Letter**. The Set Drive Letter dialog box is displayed.
2. Using the **Drive Letter** drop-down box, select a letter to assign to this snapshot, and then click **OK**. Only available, unused letters are displayed.

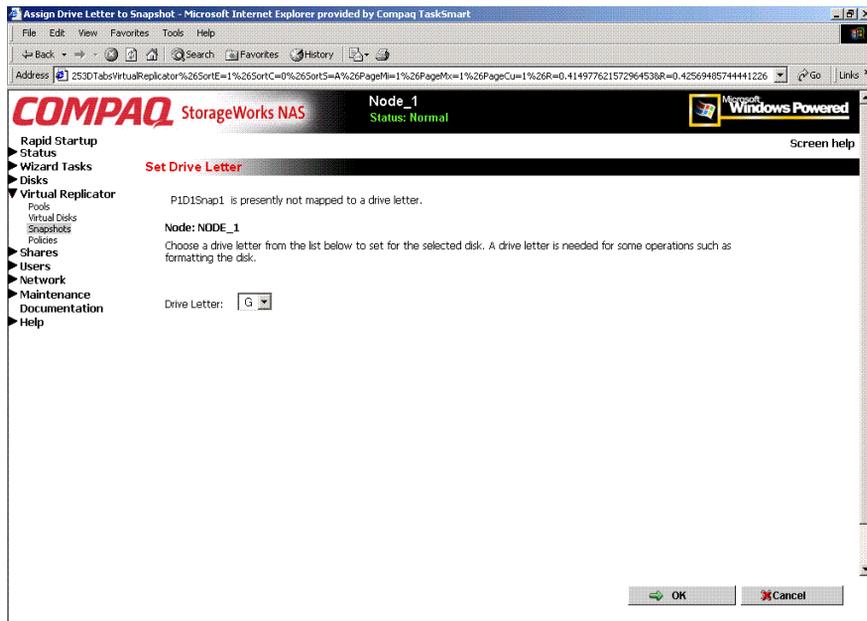


Figure 6-18: Set Drive Letter dialog box

Creating a Snapshot Schedule

A scheduled snapshot is a snapshot that that will be run at a specific date or time. These schedules can be recurring or a one-time event. To create a snapshot schedule for a virtual disk:

1. From the **Manage Virtual Disks** dialog box, select the virtual disk to work with, and click **Create Snapshot Schedule**. The Welcome screen of the Schedule a Virtual Disk Snapshot wizard is displayed.

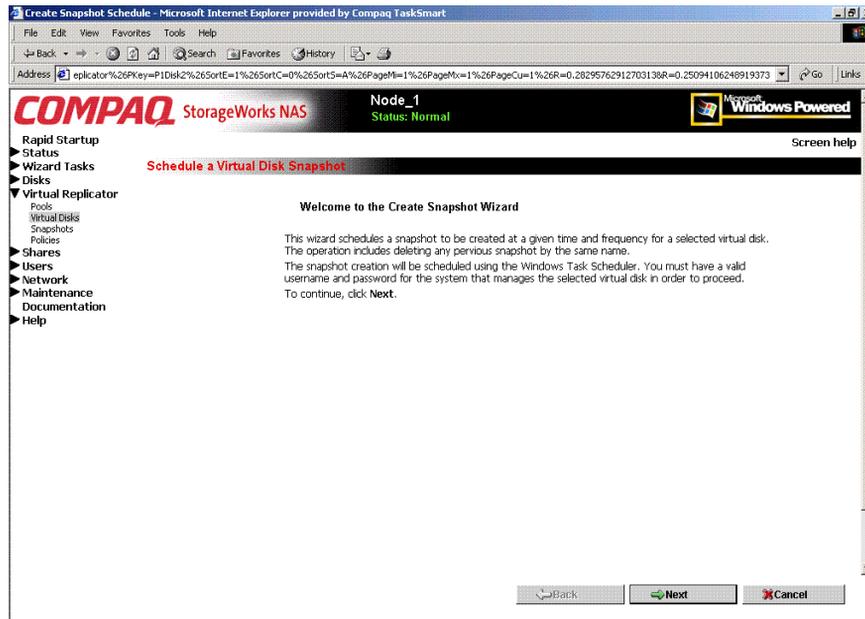


Figure 6-19: Snapshot Wizard Welcome screen

2. Click **Next**. The next screen of the wizard is displayed. The name of the selected virtual disk is displayed. In this screen, enter snapshot and task information.

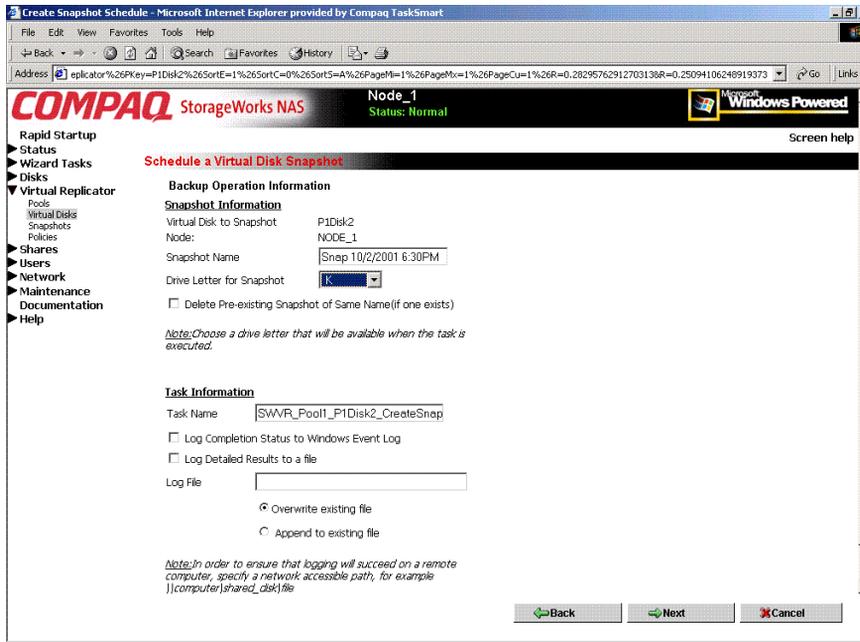


Figure 6-20: Snapshot and Task Information screen

3. Enter snapshot information, including:

- A name for the snapshot
- The letter to assign to the snapshot
- Indicate whether the system should delete any existing snapshots with the same name. If this option is not checked and a snapshot with the same name already exists, the new snapshot will not be created.

4. Enter task information.

The screen displays a default name for the task. Indicate whether the system should log the task into an audit file:

- Completion status can be posted to the Windows Event Log. Compaq recommends posting to the log to confirm that the task was begun and was completed.
- Task details can be posted to a custom file. To post the details to a file, enter the filename and also indicate whether the system should append this information onto the end of the file or overwrite information in an existing file.

5. Click **Next**. The **Schedule Information** screen is displayed.

COMPAQ StorageWorks NAS Node_1
Status: Normal

Schedule a Virtual Disk Snapshot

Log File Information

User Information

User Name: Administrator

Password: []

Verify Password: []

Schedule Information

Time: 18:30 (24 hour format: hh:mm)

Run Immediately

Every Day

On Date: 10/2/2001 (mm/dd/yyyy)

Weekly: Sunday

Monthly: Week 1, Sunday

Enable Task At Wizard Completion

Your Task Information is about to be entered in the Windows Task Scheduler. If the Check-box is checked, the task will be enabled to run.

Back Next Cancel

Figure 6-21: Schedule Information screen

6. Enter schedule information, including:
 - Time to create the snapshot (using a 24-hour clock)
 - When to create the snapshot
 - Options include:
 - Immediately
 - Daily
 - On a specific date
 - Weekly
 - Monthly
 - Whether to enable this schedule at this time.

To activate the snapshot schedule immediately, check the box. If Enable Task at Wizard Completion is not selected, this schedule will be stored, but will not be activated until you return to Windows Task Scheduler and enable the schedule.

Click **Next**. A summary screen is displayed
7. After verifying the accuracy of the schedule information in the summary screen, click **Finish**. The Manage Virtual Disks dialog box is displayed again.

Displaying and Deleting a Snapshot Schedule

Snapshot schedules can be displayed and deleted from two places in the WebUI:

- Through the **Virtual Disks** menu option
- Through the **Manage Snapshots** menu option

Displaying and Deleting Snapshot Schedules through the Virtual Disks Menu

To view or delete snapshot schedules from the Virtual Disks menu option:

1. From the **Manage Virtual Disks** dialog box, select the virtual disk whose snapshots you want to see and then click **Display/Delete Snapshot Schedule**. The Display/Delete Snapshot Schedule dialog box is displayed.

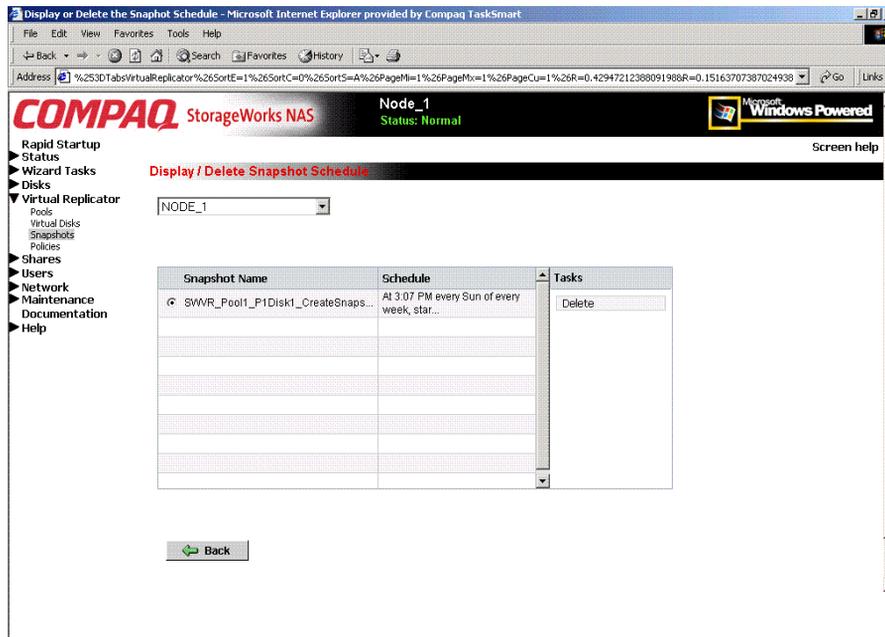


Figure 6-22: Display/Delete Snapshot Schedule dialog box

2. Select the desired schedules from the displayed list and then click **Delete**.
3. A confirmation screen is displayed. After verifying that the correct snapshot schedule was selected, click **OK**.

Displaying and Deleting Snapshot Schedules through the Manage Snapshot menu

To view and delete snapshot schedules from the Manage Snapshots menu option:

1. From the **Manage Snapshots** dialog box, select the snapshot whose schedule you want to view or delete and then click **Display/Delete Snapshot Schedule**. The Display/Delete Snapshot Schedule dialog box is displayed.
2. Select the schedule to delete and then click **Delete**.

A confirmation screen is displayed. View the display to confirm that the desired schedule was selected.

3. To proceed with the deletion, click **OK**.

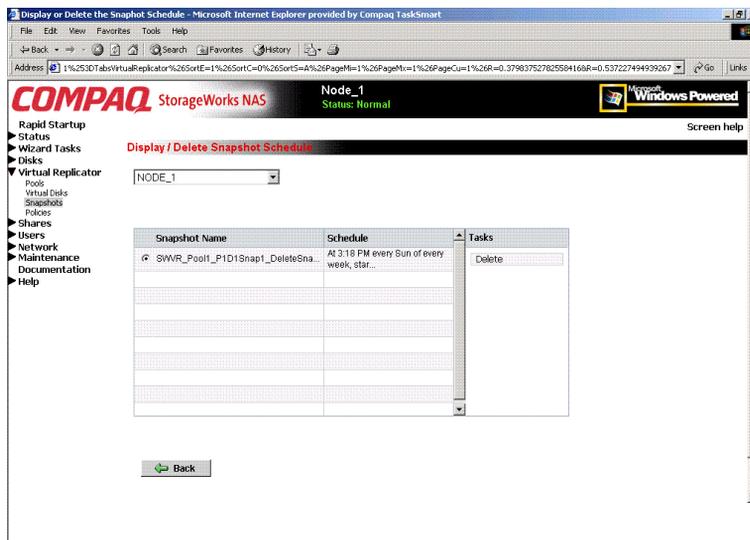


Figure 6-23: Display/Delete the Snapshot Schedule wizard screen

Scheduling Snapshot Deletions

To schedule the deletion of a snapshot:

1. From the **Manage Snapshots** dialog box, select the snapshot to delete from the display list. Then click **Schedule Snapshot Deletion**.

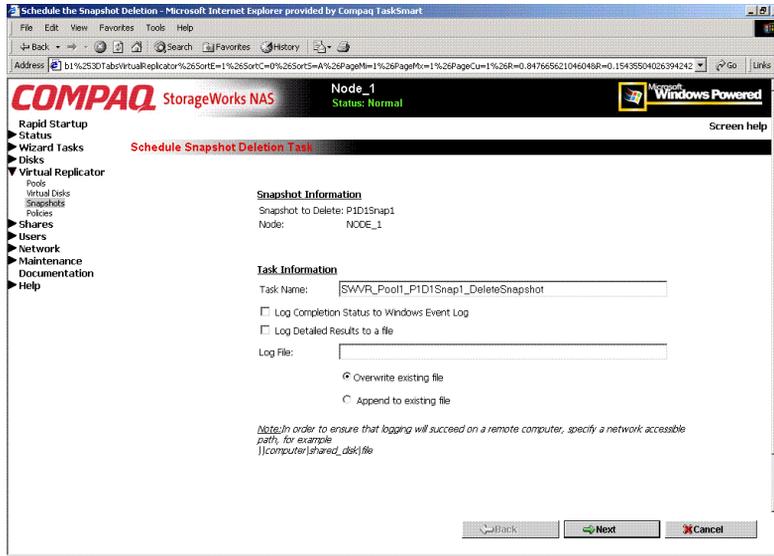


Figure 6-24: Schedule NAS Snapshot Deletion Task dialog box

Snapshot information, including the snapshot name, is displayed.

2. Enter task information, including whether the system should log the task into an audit file:
 - Completion status can be posted to the Windows Event Log.
 - Task details can be posted to a custom file. To post the details to a custom file, enter the filename and indicate whether the system should append this information onto the end of the file or overwrite the existing file.
3. Click **Next**. The **Schedule Information** screen is displayed.

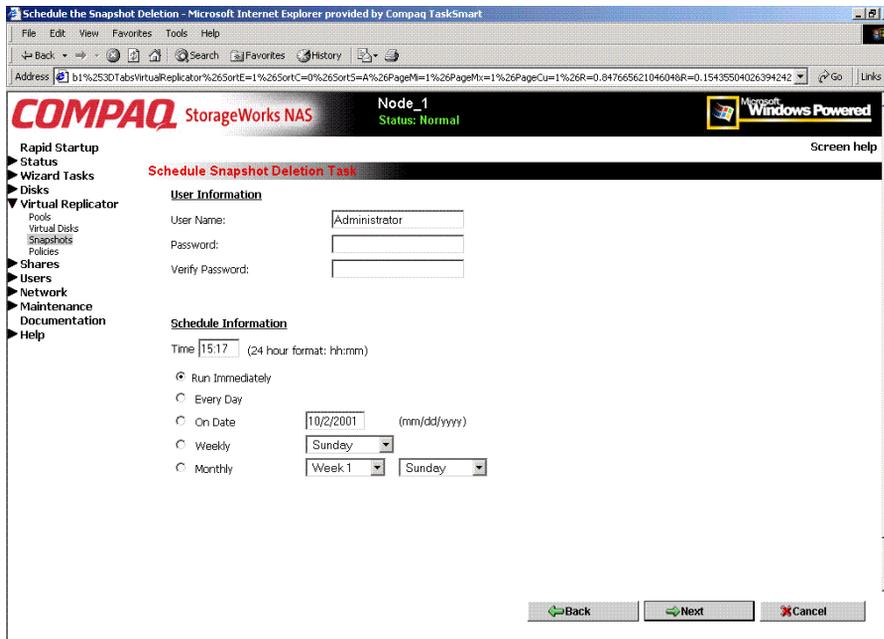


Figure 6-25: User Information and Schedule Information screen

4. Enter schedule information, including:
 - Time to delete the snapshot (using a 24-hour clock)
 - When to delete the snapshotOptions include:
 - Immediately
 - Daily
 - On a specific day
 - Weekly
 - Monthly
5. Click **Next**. A completion screen is displayed. Click **Finish**.

Restoring a Virtual Disk from a Snapshot

Snapshots can be used to return a virtual disk to a previous state. If the virtual disk or the data is corrupted, the snapshot may not be able to recover all necessary blocks and the disk must be restored using normal tape backup.

NOTE: Snapshots are not a replacement for normal tape backup.

This process is also called SnapBack. SnapBack uses the contents of a snapshot to restore the contents of a parent virtual disk. Snapshots can be used to recover data without performing a file copy to an intermediate location and consuming disk space in the process. The restoration is automatic and there is no need to map and remap drives, copy files, or delete snapshots or virtual disks. SnapBack performs all of these steps. During the operation, the drive letters of the virtual disk and snapshot are momentarily unavailable. As part of the process, the source snapshot is retained, but all snapshots that are older than the source for the SnapBack are deleted. Deleting earlier snapshots ensures the integrity of the data on the restored virtual disk.



CAUTION: When restoring a virtual disk from a snapshot, there must be sufficient free pool space; otherwise, a pool-full condition may occur and Lifeguard may delete the snapshot from which it is restoring. If the only snapshot is deleted, the virtual disk must be restored from tape backup. See “VR Lifeguard”



CAUTION: Before performing the SnapBack operation, verify that earlier snapshots, which will be deleted in the process, are not needed.



CAUTION: If the SnapBack operation is interrupted, the virtual disk will not be fully restored and will be in an unpredictable state. Compaq recommends restarting the SnapBack restore process as soon as possible.

IMPORTANT: The SnapBack process can be lengthy, depending on a variety of factors, such as the amount of data on the snapshot, the number of storage units in the pool, and hardware configuration.

The Restore From Snapshot wizard schedules the recreation of a virtual disk from a snapshot. If specified, upon successful completion of the schedule, the previous virtual disk and snapshot are both deleted.

Virtual disks can be restored from a snapshot through two access points on the WebUI:

- Through the Virtual Disks menu
- Through the Snapshots menu

Each of these methods is discussed in the following sections of this chapter.

Preparing to Restore a Virtual Disk from a Snapshot

Before executing the restore:

1. Remove all shares associated with the virtual disk that is being restored.
2. Delete all snapshots of the virtual disk that is being recovered, except for the snapshot that will be used for the recovery. Additionally, to free up pool space, delete any unnecessary snapshots of other virtual disks in the same pool.
3. Make sure there is enough free pool space to accomplish the recovery and if necessary, add logical disks to increase pool free space.

Typically, pool free space equivalent to the current size of the virtual disk that is being recovered is needed, plus the space currently used by the snapshot being used for the recovery. To determine the amount of space used by this snapshot, access the Snapshot Manager, and then click the virtual disk listing. Locate the snapshot name in the right pane. The amount of the space used by the snapshot is in the **Delspace** column.

Restoring a Virtual Disk from a Snapshot through the Snapshots Menu

To use the Snapshot menu to restore a virtual disk from a snapshot:

1. From the **Manage Snapshots** dialog box, select the desired snapshot. Then click **Restore Virtual Disk**. The Restore Virtual Disk from Snapshot screen is displayed.

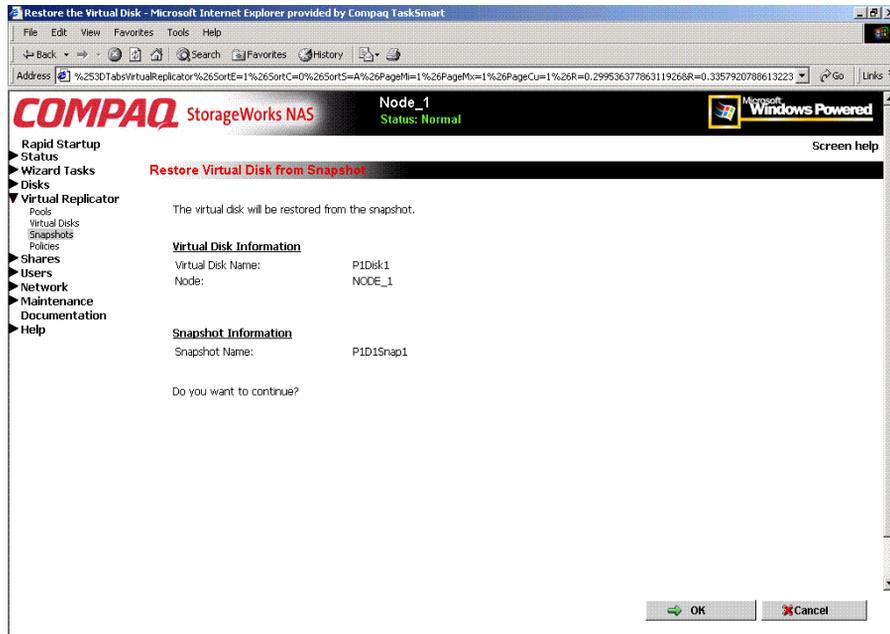


Figure 6-26: Restore Virtual Disk from Snapshot screen

2. Verify the accuracy of the data on the screen.

NOTE: There is no additional warning prompt.

3. To proceed, click **OK**.

The disk is restored, and the Manage Snapshots dialog box is displayed again.

Restoring a Virtual Disk from a Snapshot through the Virtual Disks Menu

To use the Virtual Disks menu to restore a virtual disk from a snapshot:

1. From the **Manage Virtual Disks** dialog box, select the desired virtual disk and then click **Restore Virtual Disk**.

The Restore Virtual Disk from Snapshot screen is displayed.

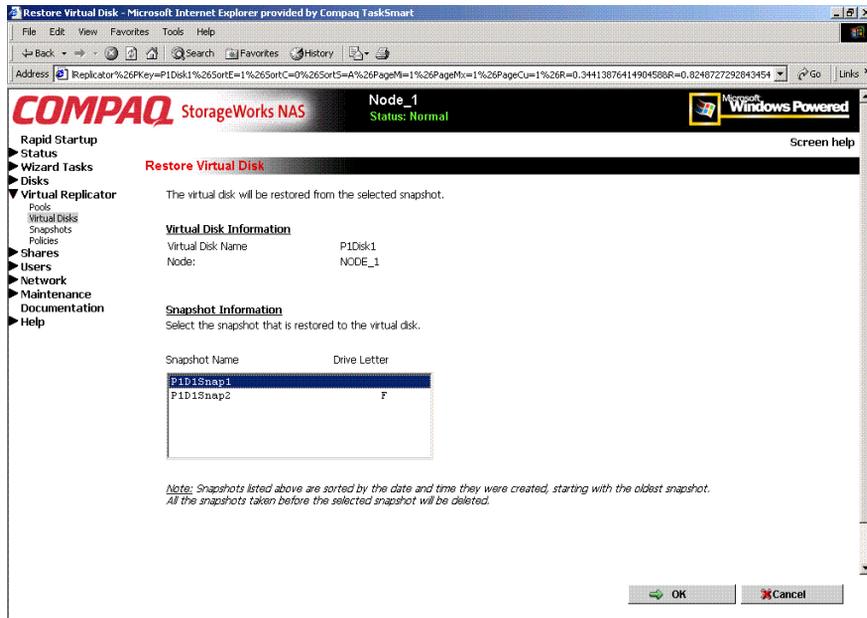


Figure 6-27: Restore Virtual Disk screen

2. In the snapshot box, select the snapshot to use.
3. Verify the accuracy of the data on the screen.

NOTE: There is no additional warning prompt.

4. To proceed, click **OK**.

The disk is restored, and the Manage Virtual Disks dialog box is displayed again.

Enabling Incremental Backup Support

Some backup utilities use an archive bit to indicate whether a file has been backed up. The Incremental Backup Support feature allows normal incremental backups by turning off the archive bits of the files on the virtual disk. The next snapshots that are taken will then have the archive bits correctly set for incremental backups.

1. From the **Manage Virtual Disks** dialog box, select the desired virtual disks, and click **Enable Incremental Backup Support**. The Welcome screen of the Enable Incremental Backup Support wizard is displayed.
2. Click **Next** to continue. The next screen of the wizard is displayed.

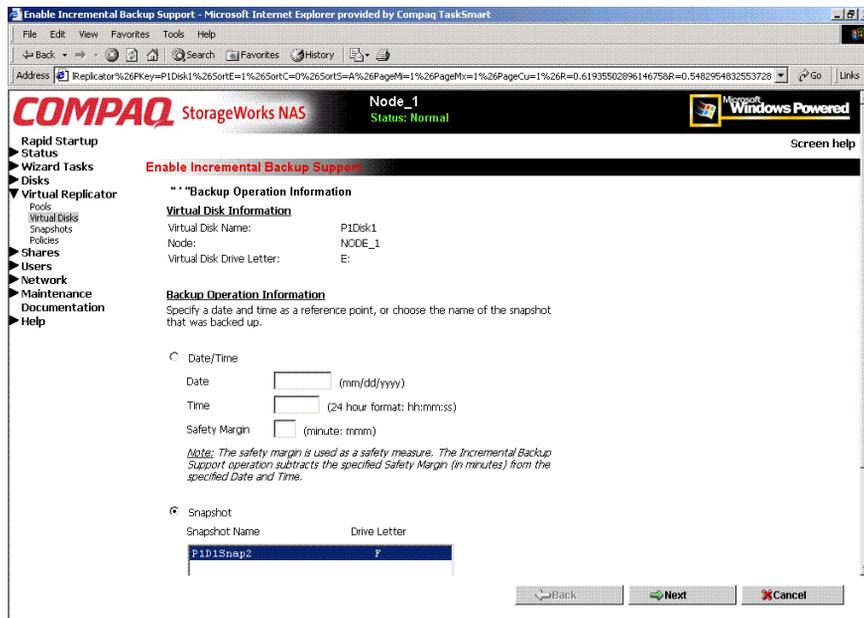


Figure 6-28: Backup Operation Information screen

3. Enter backup operation information. Select one of the following reference points:
 - Date and time
 - Snapshot name

4. Click **Next**. The **Log File Information** screen is displayed.

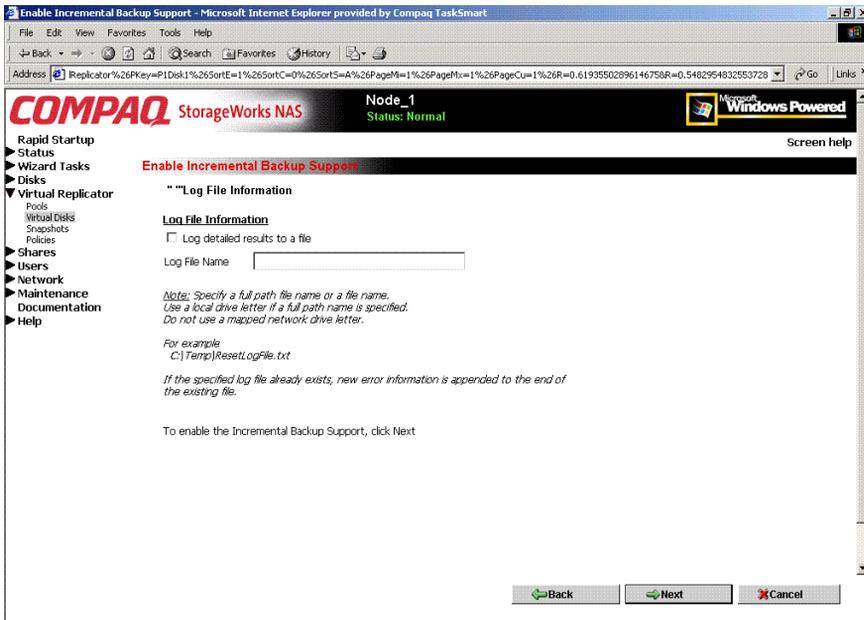


Figure 6-29: Log File Information Screen

5. To log results about the incremental backups, click **Log detailed results to a file**, and enter the file name.
6. Click **Next**. A confirmation screen is displayed.
7. After verifying the accuracy of the incremental backup information, click **Finish**.

Global Pool Policy Settings

Virtual Replicator parameters are preset on the NAS device, with one primary set of parameters controlling all pools. If a pool needs policy settings that are unique to it, the administrator can customize the policy settings for that individual pool. The global set of parameters continues to manage all pools that do not have customized policies.

This menu option allows the administrator to customize the global set of parameters. Only experienced users should modify these global and pool-specific policy settings.



CAUTION: Use caution when changing these parameters. Some choices can result in system inefficiencies, difficulties, and data loss.

This section discusses changing the global pool policies. To change the pool policies for a specific pool, see the “Setting Pool Policies for a Specific Pool” section earlier in this chapter.

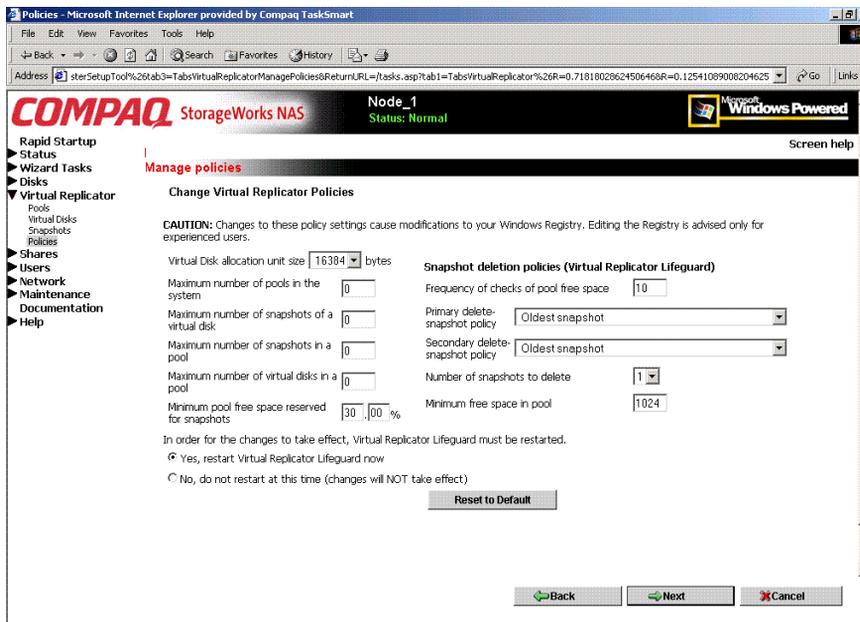


Figure 6-30: Manage Policies dialog box

1. From the **VR Main menu**, select **Policies**. An introductory **Manage policies** screen is displayed. Click **Next** to continue. The next Manage Policies dialog box is displayed.
2. In the **Manage policies** dialog box, enter the new parameter settings.

3. After all settings are entered, click **Next** and then click **Finish**.

Table 6-2: Pool Policy Default Settings (duplicated table)

Policy	Default Setting
Virtual disk allocation unit size	16 KB
Maximum number of pools in system	
Maximum number of snapshots of a virtual disk	
Maximum number of snapshots in a pool	
Maximum number of virtual disks in a pool	8
Maximum pool free space reserved for snapshots	10%
Restart Lifeguard now or later	Now
Frequency of pool free space checks	10 seconds
Primary delete snapshot policy	Oldest snapshot first
Secondary delete snapshot policy	Oldest snapshot first
Number of snapshots to delete	1
Minimum free space in pool	1024 KB

Namespace Recovery

Namespace Recovery is a utility that can be used whenever it may be necessary to re-detect the virtual storage configuration of the pools, virtual disks, and snapshots.

Figure 6-31 is an example of the Namespace Recovery screen.

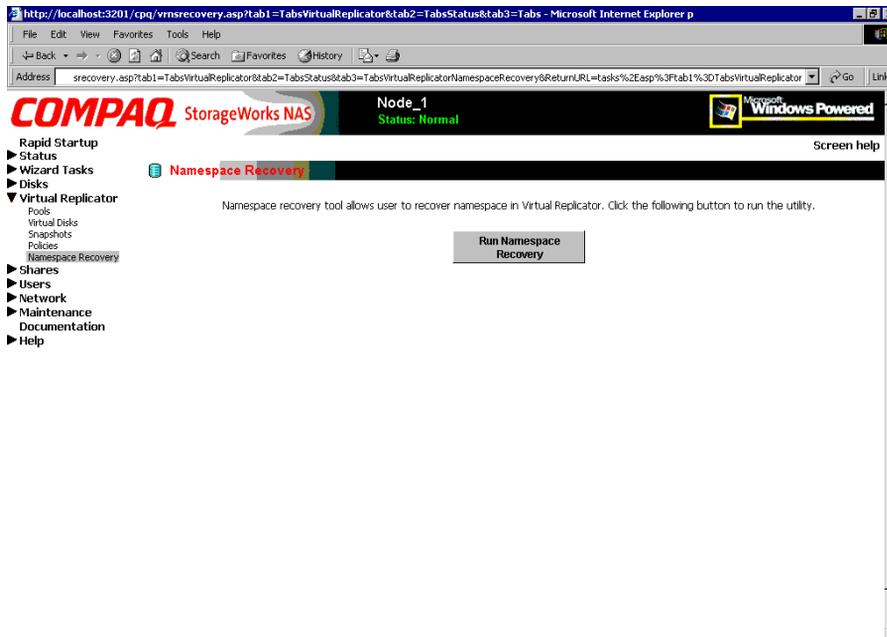


Figure 6-31: Namespace Recovery screen

Drive Quotas

Drive quotas let administrators control the allocation of drive space to individual users or groups of users. When quotas are enabled and properly configured, it is impossible for one person or group to consume all of the available space on a disk.

When quotas are enabled on a volume that already contains files, the system calculates the drive space used by all users on the volume. The quota limit and warning level are then applied to all current users. Administrators can then modify quotas as needed. By enabling and then disabling quotas, administrators take advantage of the auditing capabilities provided by quotas, without reducing server performance.

Usage quotas can be managed through two interfaces:

- Managing Quotas Using the Quota Management Wizard
- Managing Quotas (Details)

Managing Quotas Using the Quota Management Wizard

1. From the WebUI, select **Wizard Tasks**, and **User Management**. The Welcome screen of the wizard is displayed. Click **Next** to continue.
2. In the primary menu of the User Management wizard screen, select **User Quota**.
3. In the quota subscreen, existing volumes and folders are displayed. Navigate to the desired folder and then select it.
4. A sub-screen of the wizard is displayed, listing several options:
 - *To create new quota settings*, use the display box to navigate to the volume that needs a quota set up. Then, enter the quota information data fields. After all information is entered, click **Create New Quota Entry**.

A user list is displayed. Select the user to monitor, and then click **Next**.

- *To delete a quota*, navigate to the desired user and click **Delete**.
- *To modify information for a quota*, use the display box to navigate to the desired volume or user. Enter the new quota information, and then click **Save Settings**.

Managing Quotas (Details)

Managing quotas includes:

- Enabling and disabling quota management on a virtual disk
- Creating new quota entries for a user or group
- Deleting quota entries for a user or group
- Modifying quota entries for a user or group

Each of these tasks is discussed in the following sections.

Quota management tasks are performed from the **Disks, Disk Quota** selection from the WebUI menu. Figure 6-34 is an illustration of the disk quota dialog box.

NOTE: If the volume is not formatted with the NTFS file system, or if you are not a member of the administrators group, the Disk Quota option is not displayed (not accessible).

NOTE: For more information about quotas, refer to online help for NAS device quota help.

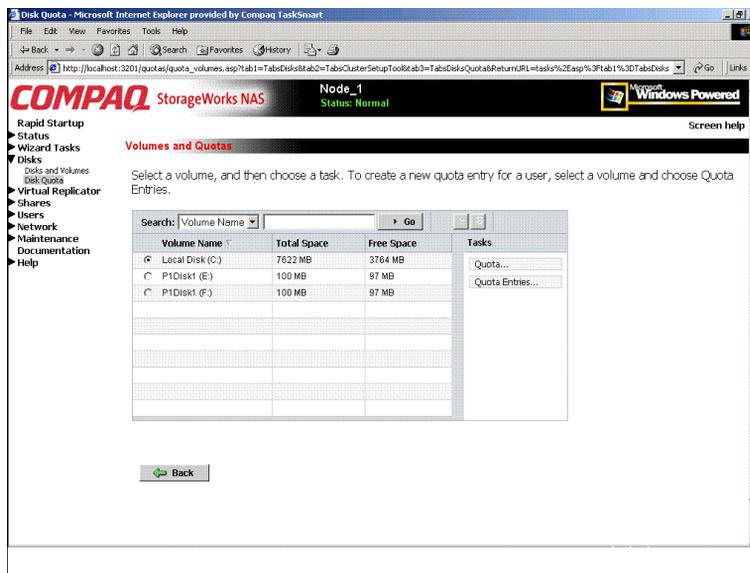


Figure 6-32: Disk Quota dialog box

Enabling and Disabling Quota Management on a Virtual Disk

To enable drive quotas:

1. From the WebUI, select **Disks**, **Disk Quota**. From the **Volumes and Quotas** dialog box, select a volume, and then click **Quota**. The **Default Quota** dialog box for the specified volume is displayed.

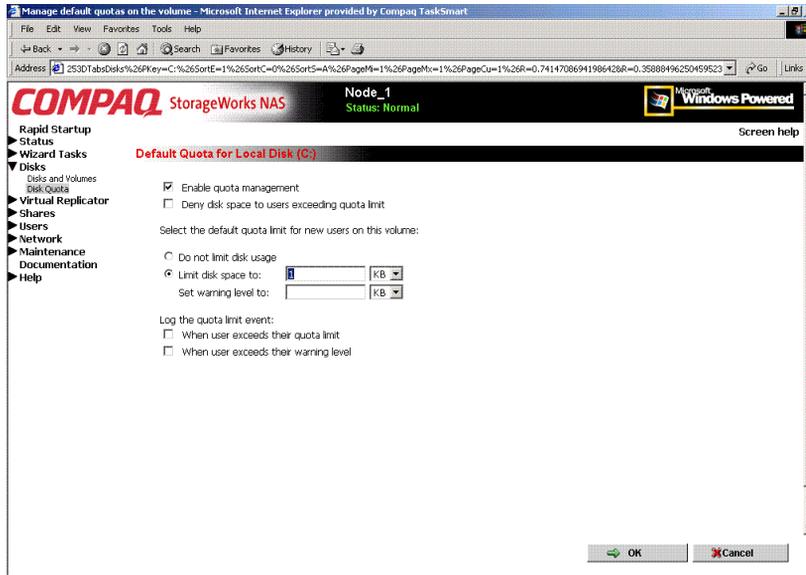


Figure 6-33: Default Quota Dialog box

2. To enable quotas on the selected disk, select **Enable quota management**. Complete the additional data fields on the screen, including disk space and warning level limits and auditing settings.
3. To disable quotas on the selected disk, de-select **Enable quota management**.
4. After completed all field entries, click **OK**. The Volume and Quotas dialog box is displayed again.

Creating New Quota Entries for a User or Group

To create new quotas for a user or group:

1. From the WebUI, select **Disks**, **Disk Quotas**. In the **Volumes and Quotas** dialog box, select a volume and then click **Quota Entries**. The **Quota Entries** dialog box is displayed.

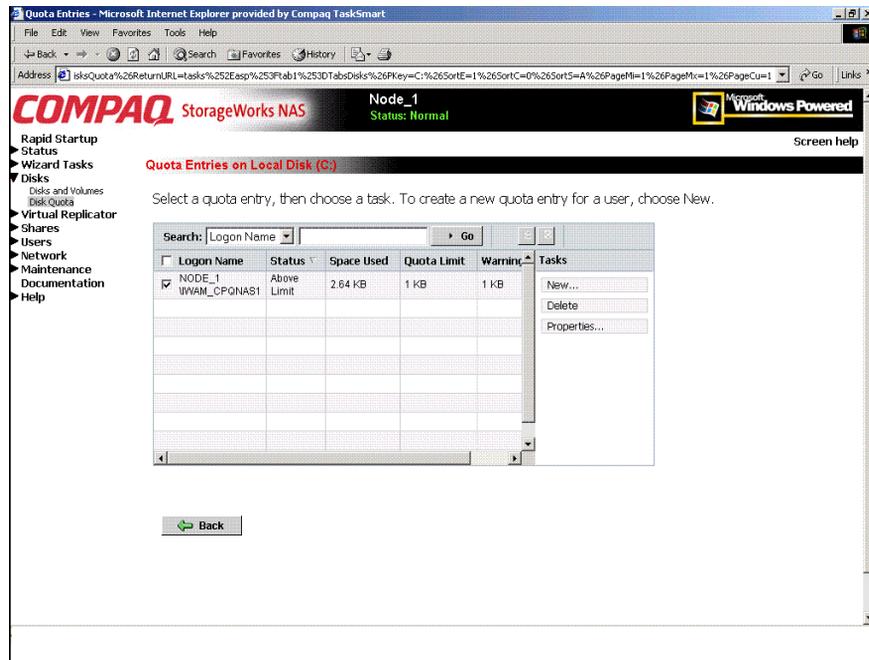


Figure 6-34: Quota Entries dialog box

2. All users and groups with established quotas are displayed. To create a new quota for a user or group, click **New**. The **New Quota Entry** dialog box is displayed.

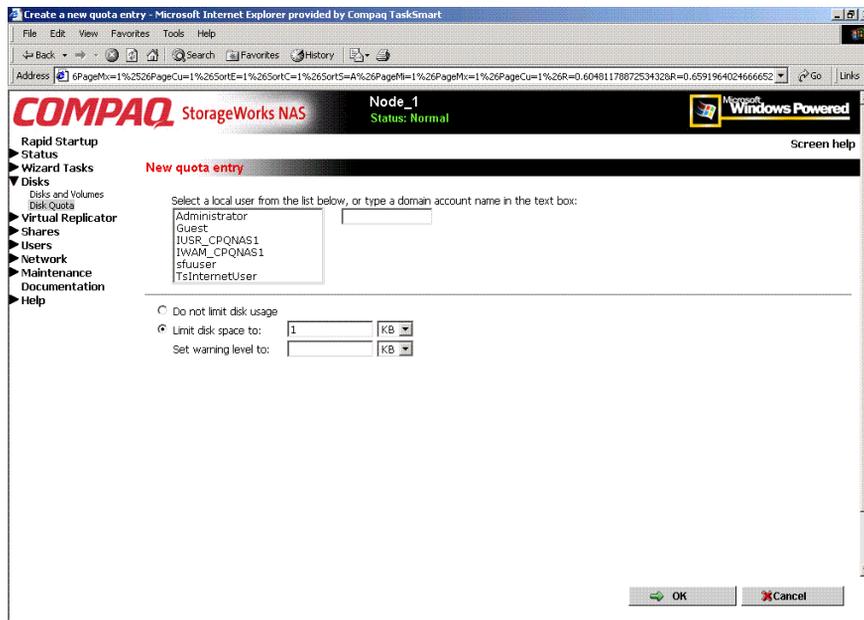


Figure 6-35: New Quota Entry dialog box

3. Indicate the user that the quota is for. For local users and groups, select the desired user from the Select a local user box. For users on the domain, enter the user's domain account name in the indicated box.
4. Enter a disk space limit.
5. Verify the accuracy of the field entries, and then click **OK**. The Quota Entries dialog box is displayed again.

Deleting Quota Entries for a User or Group

To delete quotas for a user or group:

1. From the WebUI, select **Disks, Disk Quotas**. In the Volumes and Quotas dialog box, select a volume and then click **Quota Entries**. The Quota Entries dialog box is displayed.
2. All users and groups with established quotas are displayed. To delete a quota for a user or group, click **Delete**. A verification dialog box is displayed.
3. Verify that this is the correct user, and then click **OK**. The Quota Entries dialog box is displayed again.

Modifying Quota Entries for a User or Group

Usage limit parameters for a user's quota can be changed. To modify these quota settings for a user:

1. From the WebUI, select **Disks, Disk Quotas**. In the Volumes and Quotas dialog box, select a volume and then click **Quota Entries**. The Quota Entries dialog box is displayed.
2. All users and groups with established quotas are displayed. To modify quota for a user or group, select a user, and then click **Properties**. The **Quota Entry** dialog box for that user is displayed.

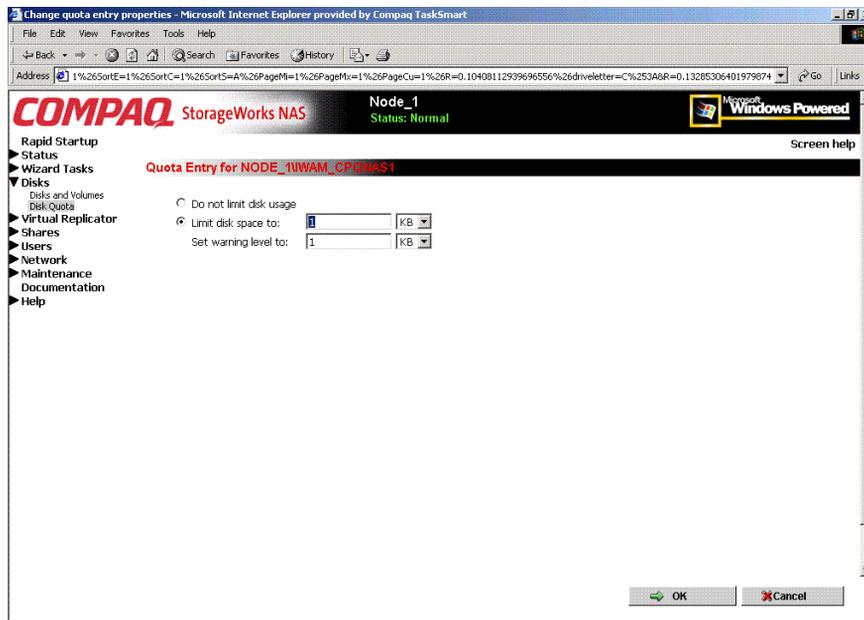


Figure 6-36: Quota Entry dialog box for a user

3. Enter the new disk limit information, and then click **OK**. The Quota Entries dialog box is displayed again.

User and Group Management

The *StorageWorks* NAS B2000 supports a variety of file sharing protocols for file access over a network, including:

- Common Internet File System (CIFS)
- Network File System (NFS)
- Novell Core Protocol (NCP)
- AppleTalk (AFP)

Access to shares requires a network logon (username and password). It follows that a fundamental part of managing shares involves managing the users and groups that have access.

There are two system environments for users and groups: workgroup and domain. Because users and groups in a domain environment are managed through standard Windows NT or Windows 2000 domain administration methods, this document discusses only local users and groups, which are stored and managed on the NAS device. For information on managing users and groups on a domain, refer to the domain documentation.

The following topics are addressed in this chapter:

- Domain Compared to Workgroup Environments
- User and Group Name Planning
 - Managing User Names
 - Managing Group Names
- Workgroup User and Group Management
 - Managing Local Users and Groups using the Wizard
 - Managing Local Users (Details)
 - Managing Local Groups (Details)

Domain Compared to Workgroup Environments

NAS B2000 devices can be deployed in workgroup or domain environments. When in a domain environment, the server is a member of the domain. The domain controller is a repository of accounts and account access for the NAS B2000. Client machines are also members of the domain, and users log on to the domain through their Windows clients. The domain controller also administers user accounts and appropriate access levels to resources that are a part of the domain.

In a CIFS environment, when mapping a network drive or a client machine, a user sends a logon credential to the server. This credential includes the username, password, and if appropriate, domain information. Using the credential, the server authenticates and provides the corresponding access to the user.

When a NAS B2000 is deployed into a workgroup environment, all user and group account access permissions to file resources are stored locally on the server.

By contrast, when a NAS B2000 is deployed into a domain environment it uses the account database from the domain controller, with user and group accounts stored outside the server. The server integrates with the domain controller infrastructure.

NOTE: The NAS B2000 cannot act as a domain controller for other servers on the network. If user and group account information is stored locally, those accounts may be used only to authenticate logons to the NAS B2000, resulting in a workgroup configuration.

Administering users and groups in a domain environment is similar in a mechanical sense to administering them in a workgroup environment. If using an Active Directory domain controller, the Computer Management tool allows for adding, modifying, and removing users in the same context as in a workgroup environment. The concepts, however, are very different.

Additional information about planning for domain environments can be found at the Compaq *ActiveAnswers*[™] website in the eBusiness Infrastructure solutions area:

www.compaq.com/ActiveAnswers/

The configuration of the domain controller is reflected on the NAS B2000 because it obtains user account information from the domain controller when deployed in a domain environment. As mentioned previously, the server cannot act as a domain controller itself.

User and Group Name Planning

Effective user and group management is dependent upon how well the user and group names are organized. Administrators typically create a small number of groups on the network and then assign users to the appropriate group or groups. File system and share permissions can then be applied at the group level, rather than at the user level. If the number of groups is small, assigning the appropriate permissions to selected group, or groups, is more efficient than assigning permissions to each user.

Although each organization has specific conventions, following general guidelines makes administration simpler and more efficient. Because CIFS is dependent on users and groups to grant appropriate access levels to file shares, CIFS administration benefits from a consistent user and group administration strategy.

Managing User Names

Usernames should reflect a logical relationship between the username and the person who uses the account. It is important that rules are established to ensure that usernames are:

- Systematic
- Easy to follow and implement
- Easy to remember

Using a combination of the user's first name, middle initial, and last name results in systematic usernames for every member of a particular organization. Common examples include:

- First initial followed by last name (jdoe for John Doe)
- First initial followed by middle initial and last name (jqpublic for John Q. Public)
- First name followed by last name, separated by a period (john.smith for John Smith)
- Last name followed by first initial (doej for Jane Doe)

Guidelines must be in place for instances when two users have the same initials or name. For example, a number can be added to the end of the username (jdoe1 and jdoe2).

Other conventions can be applied. Just ensure that conventions are both systematic and consistent.

Managing Group Names

Group management follows many of the same principles as user management.

It is recommended that group naming conventions be systematic and easy to understand. Make the group name convey some logical information about the function or purpose of the group. Table 7-1 provides examples of group names.

Group Name	Description
Administrators	All designated administrators on the server
Users	All standard server users
Power users	All standard server users requiring advanced access levels

Using tags is a helpful convention that indicates the specific access that a particular user has to a network resource. For example, if there is a data share on the device, the network administrator can create a “Data Users ROnly” group and a “Data Users RWrite” group to contain users that have read-only or read-write access on the share, respectively.

Workgroup User and Group Management

In a workgroup environment, users and groups are managed through the WebUI of the NAS B2000. Within the **Users** option, there are three choices:

- Managing users and groups using the wizard
- Managing local users
- Managing local groups

User and group administrative tasks include adding, deleting, and modifying user and group information. Managing local users and managing local groups are discussed in the following paragraphs.

Managing Users and Groups Using the Wizard

The **User Management** wizard is one of three major wizards included in the WebUI. The wizard contains most of the functionality of managing users and groups as the **Users** menu option of the WebUI. The wizard provides the easiest interface for adding, modifying, and deleting users and groups. Because the wizard is so easy to use, little information is included in this document. Online help is available.

For detailed information on managing users and groups, see “Managing Local Users (Details)” and “Managing Local Groups (Details),” later in this chapter.

For detailed information on managing NFS user name mapping, see the “UNIX File System Management” chapter.

1. From the WebUI, select **Wizard Tasks, User Management**. The Welcome screen of the wizard is displayed. Click **Next** to continue.
2. The primary menu of the User Management wizard is displayed. From this screen, select one of the following areas to work:
 - Local User
 - Local Group
 - User Quota
 - NFS Mapping

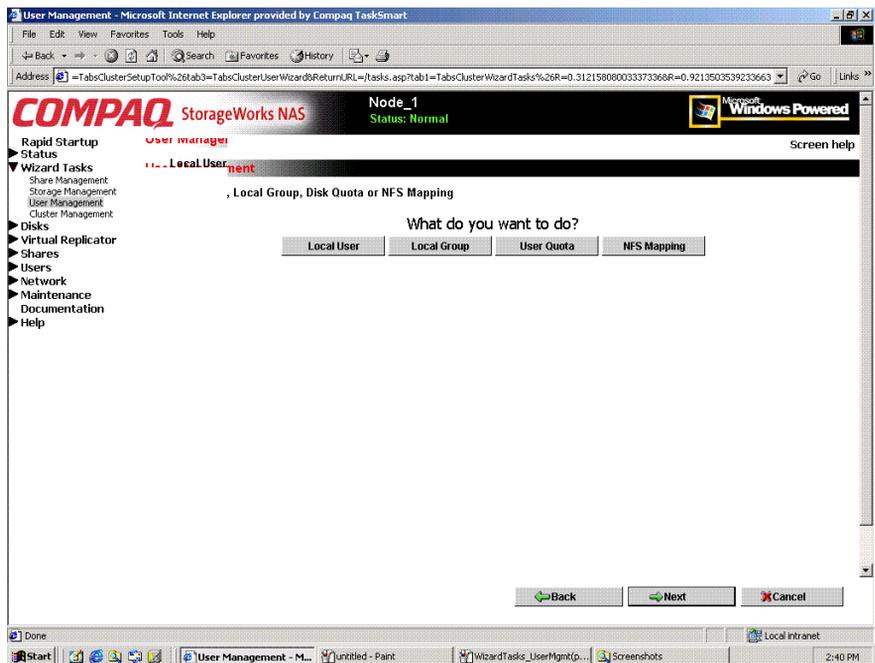


Figure 7-1: User Management wizard

Local Users

To manage user accounts through the wizard:

1. Click **Local User**. A subscreen is displayed, with the following options:
 - Create
 - Delete
 - Modify
 - Set Password
2. *To create a new user account*, enter the user data fields and then click **Create**.
3. *To delete a user account*, select a user from the displayed list and then click **Delete**.
4. *To modify information for a user*, select the user from the displayed list. User information is displayed. Enter the new values and then click **Modify**.
5. *To set the password for a user*, select the user from the displayed list. User information is displayed. Enter the new password and then click **Set Password**.

Local Groups

To manage local groups through the wizard:

1. Click **Local Group**. A sub-screen is displayed, with the following options:
 - Create
 - Delete
 - Modify
2. *To create a new group account*, enter the group information data fields and then click **Create**.
3. *To delete a group account*, select a group from the displayed list and then click **Delete**.
4. *To modify information for a group*, select the group from the displayed list, enter the new data, and then click **Modify**.

User Quota

To manage disk quotas for users and groups through the wizard:

NOTE: For detailed information on managing disk quotas, see the “Disk Quota” section of the “Virtual Storage Management” chapter.

1. Click **User Quota**.
2. Available volumes and folders are displayed. Navigate to the desired folder and then select it.

3. A sub-screen of the wizard is displayed, listing the following options:
 - Create new entry
 - Save settings
 - Delete
4. *To create quota settings*, use the display box to navigate to the volume that needs a quota set up. Then, select “Enable Quota Management” and enter the quota information data fields. After all information is entered, click **Create New Quota Entry**.

A user list is displayed. Select the user to monitor, and then click **Next**.
5. *To delete a quota*, navigate to the desired user and click **Delete**.
6. *To modify information for a quota*, use the display box to navigate to the desired volume or user. Enter the new quota information, and then click **Save Settings**.

NFS Mapping

To set up NFS user name mapping through the wizard:

1. Click **NFS Mapping**. A sub-screen of the wizard is displayed, listing four tabs of information, including:
 - General
 - Simple mapping
 - Explicit user mapping
 - Explicit group mapping
2. *In the General tab*, indicate whether NIS or password and group files are being used.
3. *In the Simple Mapping tab*, indicate whether simple mapping will be used in addition to explicit mapping.
4. *In the Explicit User Mapping tab*, link the Windows user and the UNIX users.
5. *In the Explicit Group Mapping tab*, link the Windows group and the UNIX group.

Managing Local Users (Details)

Managing users includes the following tasks:

- Adding a new user
- Deleting a user
- Setting a user password
- Modifying user properties

In the WebUI, under **Users**, **Local Users** is the **Local Users on Server Appliance** dialog box. All workgroup user administration tasks are performed in the Local Users dialog box.

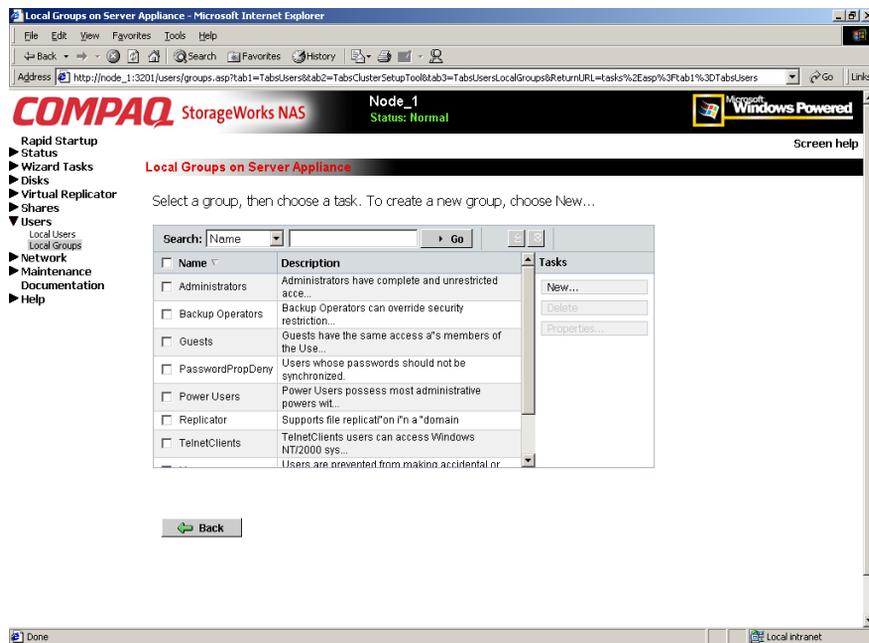


Figure 7-2: Local Users dialog box

All available options include: **New**, **Delete**, **Set a Password**, and **Properties**. When the Local Users dialog box is initially displayed, only the New option is available. After an existing user is selected, the additional actions are displayed. Each of these options is discussed in the following paragraphs.

Existing user records can be retrieved in one of two ways:

- By entering the user's **User Name** or **Full Name** in the **Search** fields to retrieve a specific user record. To redisplay the complete user list, space out the **Search** field.
- By selecting the user from the list of displayed users in the dialog box. The sort order of the display is controlled by clicking the **Name** field heading. The names are displayed in alphanumeric order or reverse alphanumeric order.

Adding a New User

To add a user:

1. From the **Local Users** dialog box, click **New**. The **Create New User** dialog box is displayed.

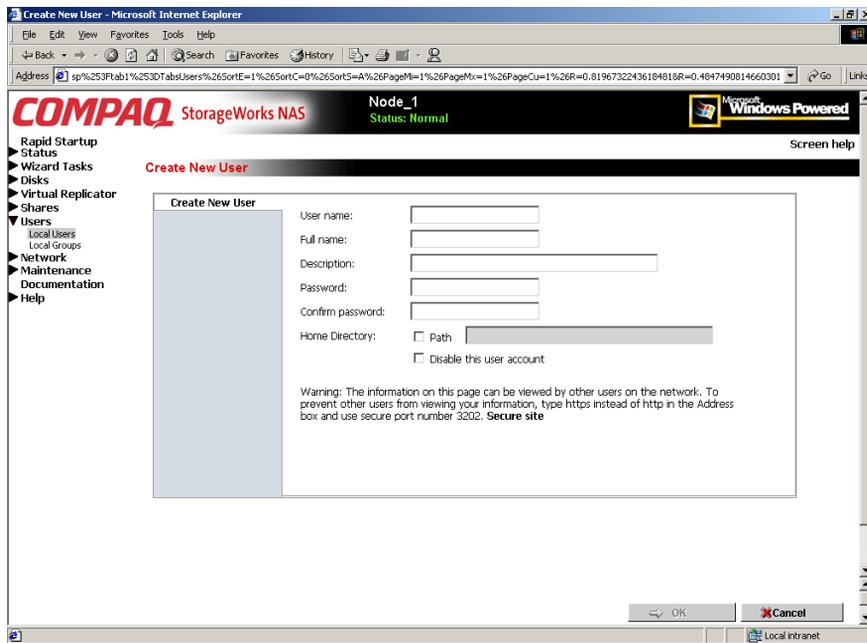


Figure 7-3: Create New User dialog box

2. Enter the user information and then click **OK**. The user is added and the **Local Users** dialog box is displayed again

Deleting a User

To delete a user:

1. In the **Local Users** dialog box, select the user to delete, and then click **Delete**.

The **Delete User** dialog box is displayed, including a warning note about deleting users.

2. To delete the user, click **OK**. The user is deleted and the Local Users dialog box is displayed again.

Modifying a User Password

Follow these steps to modify a user's password:

1. In the **Local Users** dialog box, select the user whose password needs to be changed. Then, click **Set a Password**.

The **Set Password** dialog box is displayed.

2. Enter the password and click **OK**. The Local Users dialog box is displayed again.

Modifying User Properties

To modify other user properties:

1. From the **Local Users** dialog box, select the user whose record needs to be modified. Then, click **Properties**.

The General information page of the **Properties** dialog box is displayed. Figure 7-4 is an illustration of the user Properties dialog box.

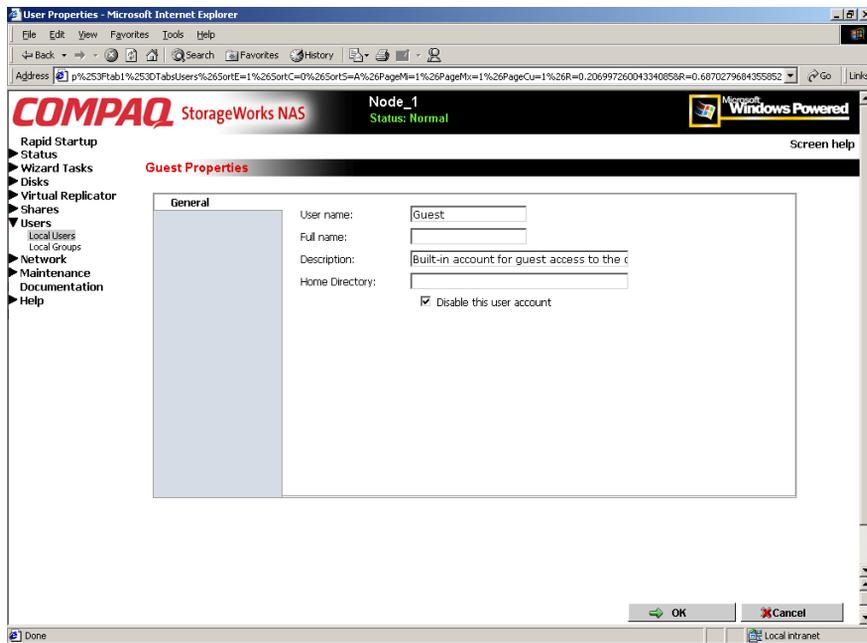


Figure 7-4: User Properties dialog box

2. The following information can be changed or set:
 - User name
 - Full name
 - Description
 - Home Directory
 - Disable this account
3. After completing the changes, click **OK**. The Local Users dialog box is displayed again.

Managing Local Groups (Details)

Managing groups includes the following tasks:

- Adding a new group
- Deleting a group
- Modifying group properties, including user memberships

Local groups in a workgroup environment are managed through the **Users** option in the WebUI.

In the WebUI, under **Users**, **Local Groups** is the **Local Groups on Server Appliance** dialog box. All workgroup group administration tasks are performed in the Local Groups on Server Appliance dialog box.

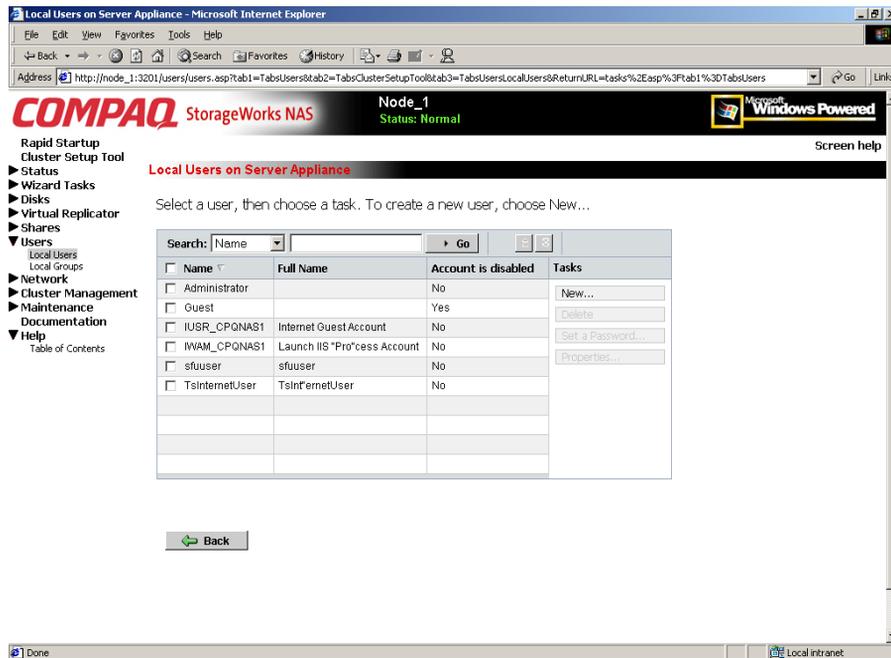


Figure 7-5: Local Groups dialog box

Adding a New Group

To add a group:

1. In the **Local Groups** dialog box, click **New**.

The **Create New Group** dialog box is displayed.

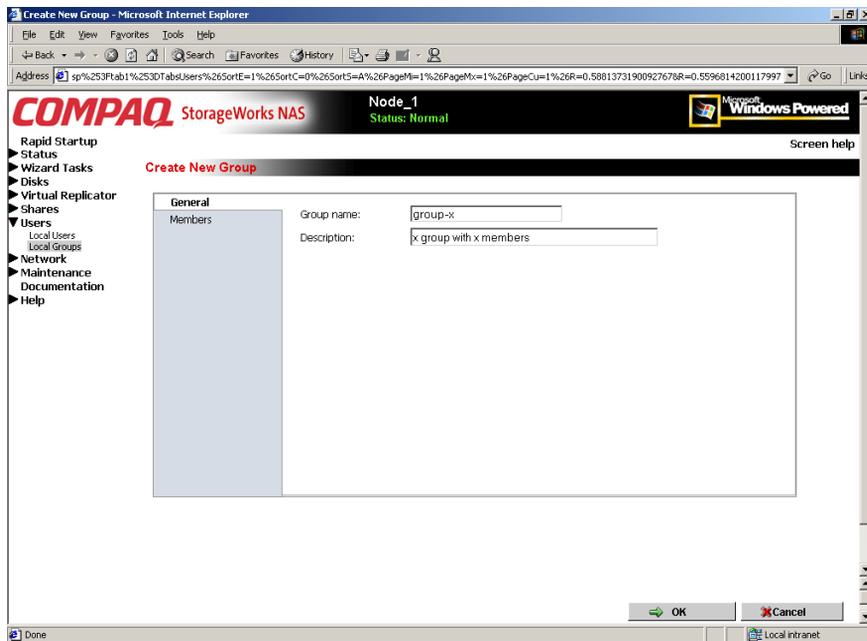


Figure 7-6: Create New Group dialog box, General tab

2. Enter the group name and description.
3. To indicate the user members of this group, click **Members**. See “Modifying Group Properties” for procedural instructions on entering group members.
4. After all group information is entered, click **OK**. The group is added, and the **Local Groups** dialog box is displayed again.

Deleting a Group

To delete a group:

1. From the **Local Groups** dialog box, select the group to delete, and then click **Delete**.
2. The **Delete Group** dialog box is displayed. Verify that this is the intended group and then click **OK**. The **Local Groups** dialog box is displayed again.

Modifying Group Properties

To modify other group properties:

1. From the **Local Groups** dialog box, select the desired group and then click **Properties**. The **Properties** dialog box is displayed.

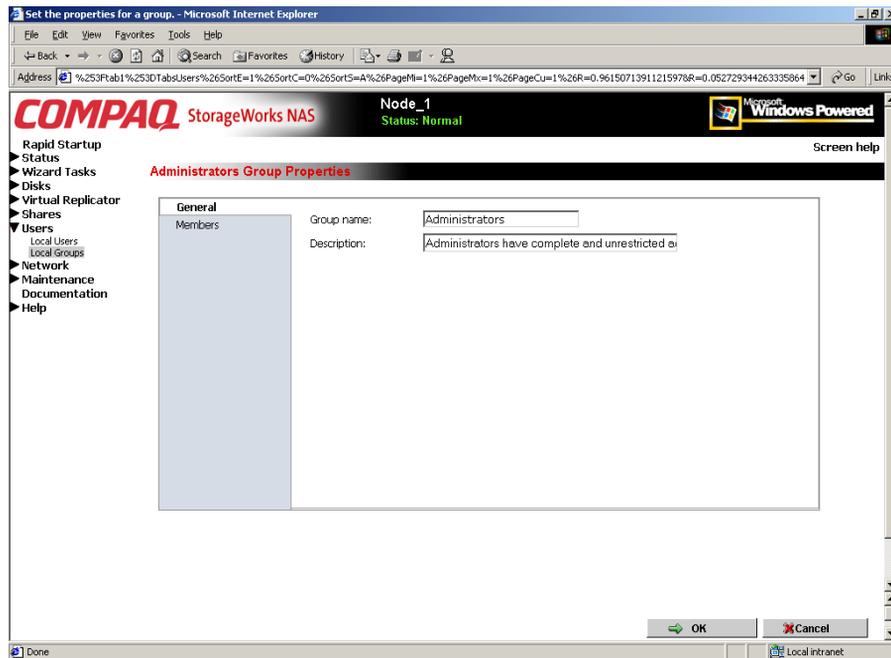


Figure 7-7: Group Properties dialog box, General tab

Within the Properties dialog box are two tabs:

- General tab
- Members tab

Each of these tabs is discussed in the following paragraphs.

2. Enter the desired changes in each of the tabs. Then, click **OK**. The Local Groups dialog box is displayed again.

General Tab

Within the General tab, basic group information can be changed, including:

- Group name
- Description

Members Tab

To indicate or change the members of a group, click the **Members** tab. Within this dialog box, users are added and removed from a group.

Two boxes are displayed: **Members** and **Add User or group**. Current members of that group are listed in the Members box. All users are listed in the Add user or group box.

- To add an existing local user to a group, select the desired user from the **Add Users** box and then click the **Add** button.
- To remove an existing local user from a group, select the desired user from the **Members** box, and then click the **Remove** button.
- To add user or group from a domain to this group, the scroll bar at the right of the screen may need to be used to scroll up the screen display. Enter the user or group name to include in the indicated format (domain/user).

Figure 7-8 is an example of the Members tab.

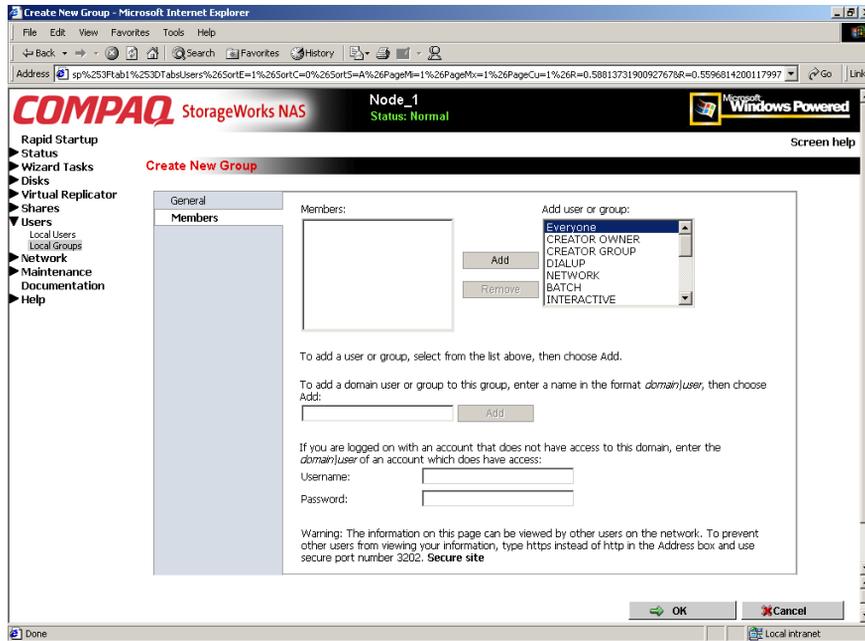


Figure 7-8: Group Properties dialog box, Members tab

Folder and Share Management

The *StorageWorks* NAS B2000 supports several file-sharing protocols, including CIFS, NFS, FTP, HTTP, NCP, and AFP (AppleTalk). This chapter discusses overview information as well as procedural instructions for the setup and management of the file shares for the supported protocols. In addition, discussions on security at the file level and at the share level are included in this chapter.

Abbreviated information on creating NFS file shares is included in this chapter; for detailed information on setting up and managing NFS file shares, see the “UNIX File System Management” chapter.

NCP shares must be set up and managed through the Microsoft Management Console (MMC) user interface. For information on managing NCP file shares, see the “NCP File System Management” chapter.

More information about Windows file system security is available on the Microsoft website:

www.microsoft.com/

The following topics are discussed in this chapter:

- Folder Management
 - Navigating to a Specific Volume or Folder
 - Creating a New Folder
 - Deleting a Folder
 - Modifying Folder Properties
 - Creating a New Share for a Volume or Folder
 - Managing Shares for a Volume or Folder
 - Managing File-Level Permissions, Open Sessions, and Open Files
- Share Management
 - Defining Access Control Lists
 - Integrating Local File System Security into Windows Domain Environments
 - Comparing Administrative and Standard Shares
 - Planning for Compatibility between File-Sharing Protocols
 - Managing Shares Using the Shares Management Wizard
 - Managing Shares (Details)
- Protocol Settings

All procedures in this chapter are documented using the WebUI. In addition to this guide, you may use the WebUI online help.

Folder Management

Volumes and folders on any system are used to organize data. Regardless of system size, systematic structuring and naming conventions of volumes and folders eases the administrative burden. Moving from volumes to folders to shares increases the level of granularity of the types of data stored in the unit and the level of security access allowed.

Although a variety of methods can be used to create and manage file folders on the NAS B2000, this document discusses using the NAS Web-based user interface (WebUI.)

Managing system volumes and file folders includes the following tasks:

- Navigating to a specific volume or folder
- Creating a new folder
- Deleting a folder
- Modifying folder properties
- Creating a new share for a volume or folder
- Managing shares for a volume or folder
- Managing file-level permissions

Navigating to a Specific Volume or Folder

When you work with volumes and folders, the first task is to gain access to the desired volume or folder.

The steps are the same, whether navigating to a volume or a folder:

1. To navigate to a specific volume or folder, from the WebUI, select **Shares** and then **Folders**. Initially, the **Volumes** dialog box is displayed.

This initial dialog box displays all system volumes and snapshots.

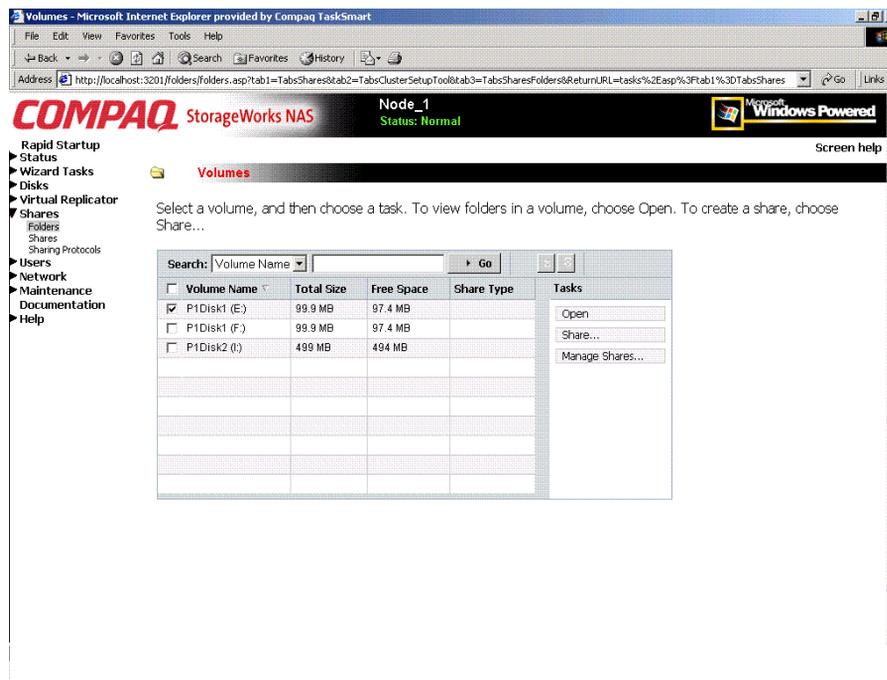


Figure 8-1: Volumes dialog box

2. From this dialog box, navigate to a specific folder by selecting the appropriate volume and then clicking **Open**. The **Folders** dialog box is displayed, with a list of all of the folders within that volume.

- To navigate to a subfolder, select the folder in which the subfolder resides, and then click **Open**. Repeat this searching and opening process until the desired folder is opened. See Figure 8-2 for an example of **Folders** dialog box.

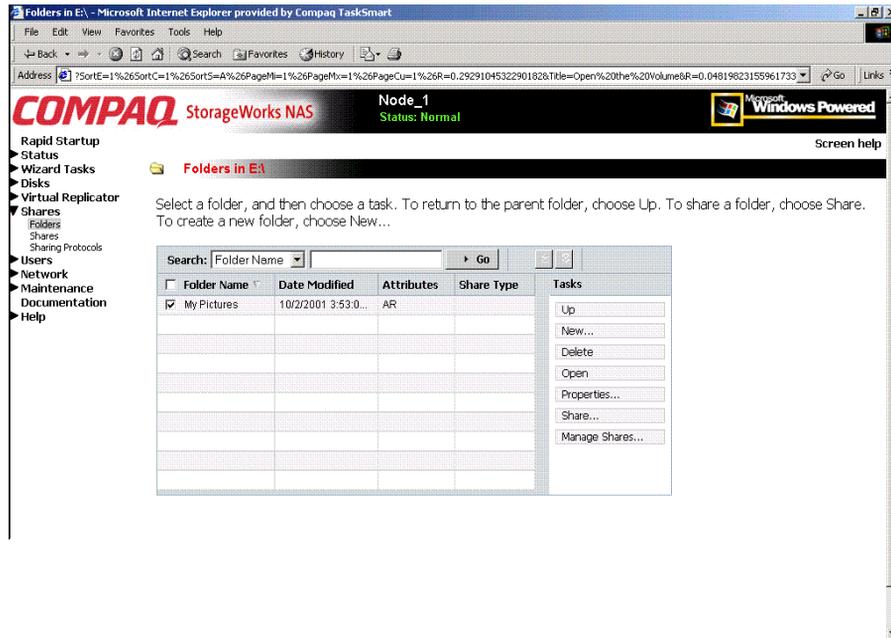


Figure 8-2: Folders dialog box

After accessing the desired folder, the following actions can be performed:

- Creating a new folder
- Deleting a folder
- Modifying folder properties
- Creating a new share for the volume or folder
- Managing shares for the volume or folder

Creating a New Folder

To create a new folder:

1. From the **Shares** directory, navigate to the **Folders** menu and then select **New**. The **Create New Folder** dialog box is displayed.

Two tabs are displayed: **General** and **Compress**. Use these two tabs to enter the parameters for the new folder.

2. In the **General** tab, enter a name for the folder and specify the folder attributes.

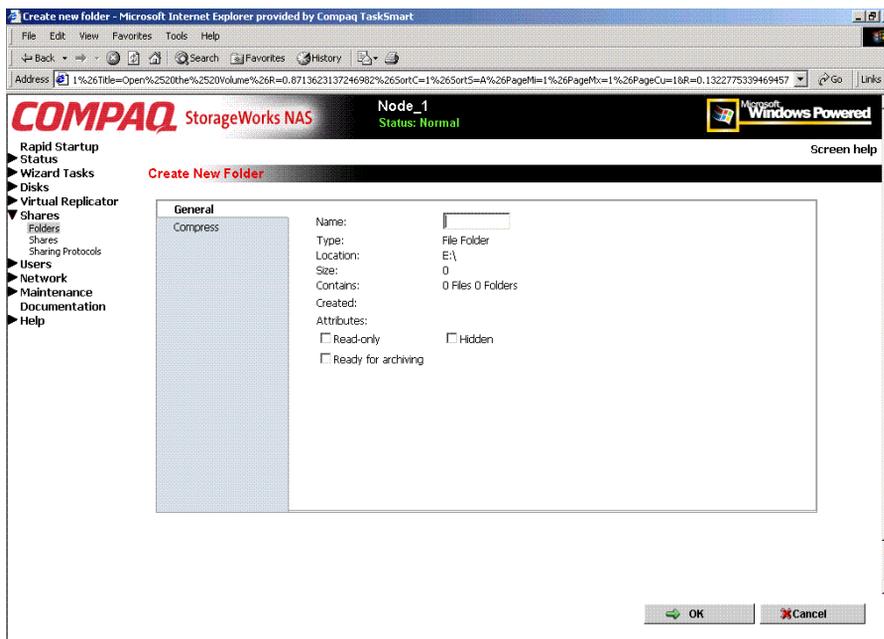


Figure 8-3: Create a New Folder dialog box, General tab

3. In the **Compress** tab, indicate whether and how this folder and its contents are to be compressed.
4. After all information for the new folder is entered, click **OK**.

Deleting a Folder

To delete a folder:

1. From the **Shares** directory, navigate to the folder to delete. Select the folder and then click **Delete**. The **Delete Folder** dialog box is displayed.

Summary information about the deletion is displayed.

IMPORTANT: View the summary information to confirm that this is the intended share.

2. Verify that the displayed folder is the folder to delete and then click **OK**.
The folder and all of its subfolders are deleted and the main dialog box is displayed again.

Modifying Folder Properties

To modify folder properties:

1. From the **Shares**, directory, navigate to the folder whose properties need to be edited. Then click **Properties**. The **Properties** dialog box is displayed.

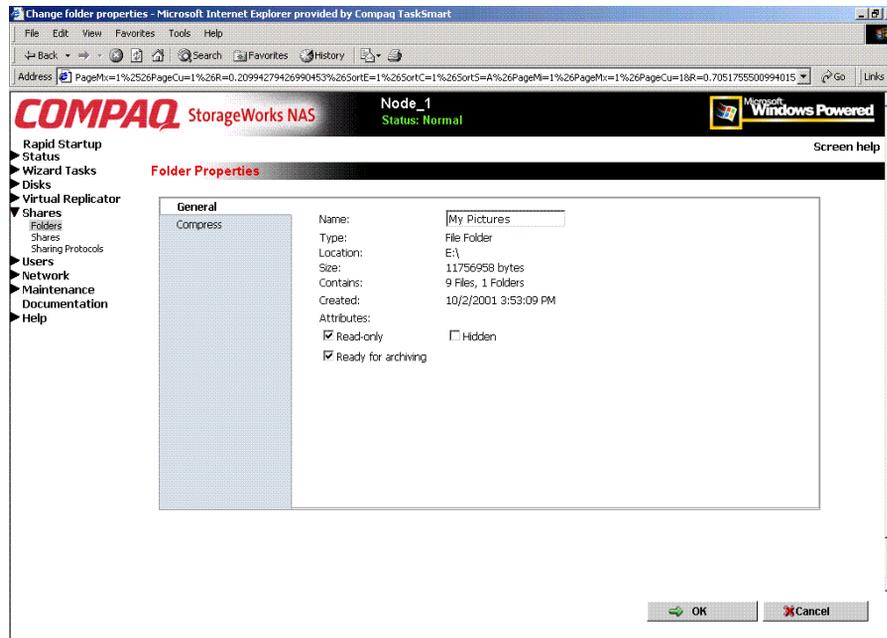


Figure 8-4: Folder Properties dialog box, General tab

2. In the **General** tab, enter the new information for the folder, which may include:
 - Folder Name
 - Folder Attributes
3. In the **Compress** tab, indicate whether and how this folder and its contents are to be compressed.
4. After all changes have been completed, click **OK**. The **Folders** dialog box is displayed again.

Creating a New Share for a Volume or Folder

Within the WebUI, there are two access points to the same screens used to create file shares:

- A share can be created for a folder while working with that folder *in the **Folders** screens*.
- A share can be created and, if necessary, new folders can be created, while working with file shares *in the **Shares** screens*.

This section discusses creating shares from the **Folders** menu, and is an overview of the procedures. Complete details on the process of creating shares are included in the discussion that documents creating shares through the **Shares** menu. See the “Managing Shares” section of this chapter for these details.

To create a new share for a specific volume or folder, *while in the **Folders** menu*:

1. Navigate to the desired volume or folder and click **Share**. The **Create New Share** dialog box is displayed.

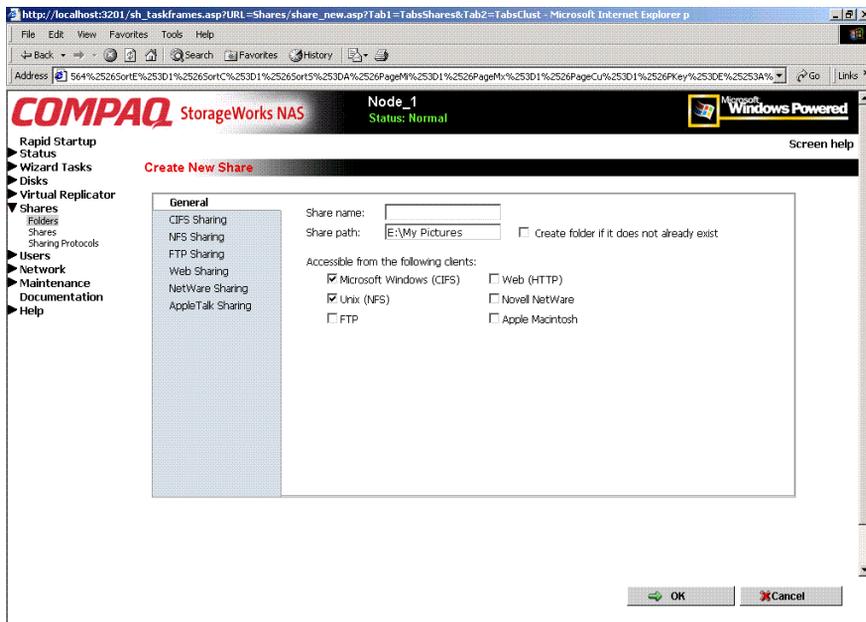


Figure 8-5: Create New Share dialog box, General tab

2. Enter the information for the share, including the name of the share, the allowed protocols, and corresponding permissions.

NOTE: The Share Path is the path of the previously selected volume or folder. This field is automatically completed by the system.

3. Select the appropriate tab to enter protocol-specific information.
See “Managing Shares” in the next section for detailed information about these entries.
4. After entering all share information, click **OK**.

Managing Shares for a Volume or Folder

Within the WebUI, there are two access points to the same screens used to manage file shares:

- While working with a folder *in the **Folders** dialog boxes*, the administrator can create, delete, and modify shares for that folder.
- While working with file shares *in the **Shares** dialog boxes*, the administrator can create, delete, and modify shares (and if necessary, create new folders).

NOTE: This section discusses managing shares from the **Folders** menu, and is an overview of the procedures. Complete details on the process of managing shares are included in the discussion that documents creating shares through the **Shares** menu. See the “Managing Shares” section later in this chapter for these details.

To create, delete, and manage shares for a particular volume or folder *while in the **Folders** menu*:

1. From the **Folders** directory, navigate to the target volume or folder and click **Manage Shares**. The **Shared Folders** dialog box is displayed.

All associated shares for that folder or volume are listed.

2. To create a new share, click **New**. The **Create a New Share** dialog box is displayed.

Because the screens are the same whether shares are managed through the Folders menu or the Shares menu, the procedures are only documented once. See “Creating a New Share” in the “Share Management” section for detailed procedural instructions on creating new file shares.

3. To delete a share, select the share to delete and click **Delete**. The **Delete Share** dialog box is displayed.

Because the screens are the same whether shares are managed through the Folders menu or the Shares menu, the procedures are only documented once. See “Creating a New Share” in the “Share Management” section for detailed procedural instructions on deleting file shares

4. To modify share properties, select the share to modify, and click **Properties**. The **Share Properties** dialog box is displayed.

Because the screens are the same whether shares are managed through the Folders menu or the Shares menu, the procedures are only documented once. See “Creating a New Share” in the “Share Management” section for detailed procedural instructions on modifying shares.

Managing File-Level Permissions

The WebUI of the NAS B2000 provides security at the share level and is discussed later in this chapter. Security at the file level is managed using Windows Explorer available from the desktop of the NAS B2000. To access the NAS B2000 desktop from the WebUI, go to the Maintenance menu, and select Terminal Services.

File-level security includes settings for permissions, ownership, and auditing for individual files.

To enter file permissions:

1. Using Windows Explorer, navigate to the folder or file that needs to be changed and then right click the folder.
2. Select **Properties** and then select the **Security** tab. Figure 8-6 illustrates the properties available on the Security tab.

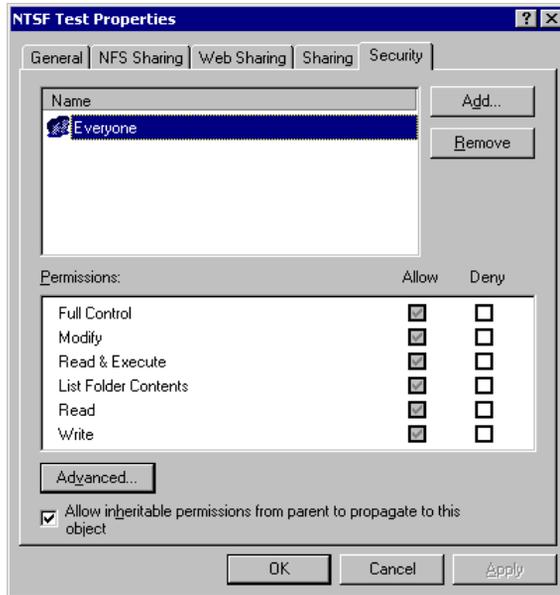


Figure 8-6: Security Properties dialog box for folder name NTSF Test

Several options are available in the Security tab dialog box:

- To add users and groups to the permissions list, click **Add**. Then follow the dialog box instructions.
- To remove users and groups from the permissions list, highlight the desired user or group and then click **Remove**.
- If the “Allow inheritable permissions from parent to propagate to this object” box at the bottom of the screen is checked, the file or directory inherits permissions from the parent directory. In this case, existing user and group permissions cannot be changed; however, additional users or groups can be added.
- The center section of the Security tab provides a listing of permission levels. When new users or groups are added to the permissions list, select the appropriate boxes to configure the common file access levels.

NOTE: Selections can be made when the “Allow inheritable permissions from parent to propagate to this object” box is disabled.

- To modify ownership of files or to modify individual file access level permissions, click the **Advanced** button.

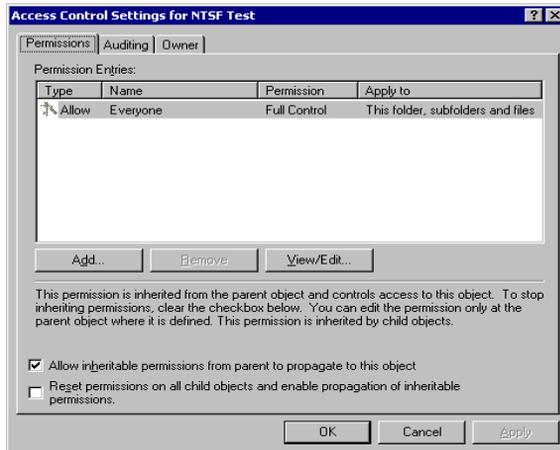


Figure 8-7: Access Control Settings dialog box for folder name NTSF Test, Permissions tab

To modify specific permissions assigned to a particular user or group for a selected file or folder in the **Advanced** screen:

1. Select the desired user or group.
3. Click **View/Edit**.
4. Check all the permissions that you want to enable, and clear the permissions that you want to disable. Enable or disable permissions by selecting the Allow box to enable permission or the Deny box to disable permission. If neither box is selected, permission is automatically disabled. Figure 8-8 illustrates the View/Edit screen and some of the permissions.

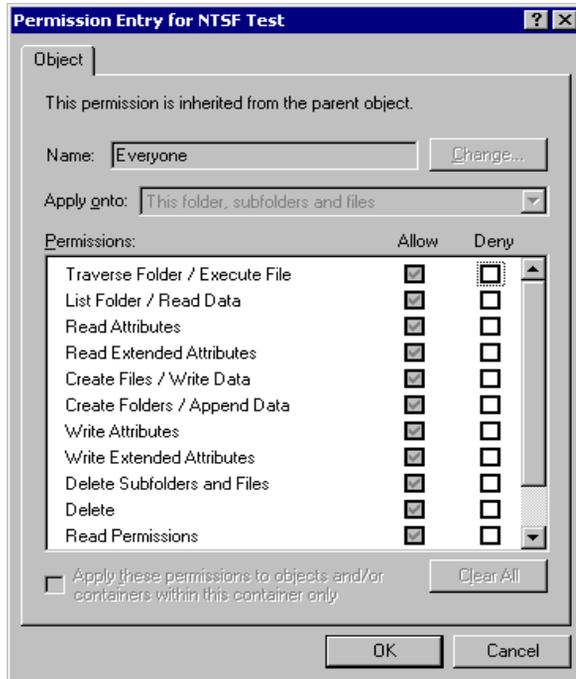


Figure 8-8: User or Group Permission Entry dialog box for folder name NTFS Test

Other functionality available in the Advanced Access Control Permissions tab is illustrated in Figure 8-8 and includes:

- Add a new user or group. Click **Add**, and then follow the dialog box instructions.
- Remove a user or group. Click **Remove**.
- Inherit permissions from the parent folder. Enable the “Allow inheritable permissions from parent to propagate to this object” box.
- Reset permissions. If the object being configured is a folder, check the “Reset permissions on all child objects and enable propagation of inheritable permissions” box, which allows all child folders and files to inherit the current folder permissions by default.

Another area of the Advanced Access Control is the Auditing tab. Auditing allows you to set rules for the auditing of access, or attempted access, to files or folders. Users or groups can be added, deleted, viewed, or modified through the advanced Access Control Settings Auditing tab. The Auditing tab dialog box is illustrated in Figure 8-9.

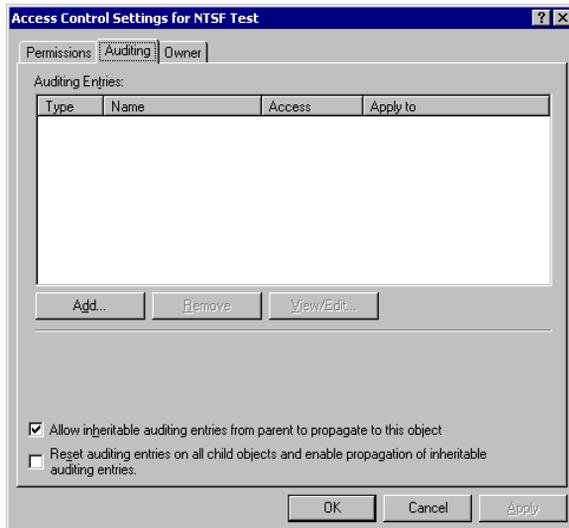


Figure 8-9: Access Control Settings, Auditing tab dialog box for folder name NTSF Test

Figure 8-10 illustrates the screen that is displayed when a user or group to be audited is added.

5. Select the appropriate domain or machine name from the “Look in:” drop-down list box at the top of the screen.

NOTE: A list of users and groups from the desired domain can be viewed if the current user has permission to view the information on the domain.

6. Select the user or group.

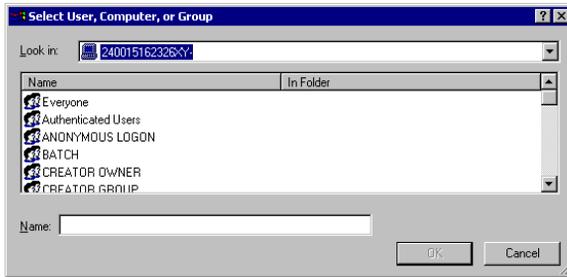


Figure 8-10: Select User, Computer, or Group dialog box

7. Click **OK**. Figure 8-11 illustrates the Auditing Entry screen that is displayed.

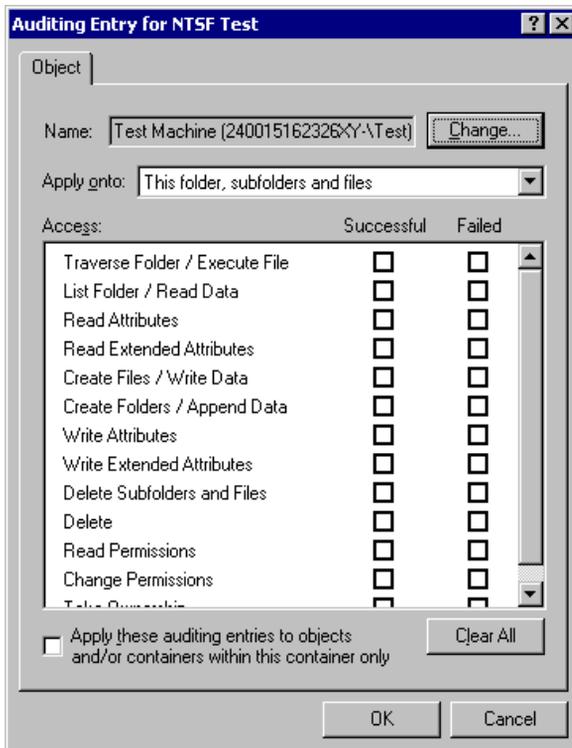


Figure 8-11: Auditing Entry dialog box for folder name NTSF Test

8. Select the desired Successful and Failed audits for the user or group as shown in Figure 8-11.
9. Click **OK**.

NOTE: Auditing must be enabled to configure this information. Use the local Computer Policy Editor to configure the audit policy on the NAS B2000.

The final tab in the advanced Advanced Access Control Settings security configuration is the Owner tab. This tab allows for taking ownership of files. Typically, administrators use this area to take ownership of files when the file ACL is incomplete or corrupt. By taking ownership, you gain access to the files and then manually apply the appropriate security configurations. Figure 8-12 illustrates the Owner tab.

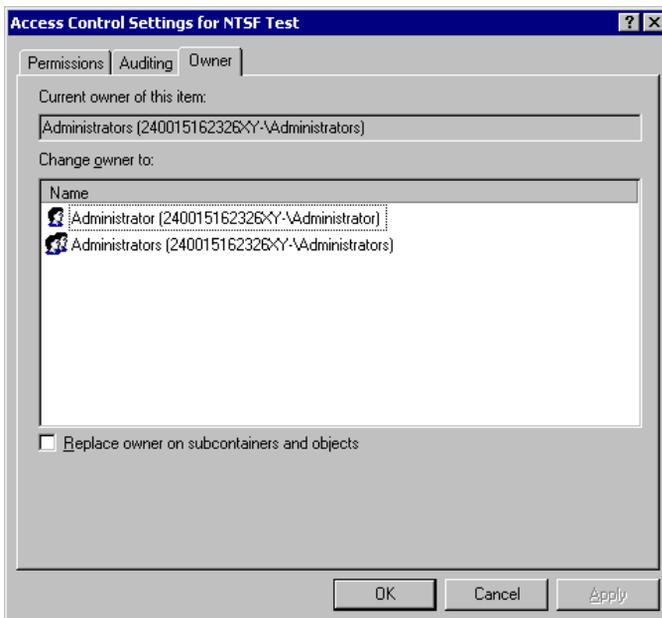


Figure 8-12: Access Control Settings, Owner tab dialog box for folder name NTSF Test

The current owner of the file or folder is listed at the top of the screen. To take ownership:

1. Select the appropriate user or group from the “Change owner to” list.
2. If it is also necessary to take ownership of subfolders and files, enable the “Replace owner on subcontainers and objects” box.
3. Click **OK** to execute the commands.

Share Management

There are several ways to set up and manage shares. The WebUI provides screens for setting up and managing shares. Additional methods include using a command line interface, Windows Explorer, or Microsoft Management Console (MMC). This guide demonstrates using the WebUI to set up and manage shares.

As previously mentioned, the file-sharing security model of the NAS device is based on the NTFS file-level security model. Share security seamlessly integrates with file security. In addition to discussing share management, this section discusses share security. See “Managing File-Level Security” earlier in this chapter for information on file security.

Shares management topics include:

- Defining Access Control Lists
- Integrating Local File System Security into Windows Domain Environments
- Comparing Administrative and Standard Shares
- Planning for Compatibility between CIFS and NFS
- Managing Shares Using the Shares Management Wizard
- Managing Shares (Details)

Defining Access Control Lists

The Access Control List (ACL) contains the information that dictates which users and groups have access to a share, as well as the type of access that is permitted. Each share on an NTFS file system has one ACL with multiple associated user permissions. For example, an ACL can define that User1 has read and write access to a share, User2 has read-only access, and User3 has no access to the share. The ACL also includes group access information that applies to every user in a configured group. ACLs are also referred to as permissions.

Integrating Local File System Security into Windows Domain Environments

ACLs include properties specific to users and groups from a particular workgroup server or domain environment. In a multidomain environment, user and group permissions from several domains can apply to files stored on the same device. Users and groups local to the NAS B2000 can be given access permissions to shares managed by the device. The domain name of the NAS B2000 supplies the context in which the user or group is understood. Permission configuration depends on the network and domain infrastructure where the server resides.

File-sharing protocols (except NFS) supply a user and group context for all connections over the network. (NFS supplies a machine-based context.) When new files are created by those users or machines, the appropriate ACLs are applied.

Configuration tools provide the ability to share permissions out to clients. These shared permissions are propagated into a file system ACL and when new files are created over the network, the user creating the file becomes the file owner. In cases where a specific subdirectory of a share has different permissions from the share itself, the NTFS permissions on the subdirectory apply instead. This method results in a hierarchical security model where the network protocol permissions and the file permissions work together to provide appropriate security for shares on the device.

NOTE: Share permissions and file-level permissions are implemented separately. It is possible for files on a file system to have different permissions from those applied to a share. When this situation occurs, the file-level permissions override the share permissions.

Comparing Administrative (Hidden) and Standard Shares

CIFS supports both administrative shares and standard shares. Administrative shares are shares with a last character of \$. Administrative shares are not included in the list of shares when a client browses for available shares on a CIFS server. Standard shares are shares that do not end in a \$ character. Standard shares are listed whenever a CIFS client browses for available shares on a CIFS server.

The NAS B2000 supports both administrative and standard CIFS shares. To create an administrative share, end the share name with the \$ character when setting up the share. Do not type a \$ character at the end of the share name when creating a standard share.

Planning for Compatibility between File-Sharing Protocols

When planning for cross-platform share management on the NAS B2000, it is important to understand the different protocols and their associated constraints. Each additional protocol that is supported adds another level of constraints and complexity.

NFS Compatibility Issues

Of the file sharing protocols that are supported on the NAS B2000, NFS introduces the most constraints. When planning to manage CIFS and NFS shares, consider two specific requirements.

NOTE: Further information, including details about the NFS Service and the User Mapping service, is available in the “UNIX File System Management” chapter.

- **NFS service does not support spaces in the names for NFS file shares.**

NFS translates any spaces in an export into an underscore character. Additional translations can be set up for files. See the Compaq “OEM Supplemental Help” chapter of the SFU help, found on the NAS B2000. This feature is designed to ensure the greatest level of compatibility with NFS clients, because some do not work with NFS exports that contain a space in the export name.

If you plan to use the same name when sharing a folder through CIFS, and then exporting it through NFS, do not put spaces in the CIFS share name.

- **NFS service does not support exporting a child folder when its parent folder has already been exported.**

An NFS client can access a child folder by selecting the parent folder and then navigating to the child folder. If strict cross-platform compatibility is an administration goal, CIFS must be managed in the same way. Do not share a folder through CIFS if the parent folder is already shared.

Managing Shares Using the Shares Management Wizard

The **Share Management** wizard is one of three major wizards included in the WebUI. The wizard contains most of the functionality of managing system shares as the **Shares** menu option of the WebUI. Although all share administration tasks are not included in the wizard, it provides the easiest to use interface for adding, modifying, and deleting shares. For detailed information on managing file shares, see “Managing Shares (Details).” For detailed information on managing NFS (UNIX) file shares, see the “UNIX File System Management” chapter.

To use the wizard:

1. From the WebUI, select **Wizard Tasks**, and then **Share Management**. The Welcome screen of the wizard is displayed. Click **Next** to continue.

The primary menu of the Share Management Wizard is displayed. Figure 8-13 is an example of the Share Management Wizard screen.

2. From this screen, indicate the action to perform on system shares. Options include:
 - Create a new share
 - Delete an existing share
 - Modify share properties

Each of these options is discussed in the following paragraphs.

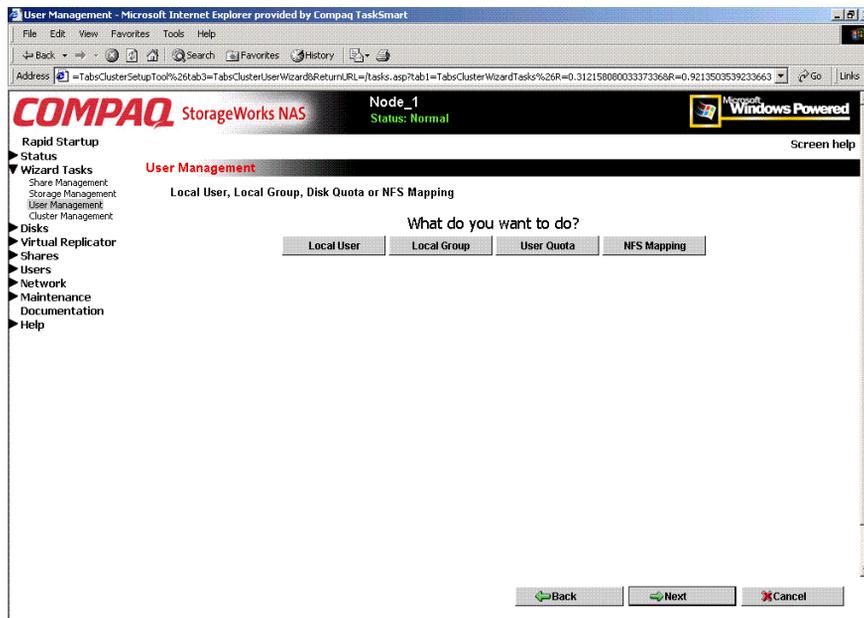


Figure 8-13: Share Management wizard

Create a New Share

After selecting the create option, the system prompts with the following questions:

1. Do you want to create a new pool for the new share?

When creating a new pool for this share, subscreens will prompt for the Storage Unit to use, a Name for the pool, and the Segment Size to use.

2. Do you want to create a new virtual disk for the new share?

When creating a new virtual disk for this share, subscreens will prompt for the Pool to use, a Name for the virtual disk, disk size, disk drive letter to assign, and allocation size. In addition, the option is available for system to format the new virtual disk at this time.

3. Select the folder to use. To create a new subfolder for this share, enter the folder name to use in the indicated space at the bottom of the dialog box.
4. Enter basic share settings, including share Name, Path, and access protocols to use.
5. In the final screen of the **Create a new share** option of the wizard, summary information is displayed. Verify the accuracy of the data.
6. To enter protocol-specific information about the share, click the **More Property** button.
7. After all settings are entered, click **Next**.

Delete an Existing Share

1. After selecting the delete option, the **List Shares** screen of the wizard is displayed.
2. Select the share to delete from the displayed list, and then click **Next**.
3. A verification screen is displayed. Confirm that this is the correct share, and then click **Yes**. The share is deleted, and the primary screen of the Share Management Wizard is displayed again.

Modify Existing Share Properties

To modify the properties of an existing share:

1. Select **Modify existing share properties**. In the **List Shares** screen, select the share to modify, and then click **Next**.
2. In the Show Share Property dialog box, enter the new information for the share.
3. Tabs are available in the Show Share Properties dialog box:
 - **General**—Use to indicate the protocols that will access this share.
 - **Protocol-specific**—Use to enter information that is unique to each sharing protocol.
4. After all changes are entered, click **Next**.

Managing Shares (Details)

In addition to the Share Management Wizard, shares can be managed through the **Shares** menu option of the WebUI. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties

Each of these tasks is discussed in this section.

Creating a New Share

To create a new share:

1. From WebUI main menu, select the **Shares** directory and then select the **Shares** option. The Shares dialog box is displayed. From the **Shares** dialog box, click **New**. The **General** tab of the **Create a New Share** dialog box is displayed.

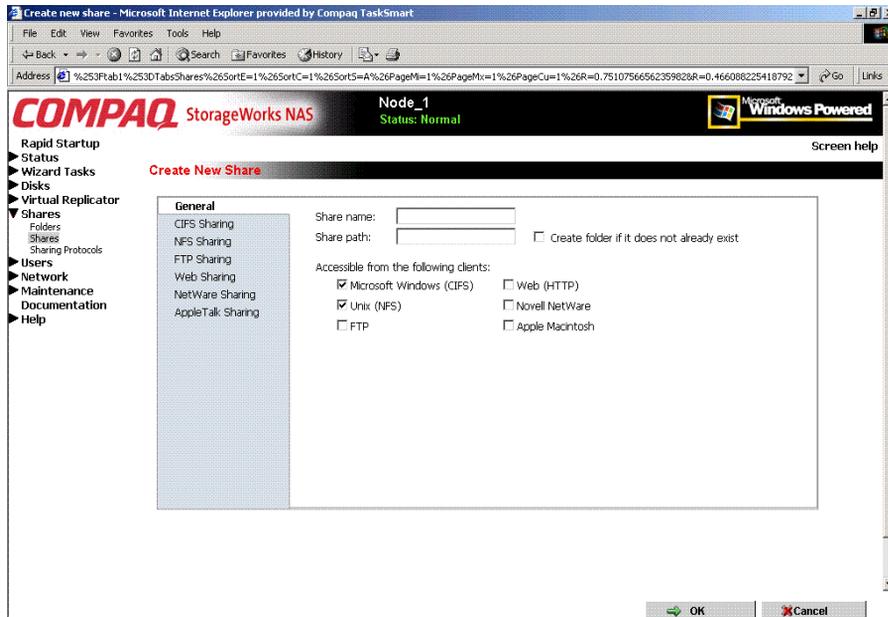


Figure 8-14: Create a New Share dialog box, General tab

2. Enter the following information:

- Share name
- Share path
- Client protocol types

To create a folder for the new share, check the indicated box and the system will create the folder at the same time it creates the share.

Protocol-specific tabs are available to enter sharing and permissions information for each sharing type. See “Modifying Share Properties” for detailed information on these tabs.

3. After entering all share information, click **OK**.

Deleting a Share

IMPORTANT: Before deleting a share, warn all users to exit that share and confirm that no one is using the share.

To delete a share:

1. From the **Shares** menu, click **Delete**.
2. Verify that this is the correct share, and then click **OK**.

NOTE: This option deletes only the share. The resource is not deleted.

Modifying Share Properties

To change share settings:

1. From the **Shares** menu, select the share to modify and then click **Properties**. The **General** tab of the **Share Properties** dialog box is displayed.

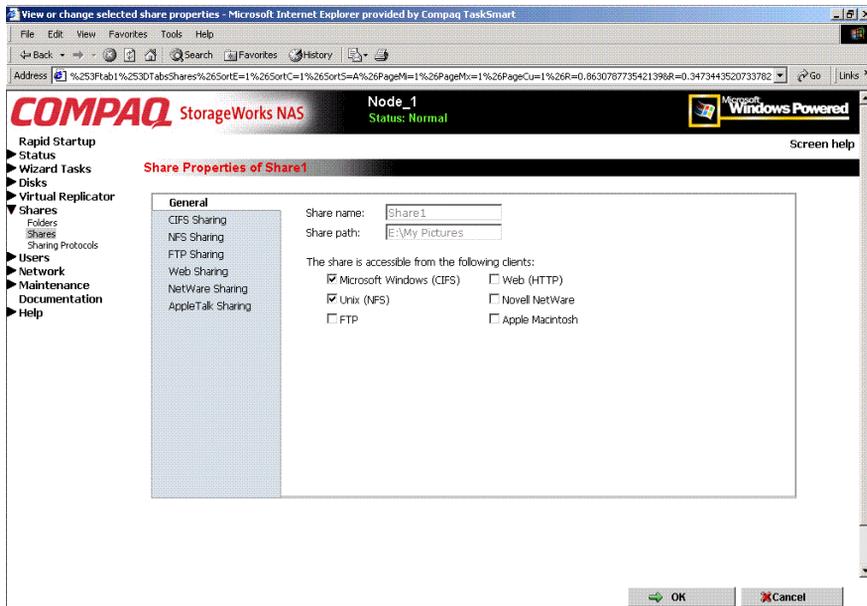


Figure 8-15: Share Properties dialog box, General tab

The name and path of the selected share is displayed.

2. To enter or change client protocol information, check the appropriate boxes and then click the corresponding tabs.
 - CIFS Sharing
 - NFS Sharing
 - FTP Sharing
 - Web Sharing (HTTP)
 - NetWare Sharing (NCP)
 - AFP (AppleTalk) Sharing

Each of these tabs is discussed in the following paragraphs.

3. After all share information has been entered, click **OK**. The **Share** menu is displayed again.

CIFS Sharing

From the **CIFS Sharing** tab of the **Share Properties** dialog box:

1. Enter a descriptive **Comment**, and the **User limit** (optional).
See Figure 8-16 for an example of the CIFS Sharing tab screen display.
2. If file caching on the client machines is allowed, click **Enable file caching on client computers accessing this share**.

Select one of the following caching policies:

- **Manual Caching for Documents**—The default setting. Recommended for folders containing user documents. Users must manually specify any files that they want available when working offline. To ensure proper file sharing, the server version of the file is always open.
- **Automatic Caching for Documents**—Also recommended for folders containing user documents. In contrast to the default setting of Manual Caching, with this option, open files are automatically downloaded and made available when working offline. Older copies are automatically deleted to make room for newer, more recently accessed files. To ensure proper file sharing, the server version of the file is always open.
- **Automatic Caching for Programs**—Recommended for folders with read-only data or run-from-the-network applications. File sharing is not ensured. Open files are automatically downloaded and made available when working offline. Older copies are automatically deleted to make room for newer, more recently accessed files.

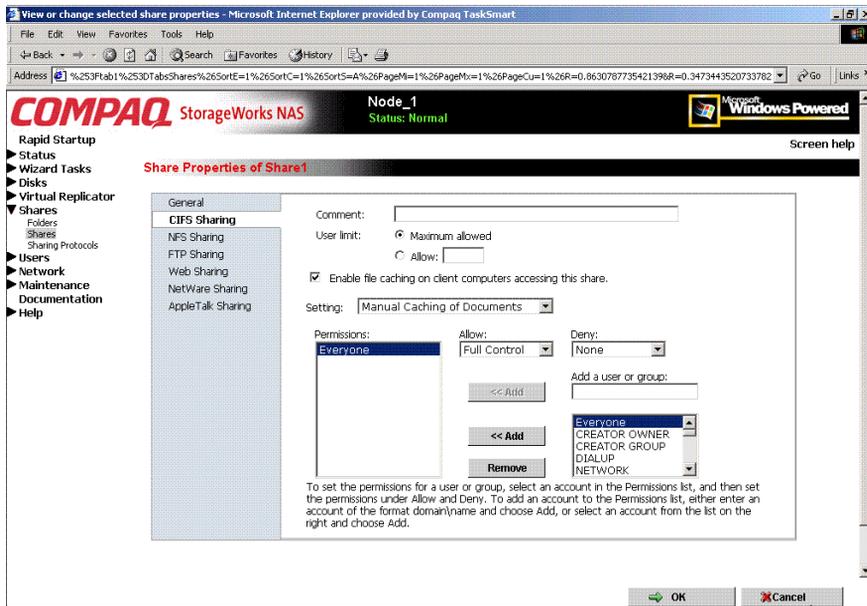


Figure 8-16: Share Properties dialog box, CIFS Sharing tab

3. Enter Permissions information:

- The **Permissions** box lists the currently approved users for this share.
- *To add a new user or group*, either select a user or group from the box at the bottom right of the screen or manually enter the user or group name in the **Add a user or group** box and then click **Add**. That user or group is added to the Permissions box.
- *To remove access to a currently approved user or group*, select the user or group from the Permissions box and then click **Remove**.
- *To indicate the type of access allowed for each user*, select the user and then expand the **Allow** and **Deny** drop-down boxes. Select the appropriate option.

4. After all CIFS Sharing information is entered, either click the next Sharing tab or click **OK**.

NFS Sharing

From the **NFS Sharing** tab of the **Create a New Share** dialog box:

1. Indicate the machines that will have access to this share.

Select the machine to include in the **Select a client or client group** box or manually enter the NFS client computer name or IP address. Then click **Add**.

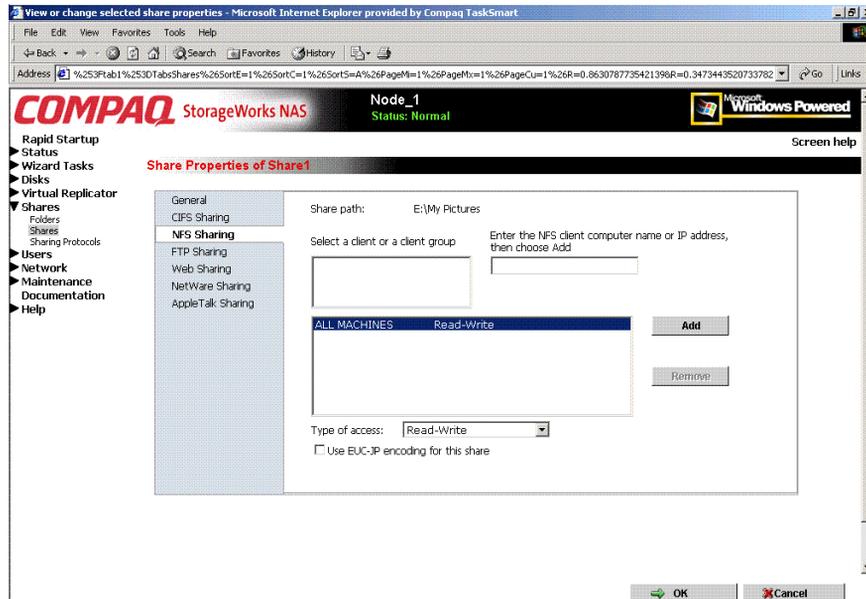


Figure 8-17: Share Properties dialog box, NFS Sharing tab

2. Indicate the permissions.

Select the machine from the main user display box, and then select the appropriate access methods from the **Type of access** drop-down box at the bottom of the screen.

3. After all NFS sharing information is entered, either click the next Sharing tab or click **OK**.

FTP Sharing

From the **FTP Sharing** tab of the **Create a New Share** dialog box:

1. Select the read and write access permissions that are allowed, and indicate whether visits should be written to the FTP log.
2. Then, either click the next Sharing tab or click **OK**.

Web Sharing (HTTP)

From the **Web Sharing** tab of the **Create New Share** dialog box:

1. Select the read and write access permissions that are allowed, and indicate whether visits should be written to the HTTP log.
2. Then, either click the next Sharing tab or click **OK**.

NetWare Sharing (NCP)

NOTE: NCP shares can be set up only after Microsoft Services for NetWare (SFN) has been installed on the NAS B2000. Procedures for installing SFN are included in the “NetWare File System Management” chapter.

From the **NetWare Sharing** tab, as illustrated in Figure 8-15, of the **Create a New Share** dialog box:

1. Enter a user limit.
2. Enter Permissions information.
 - The **Permissions** box lists the currently approved users for this share.
 - To add a new user or group, either select a user or group from the box at the bottom right of the screen or manually enter the user or group name in the **Add a user or group box**. Then click **Add**. That user or group is added to the Permissions box.
 - To remove access to a currently approved user or group, select the user or group from the Permissions box, and then click **Remove**.
 - To indicate the allowed access for each user, select the user and then expand the **Allow** and **Deny** drop-down boxes. Then, select the appropriate option.

3. After all NetWare Sharing information is entered, either click the next Sharing tab or click **OK**.

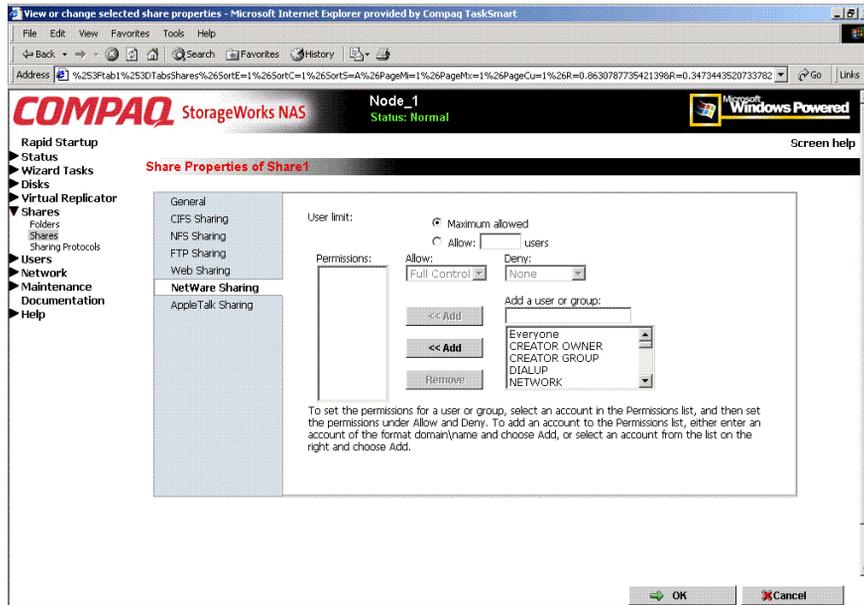


Figure 8-18: Share Properties dialog box, NetWare Sharing tab

AFP (AppleTalk) Sharing

From the **AppleTalk Sharing** tab of the **Create a New Share** dialog box:

1. Enter a user limit.
2. Enter password information.
3. Indicate whether the share has read-only permission or read-write permission.
4. After all AFP (AppleTalk) Sharing information is entered, either click the next Sharing tab or click **OK**.

Protocol Parameter Settings

As previously mentioned, the NAS B2000 supports the following protocols:

- CIFS
- NFS
- FTP
- HTTP
- NCP (NetWare)
- AFP (AppleTalk)

This section discusses the parameter settings for each protocol type.

To access and enter protocol parameter settings:

1. From the **Shares** menu, select **Sharing Protocols**. The File Sharing Protocols dialog box is displayed.

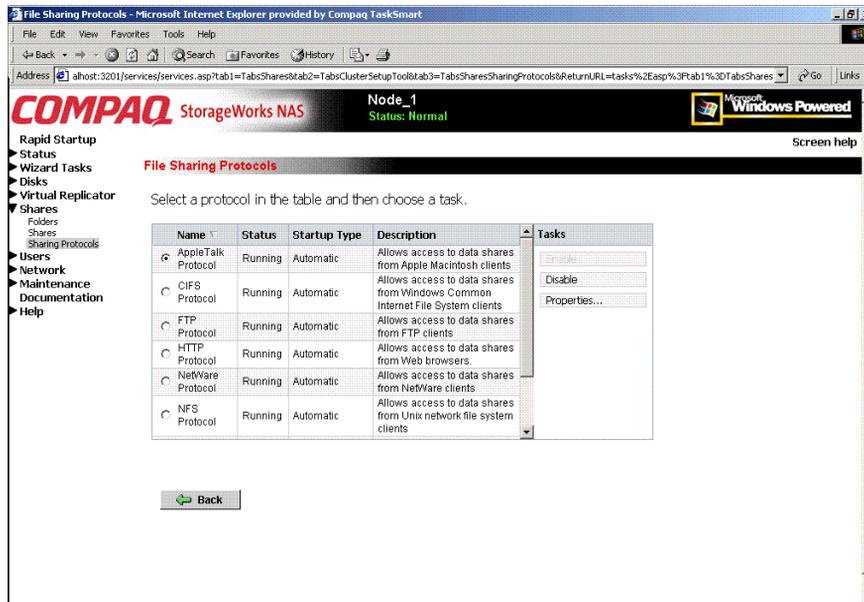


Figure 8-19: Sharing Protocols dialog box

2. Protocols and their statuses are listed. The following options are available:

- Enabling a protocol
- Disabling a protocol
- Modifying Protocol Settings

Because enabling and disabling a protocol are self-explanatory, only modifying protocol-specific settings is described in this section.

CIFS Protocol Settings

There are no user-configurable settings for CIFS.

NFS Protocol Settings

NFS is the networking protocol for exporting UNIX file systems across a network. UNIX and NFS are discussed in the “UNIX File System Management” chapter.

Some of the NFS protocol settings include:

- Async/Sync Settings
- Locks
- Client Groups
- User and Group Maps

FTP Protocol Settings

Three tabs are presented in the FTP Protocol Properties dialog box: **Logging**, **Anonymous Access**, and **Messages**.

Within these tabs:

- **Logging**—Enable logging
- **Anonymous Access**—Enable anonymous access
- **Messages**—Enter a welcome and an exit message

HTTP Protocol Settings

The following parameters can be set for Web protocols:

- Indicate which IP addresses can be used to access data shares
- Indicate which port can be used to access data shares

NCP (NetWare) Protocol Settings

There are no user-configurable settings for NCP.

AFP (AppleTalk) Protocol Settings

Several parameters can be set for AFP shares, including:

- Welcome message
- Security settings
- Limits on number of sessions

UNIX File System Management

Microsoft Services for UNIX (SFU) is a comprehensive software package designed to provide complete UNIX environment integration into a Windows NT, Windows 2000, or Active Directory domain file server. SFU manages tasks on both Windows and UNIX platforms. Tasks include creating NFS exports from Windows and administering user name mappings. SFU also includes Telnet Server and Remote Shell for remote administration.

The following SFU components are included in the NAS B2000:

- Server for NFS
- User Name Mapping
- Telnet and Remote Shell Services
- Password Synchronization.

The following topics are described in this chapter:

- Network File System (NFS)
- Server for NFS
 - Authenticating User Access
 - Indicating the Computer to Use for NFS User Mapping Server
 - Logging Events
 - Installing NFS Authentication Software on the Domain Controller
- NFS File Shares
- NFS Protocol Properties Settings
- NFS Client Groups
 - Adding a New Client Group
 - Deleting a Client Group
 - Editing Client Group Information
- NFS User and Group Mappings
 - Types of Mappings
 - User Name Mapping Best Practices
 - Creating and Managing User and Group Mappings
 - Backing up and Restoring Mappings
- NFS File Sharing Tests
- Terminal Services, Telnet Server, and Remote Shell service
 - Using Terminal Services
 - Using Telnet Server
 - Using Remote Shell Service
- Password Synchronization

Network File System

Network File System (NFS) is a networking protocol for exporting UNIX file systems across a network.

There are two versions of NFS, Version 2 and Version 3. Version 3 supports additional file operations that Version 2 did not have.

In addition, NFS has the capacity to operate with two different network protocols, Transport Control Protocol (TCP) and User Datagram Protocol (UDP).

Traditionally, NFS operates with UDP for performance purposes, but it can also operate with TCP.

There are three key design goals of NFS:

- Allow different UNIX machines to transparently export files across a network.

This feature works across different versions of UNIX and across different platforms. For example, a Linux machine can access files on a *Tru64*TM UNIX machine. Accessing these files is transparent to both the administrator and the users. The administrator and user do not notice any difference between accessing local files or files on the remote machine.

- Make the administration as easy as possible.

The remote file system connects to the local machine in the same manner that a local file system does. The administrator is able to add a remote file system in the same manner as adding another hard drive or external storage.

- Focus exclusively on file system operations.

The file system is used only for exporting file systems to remote machines. NFS supports only operations such as read, write, create, delete, and copy.

Server for NFS

Until recently, UNIX used only NFS to export files. UNIX-based platforms and Windows-based platforms were not able to share files. This restriction caused UNIX clients to require UNIX file servers and Windows clients to require Windows file servers. Windows and UNIX were separate environments, including the duplication of hardware, overhead, and effort. UNIX clients can now use Windows-based machines as file servers using Microsoft Services for UNIX (SFU).

SFU enables UNIX clients to use Windows-based machines as file servers. The SFU NFS server supports NFS Version 2 and Version 3, and supports them both on the TCP and UDP network protocols.

SFU is more fully integrated into the operating system than other third-party NFS server packages. The administrative interface for NFS exports is similar to the Common Internet File System (CIFS) sharing interface used by Windows platforms.

Authenticating User Access

NFS export access is granted or denied to clients based on client name or IP address. The server determines whether a specific client machine has access to an NFS export. No user logon to the NFS server takes place when a file system is exported by the NFS server. Permission to read or write to the export is granted to specific client machines. For example, if client machine M1 is granted access to an export but client M2 is not, user jdoe can access the export from M1 but not from M2.

Permissions are granted on a per-export basis; each export has its own permissions, independent of other exports on the system. For example, file system *a* can be exported to allow only the Accounting department access, and file system *m* can be exported allowing only the Management department access. If a user in Management needs access to the Accounting information, the *a* export permissions can be modified to let that one user's client machine have access. This modification does not affect other client access to the same export, nor does it allow the Management user or client access to other exports.

After the client machine has permission to the export, the user logon affects file access. The client machine presents the UNIX user ID (UID) and group ID (GID) to the server. When the computer accesses a file, the user logon is compared against the typical UNIX permissions of user, group, and other, and typical UNIX access is applied.

NOTE: User credentials are not questioned or verified by the NFS server. The server accepts the presented credentials as valid and correct.

If the NFS server does not have a corresponding UID or GID, or if the administrator has set other conditions to filter out the user, a process called squashing takes effect. Squashing is the conversion of an unknown or filtered user to an anonymous user. This anonymous user has very restricted permissions on the system. Squashing helps administrators manage access to their exports by allowing them to restrict access to certain individuals or groups and to squash all others down to restricted (or no) access. Squashing enables the administrator to allow permissions instead of denying access to all the individuals who are not supposed to have access. See “User and Group Mappings” later in this chapter for specific information about creating and maintaining mappings.

Indicating the Computer to Use for the NFS User Mapping Server

During the processes of starting and installing the NAS B2000, the name *localhost* is assigned by default to the computer. It is assumed that the NAS B2000 is the computer that will be used for user name mapping.

If there are other mapping servers and a machine other than the localhost will store user name mappings, the name of that computer must be indicated, as detailed below:

1. Use Terminal Services to access the MMC, click **File Sharing, Services for UNIX**. Click **Server for NFS**. Figure 9-1 is an example of the Server for NFS user interface.
2. In the **Computer name** box of the user-mapping screen, type the name of the computer designated for user mapping and authentication.
3. *Localhost* is the computer name assigned by default on the NAS B2000. To control user mapping from a different computer, enter the name of that computer.

NOTE: If a machine other than the localhost is to be used, make sure that the user name mapping service is installed and running on that machine.

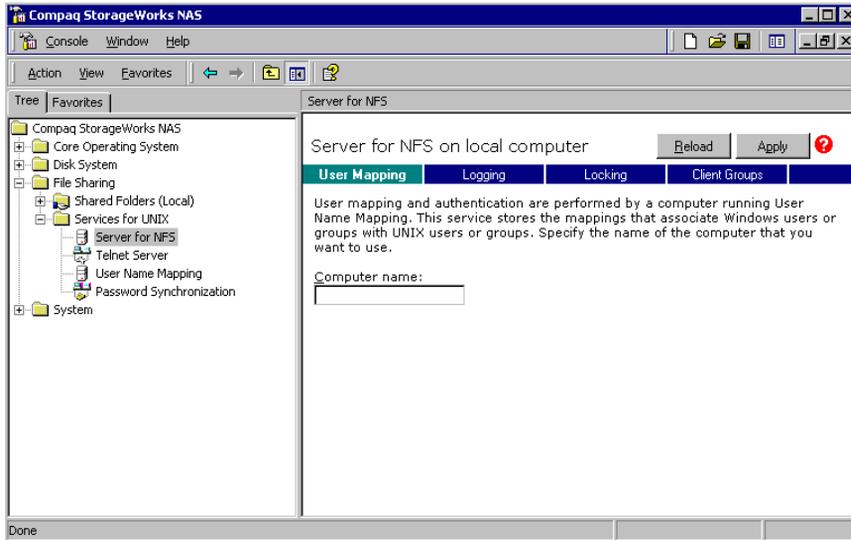


Figure 9-1: MMC Server for NFS screen, User Mapping tab

Logging Events

Various levels of auditing are available. Auditing sends SFU events to a file for later review and establishes log-setting behavior. Some behavior examples include events logged and log file size. See the online SFU help for more information.

1. Use Terminal Services to access the MMC, click **File Sharing, Services for UNIX, Server for NFS**. Click the **Logging** tab.
2. To log selected event types, click the check box for “Log events in this file” on the screen.
3. Enter a filename or use the default filename provided (*rootdrive\SFU\log\nfssvr.log*) and log file size (7-MB default). The default log file is created when the changes are applied.

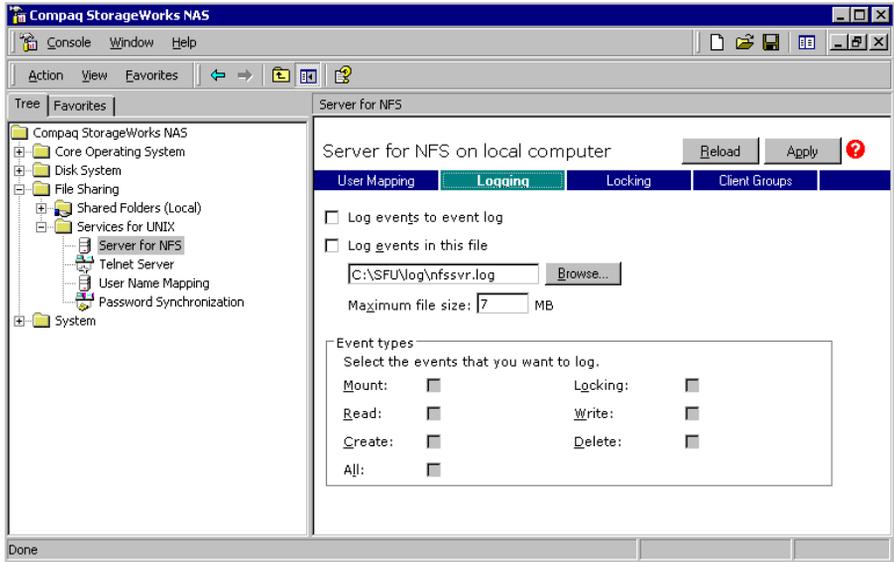


Figure 9-2: MMC Server for NFS screen, Logging tab

Installing NFS Authentication Software on the Domain Controllers

The NFS Authentication software must be installed on all Primary Domain Controllers (PDCs) and backup domain controllers (BDCs) that have Windows users mapped to UNIX users. For instructions on setting up user mappings, see “User and Group Mappings.”

To install the Authentication software on the domain controllers:

1. Locate the *sfucustom.msi* file located in the SFU directory of the NAS B2000.
2. Share out the SFU directory on the NAS B2000.
3. On the domain controller where the service is being installed, using Windows Explorer:
 - a. Connect to the SFU share on the NAS B2000.
 - b. Open the shared directory containing *sfucustom.msi*.
 - c. Double-click the file to open it. Windows Installer is opened.

NOTE: If the domain controller being used does not have Windows Installer installed, locate the file *InstMSI.exe* on the SFU directory and run it. After this installation, the Windows Installer program starts when opening *sfucustom.msi*.

- d. Click **Next** when the Welcome screen is displayed.
- e. Enter the User name and Organization and click **Next**.
- f. Accept the license agreement and click **Next**.
- g. Select **Customized Installation** and click **Next**.
- h. Mark the selections to add **Authentication Tools for NFS** and de-select **Password Synchronization**. To de-select Password Synchronization, expand the drop-down box and select the red “X” next to Password Synchronization. (The entire feature will be unavailable.) The instructions for installing both Authentication Tools for NFS and Password Synchronization are found later in this chapter.
- i. Select the installation directory and click **Next**.
- j. Click **Finish** when installation is complete.

NFS File Shares

NFS file shares are created in the same manner as other file shares, however there are some unique settings. Procedures for creating and managing NFS file shares are documented in the same sections as creating file shares for other protocols. See the “Share Management” chapter for more information.

NOTE: NFS-specific information is extracted from the “Folder and Share Management” chapter and duplicated below.

Complete share management is performed through the **Shares** menu option of the WebUI. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties

Each of these tasks is discussed in this section.

Creating a New Share

To create a new NFS file share:

1. From the WebUI main menu, select the **Shares** directory and then select the **Shares** option. The Shares dialog box is displayed. From the **Shares** dialog box, click **New**. The **General** tab of the **Create a New Share** dialog box is displayed.

NOTE: By default, anonymous access is allowed on NFS shares. To enable anonymous access on an NFS share, use Terminal Services to go to the NAS B2000 desktop. Then, use Windows Explorer, right-click the file share, select Properties, and select NFS Sharing. Select Allow Anonymous Access. Then, click OK.

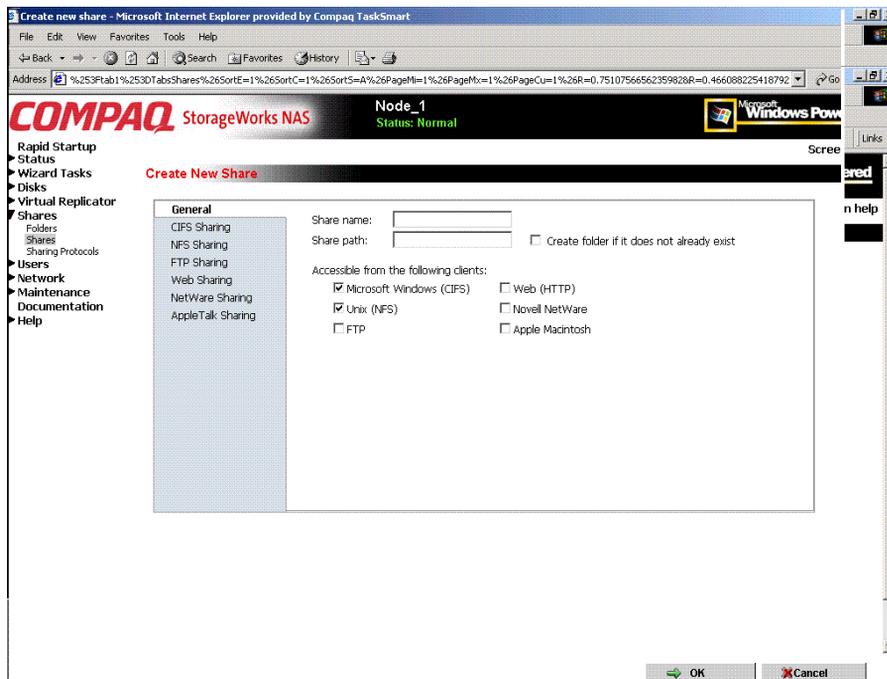


Figure 9-3: Create a New Share dialog box, General tab

2. In the **General** tab, enter the share name and path. Check the Unix (NFS) client protocol check box.

IMPORTANT: NFS service does not support the use of spaces in the names for NFS file shares. NFS translates any spaces in an export into an underscore character. If you plan to use the same name when sharing a folder through CIFS, and then exporting it through NFS, do not put spaces in the CIFS share name.

To create a folder for the share, check the indicated box and the system will create the folder at the same time it creates the share.

3. Select the **NFS Sharing** tab to enter NFS-specific information. See “Modifying Share Properties” for information on this tab.
4. After all share information is entered, click **OK**.

Deleting a Share

IMPORTANT: Before deleting a share, warn all users to exit that share. Then confirm that no one is using the share.

To delete a share:

1. From the **Shares** menu, select the share to be deleted, and then click **Delete**.
2. Verify that this is the correct share, and then click **OK**.

Modifying Share Properties

To change share settings:

1. From the **Shares** menu, select the share to modify and then click **Properties**. The **General** tab of the **Share Properties** dialog box is displayed.

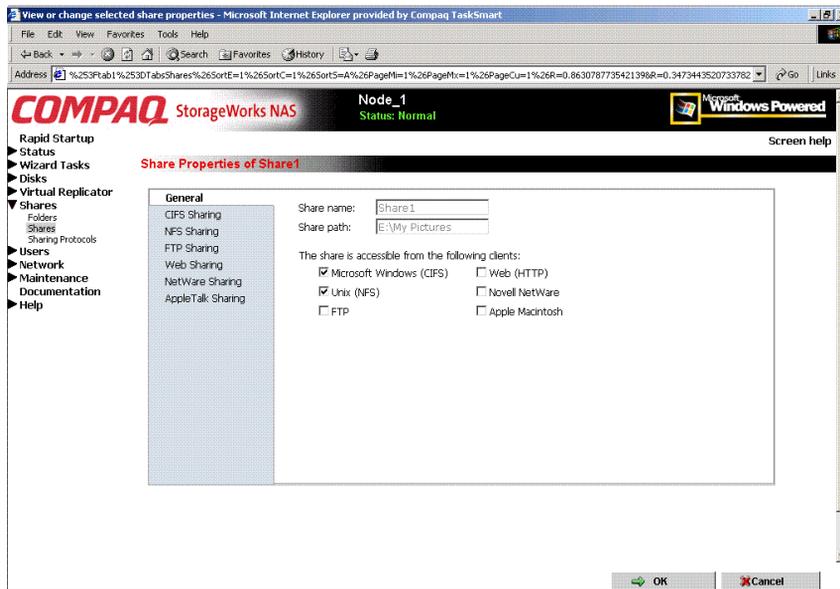


Figure 9-4: Share Properties dialog box, General tab

The name and path of the selected share is displayed.

2. To enter or change client protocol information, check the UNIX (NFS) client-type box and then click the **NFS Sharing** tab.

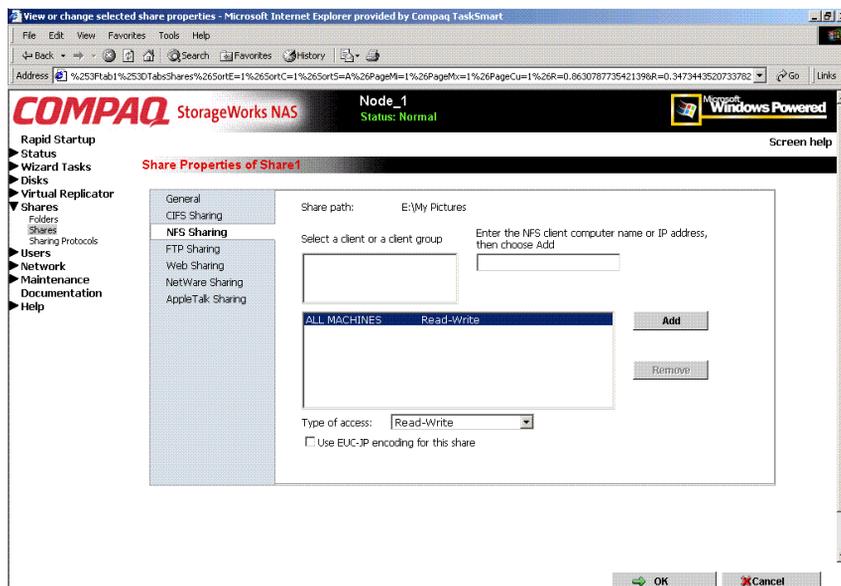


Figure 9-5: NFS Sharing tab

3. From the **NFS Sharing** tab of the **Share Properties** dialog box,

- a. Indicate the allowed clients.

Select the machine to include in the **Select a client or client group** box or manually enter the NFS client computer name or IP address. Then click **Add**.

- b. Indicate the access permissions.

Select the machine from the main user display box and then select the appropriate access method from the **Type of access** drop-down box.

The types of access are:

- **Root**—Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.
- **Read-write**—Use this permission to allow clients to read or write to the share.

- Read-write (EUC-JP)—Use this permission to allow the client to read or write to the share. The EUC is an encoding scheme that is used to represent multibyte characters in UNIX. Set this access if the NFS users use Japanese UNIX systems to access files with read-write permissions on the NAS B2000.
 - Read-only—Use this permission to restrict write access to the share.
 - Read-only (EUC-JP)—Use this permission to restrict write access to the share. The EUC is an encoding scheme that is used to represent multibyte characters in UNIX. Set this access if the NFS users use Japanese UNIX systems to access files with read-only permissions on the NAS B2000.
 - No access—Use this permission to restrict all access to the share.
4. After all NFS sharing information is entered, click **OK**.

NFS Protocol Properties Settings

Parameter settings for the NFS protocol are entered and maintained through the WebUI in the **NFS Properties** dialog box. To access the NFS Properties dialog box, select **Shares**, **Sharing Protocols**. Then, select the **NFS Protocol** radio button and click **Properties**.

The **NFS Properties** menu is displayed.

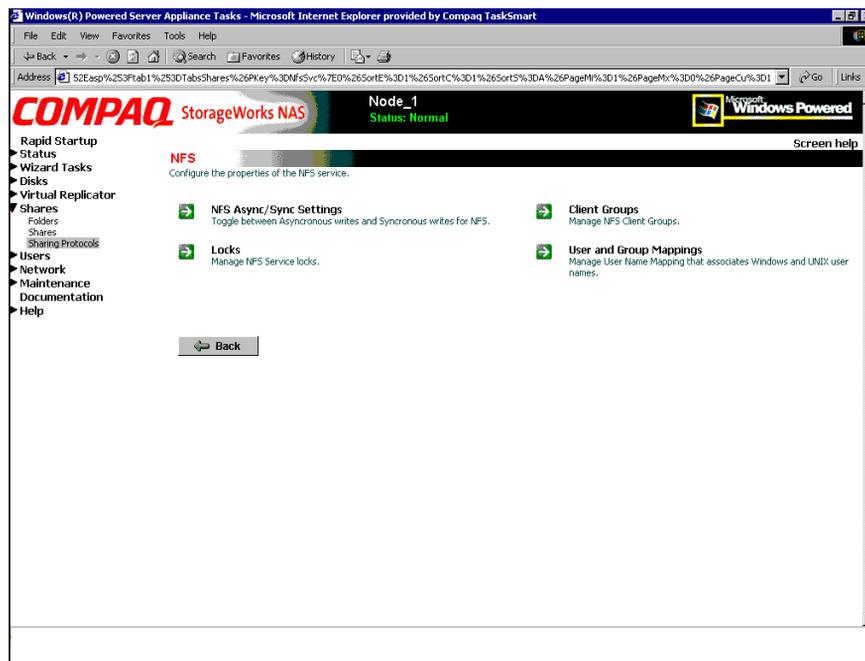


Figure 9-6: NFS Sharing Protocols menu

NFS properties include:

- Async/Sync Settings
- Locks
- Client Groups
- User and Group Mappings

Settings for asynchronous/synchronous writes and service locks are discussed together in the following paragraphs of this chapter.

Client groups and user and group mappings are each discussed in separate sections later in this chapter.

NFS Async/Sync Settings

As mentioned in a previous section, there are two versions of NFS: Version 2 and Version 3. Version 3 supports additional file operations that Version 2 did not have, such as asynchronous file operations.

To indicate whether to use asynchronous or synchronous write settings:

1. From the WebUI, access the NFS Protocol Properties menu by selecting **Shares, Sharing Protocols**. Select **NFS Protocol** and then click **Properties**. The NFS Properties menu is displayed.
2. In the NFS Properties menu, select **NFS Async/Sync Settings**. The NFS Async/Sync Settings dialog box is displayed.
3. Select the desired write setting. The default setting is Synchronous writes.

NOTE: Using synchronous writes allows for greater data integrity. Asynchronous writes will increase performance but will reduce data integrity as the data is cached before being written to disk.

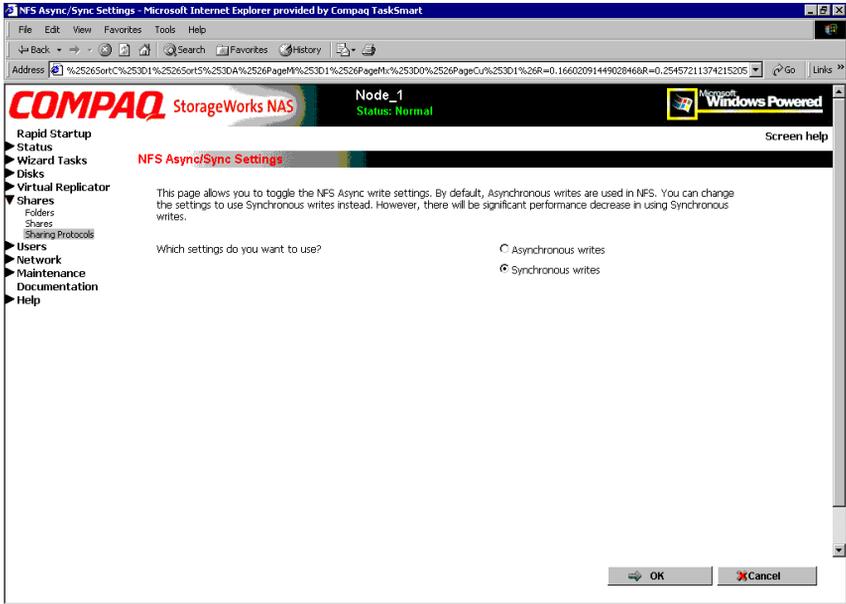


Figure 9-7: NFS Async/Sync Settings dialog box

NFS Locks

NFS supports the ability to lock files. File locking helps prevent two or more users from working with the same files at the same time.

NFS locking depends on the software application components to manage the locks. If an application does not lock a file or if a second application does not check for locks before writing to the file, nothing prevents the users from overwriting files.

To enter locking parameters:

1. From the WebUI, access the NFS Protocol Properties menu by selecting **Shares, Sharing Protocols**. Select **NFS Protocol** and then click **Properties**.

The NFS Properties menu is displayed.

2. In the NFS Properties menu, select **Locks**. The NFS Locks dialog box is displayed. Figure 9-8 is an illustration of the NFS Locks dialog box.

All clients that have locks on system files are listed in the **Current locks** box.

3. To manually clear locks that a client has on files, select the client from the displayed list, and then click **OK**.
4. To indicate the amount of time after a system failure that the locks are kept active, enter the number of seconds in the **Wait period** box.

The NAS B2000 keeps the locks active for the specified number of seconds, while querying the client to see if it wants to keep the lock. If the client responds within this time frame, the lock is kept active. Otherwise, the lock is cleared.

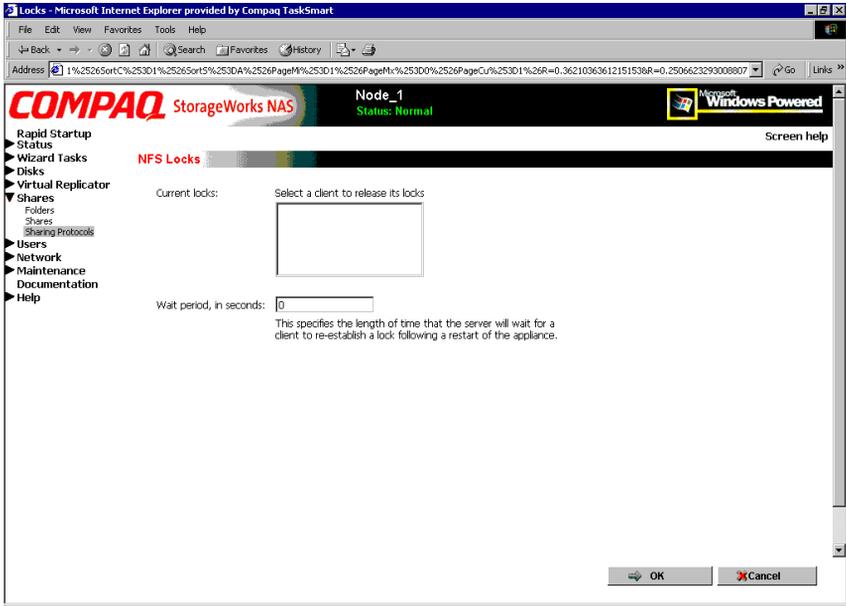


Figure 9-8: NFS Locks dialog box

NFS Client Groups

The Client Groups feature gives administrators a method of assigning access permissions to a set of clients. The administrator creates a client group, gives it a name, and then inserts clients into the group by client name or IP address. After the client group is created, the administrator adds or removes permissions for the entire group, instead of allowing or denying access for each individual client machine.

Proper planning includes control over the naming conventions of client groups and users. If the client group is given the same name as a client, the client is obscured from the view of the server. For example, assume that a client d4 exists. If a client group called d4 is created, permissions can no longer be assigned to just the client d4. Any reference to d4 now refers to client group d4.

To manage NFS client groups:

1. From the WebUI, access the NFS Protocol Properties dialog box by selecting **Shares, Sharing Protocols**. Select **NFS Protocol** and then click **Properties**. The **NFS Protocol Properties** menu is displayed.
2. In the NFS Protocol Properties menu, select **Client Groups**. The **NFS Client Groups** dialog box is displayed.

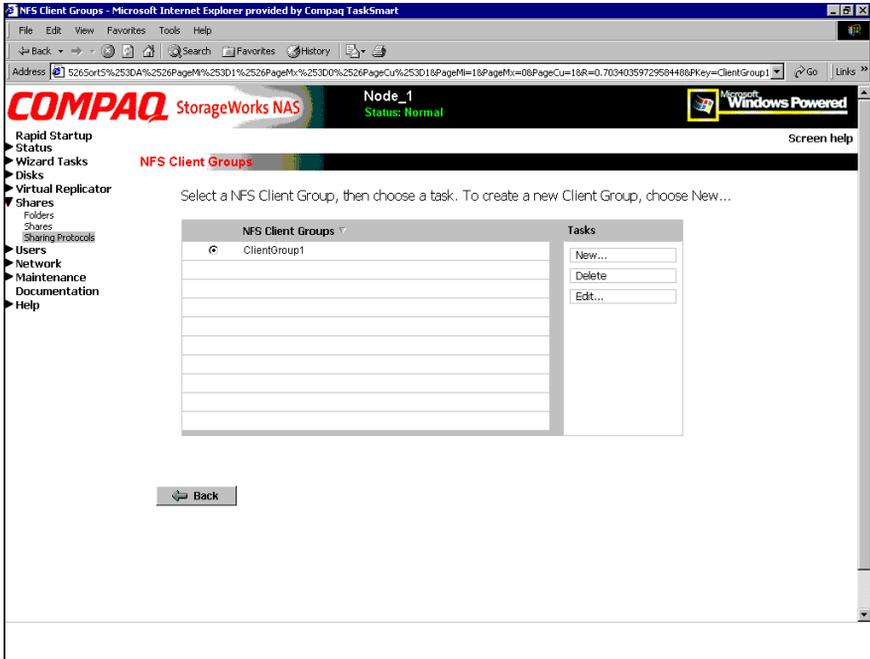


Figure 9-9: NFS Client Groups dialog box

The following tasks are available:

- Adding a new client group
- Deleting a client group
- Editing client group information

Adding a New Client Group

To add a new client group:

1. From the **NFS Client Groups** dialog box, click **New**. The **New NFS Client Group** dialog box is displayed.

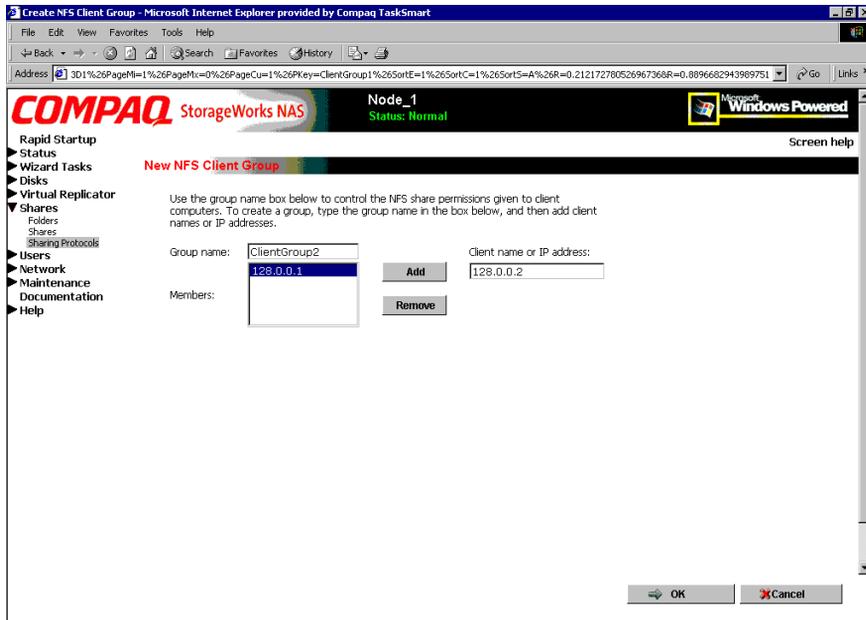


Figure 9-10: New NFS Client Group dialog box

2. Enter the name of the new group.
3. Enter the client name or their IP address.
4. Click **Add**. The system adds the client to the displayed list of members.
5. To remove a client from the group, select the client from the **Members** box and then click **Remove**.
6. After all clients have been added to the group, click **OK**. The NFS Client Groups dialog box is displayed again.

Deleting a client group

To delete a group:

1. From the **NFS Client Groups** dialog box, select the group to delete and click **Delete**.
2. A verification screen is displayed. Confirm that this is the correct group and then click **OK**.

The NFS Client Groups dialog box is displayed again.

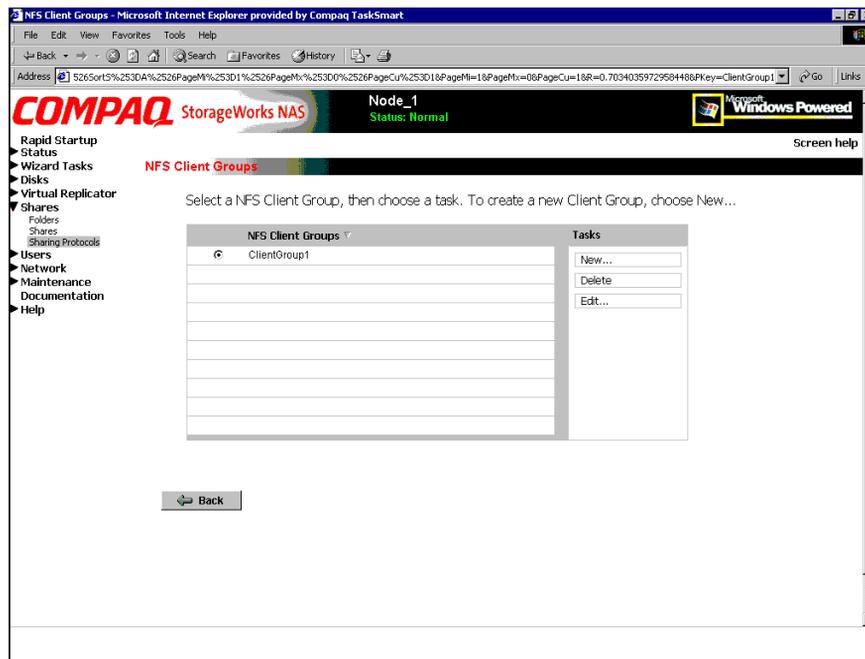


Figure 9-11: Client Groups dialog box

Editing Client Group Information

To modify the members of an existing client group:

1. From the **NFS Client Groups** dialog box, select the group to modify, and click **Edit**.

The **Edit NFS Client Group** dialog box is displayed. Current members of the group are listed in the **Members** box.

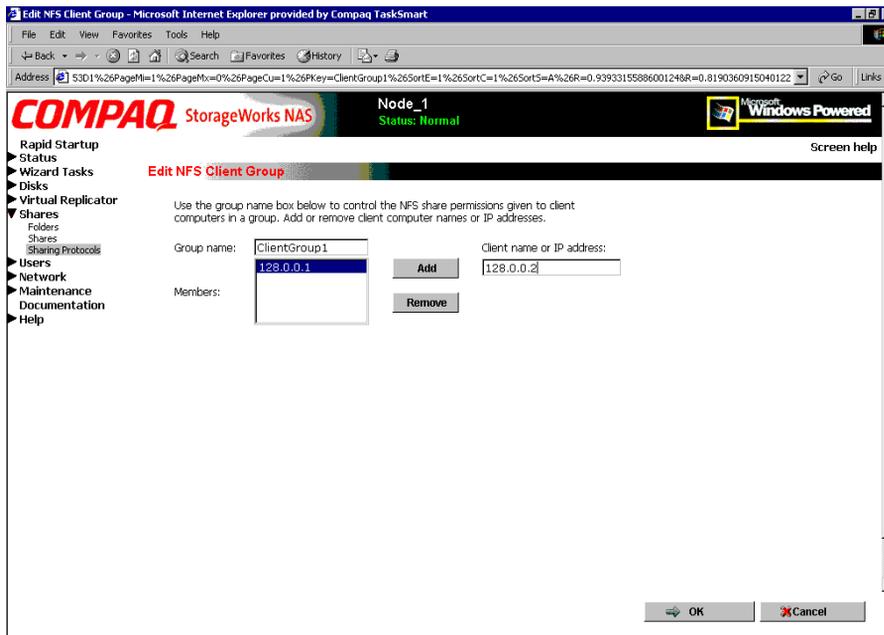


Figure 9-12: Edit NFS Client Groups dialog box

2. To add a client to the group, enter the client name or IP address in the **Client name** box, and then click **Add**. The client is automatically added to the **Members** list.

3. To delete a client from the group, select the client from the Members list, and then click **Remove**. The client is removed from the list.
4. After all additions and deletions are completed, click **OK**. The NFS Client Groups dialog box is displayed again.

NFS User and Group Mappings

When a fileserver exports files within a homogeneous environment, there are no problems with authentication. It is a simple matter of making a direct comparison to determine whether the user should be allowed access to the file, and what level of access to allow.

However, when a fileserver works in a heterogeneous environment, some method of translating user access is required. User mapping is the process of translating the user security rights from one environment to another.

User name mapping is the process of taking user and group identification from one environment and translating it into user identification in another environment. In the context of UNIX and NFS, user and group identification is a combination of a user ID (UID) and group ID (GID). In Windows environments, user identification is a Security ID (SID) or, in Windows 2000, a Globally Unique Identifier (GUID).

The server grants or denies access to the export based on machine name or IP address. However, after the client machine has access to the export, user-level permissions are used to grant or deny access to user files and directories.

The NAS B2000 is capable of operating in a heterogeneous environment, meaning that it is able to work with both UNIX and Windows clients. Because the files are stored in the native Windows NT file system, the server has to map the UNIX users to Windows users to determine the user access level of the files.

NOTE: User mapping is not designed to address existing user database problems in the existing environment. All UIDs and GIDs must be unique across all NIS (Network Information Service) domains and all user names must be unique across all Windows NT domains.

The NAS B2000 supports mappings between one or more Windows domains and one or more NIS domains. The default setup supports multiple Windows NT domains to a single NIS domain. For information about users in multiple NIS domains, refer to the Compaq Supplemental Help section in the SFU online help.

Types of Mappings

There are three types of mappings. These mappings are listed below in order of the most complex (with the greatest level of security) to the least complex (easiest to manage, but with little security):

- Explicit mappings
- Simple mappings
- Squashed mappings

Explicit Mappings

Explicit mappings are created by the administrator to link Windows and UNIX users. They override simple mappings and are used to map users on the different systems that have unique names.

Simple Mappings

Simple mapping is a direct comparison of user names on the Windows system and the UNIX system. If the names match, the user is assumed to be authentic, and appropriate share access is granted. Simple mapping is an option that the administrator must turn on if it is to be used.

Squashed Mappings

If the NFS server does not have a corresponding UID or GID or if the administrator has set other conditions to filter out the user, a process called squashing takes effect. Squashing is the conversion of an unmapped or filtered user to an anonymous user. This anonymous user has very restricted permissions on the system. Squashing helps administrators manage access to their exports by allowing them to restrict access to certain individuals or groups and to squash all others down to restricted (or no) access. Squashing enables the administrator to allow permissions instead of denying access to all the individuals who are not supposed to have access.

The default squash user on Windows is “Anonymous Logon,” but this default user can be changed. For more details on how to change the default squashing user, see the Compaq “OEM Supplemental Help” chapter of the SFU help, found on the NAS B2000.

Figure 9-13 is a diagram showing an example of how the mapping server works for an `ls -al` command.

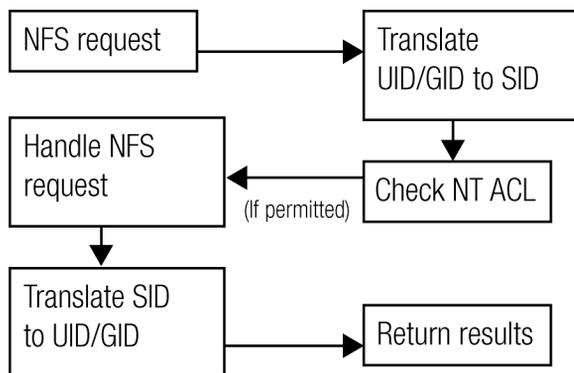


Figure 9-13: Mapping Server “ls -al” Command example

A double translation, as illustrated in Figure 9-13, is sometimes necessary because some commands return user ID information. For example, if the NFS request issued was an `ls -al` command, the return listing of files contains user information (the user and group that own the file). The `ls -al` command is a UNIX command. It returns a long or full listing of all files. Because this information is contained in a Windows NT Access Control List (ACL), it is not UNIX ready. The ACL information has to be converted back to UNIX UIDs and GIDs for the UNIX systems to understand and display the user information.

This second translation is not done for commands that do not return user information. For example, if the NFS request were just to read data from or write data to a file, the second translation would not be performed because there is no returning user information.

User Name Mapping Best Practices

Below is a brief list of suggested practices:

- **Back up user and group mappings.**

To avoid loss of complex advanced mappings in the case of a system failure, back up the mappings whenever the mappings have been edited or new mappings have been added.

- **Map consistently.**

Groups that are mapped to each other should contain the same users and the members of the groups should be properly mapped to each other to ensure proper file access.

Example using User1 and Group1:

- Make sure that the Windows User1 is mapped to the corresponding UNIX User1.
- Make sure that the Windows Group1 is mapped to the corresponding UNIX Group1.
- Make sure that User1 is a member of Group1 on both Windows and UNIX.

- **Map properly.**
 - Valid UNIX users should be mapped to valid Windows users.
 - Valid UNIX groups should be mapped to valid Windows groups.
 - Mapped Windows user must have the **Access this computer from the Network** privilege, or the mapping will be squashed.
 - The mapped Windows user must have an active password, or the mapping will be squashed.

Creating and Managing User and Group Mappings

To set up and manage user name mappings:

1. From the WebUI, select **Shares, Sharing Protocols**. Select **NFS Protocol** and then click **Properties**. The **NFS Properties** menu is displayed.
2. In the NFS Properties Menu, select **User and Group Mappings**. The User and Group Mappings dialog box is displayed.

There are four tabs in the User and Group Mappings dialog box:

- **General information**—Sets the mapping information source, which is either NIS or password and group files.
- **Simple Mapping**—Indicates whether simple mappings are being used.
- **Explicit User Mapping**—Lists exceptional user mappings that will override the simple user mappings.
- **Explicit Group Mapping**—Lists exceptional group mappings that will override the simple group mappings.

Each of these tabs is discussed in the following sections.

3. Enter mapping information on the appropriate tabs, then click **OK**.

NOTE: A wizard is available for entering user name mappings. The wizard is located in the **Users** menu option of the WebUI. See the “User and Group Management” chapter for brief information about the wizard.

General Information

The NAS B2000 stores the mapping data in an NTFS file system. The user name mapping server translates the UNIX users into Windows users so that the server can determine user access rights to the data.

Within this initial screen, indicate whether the source of mapping information is an NIS server or is a special file with password and group information.

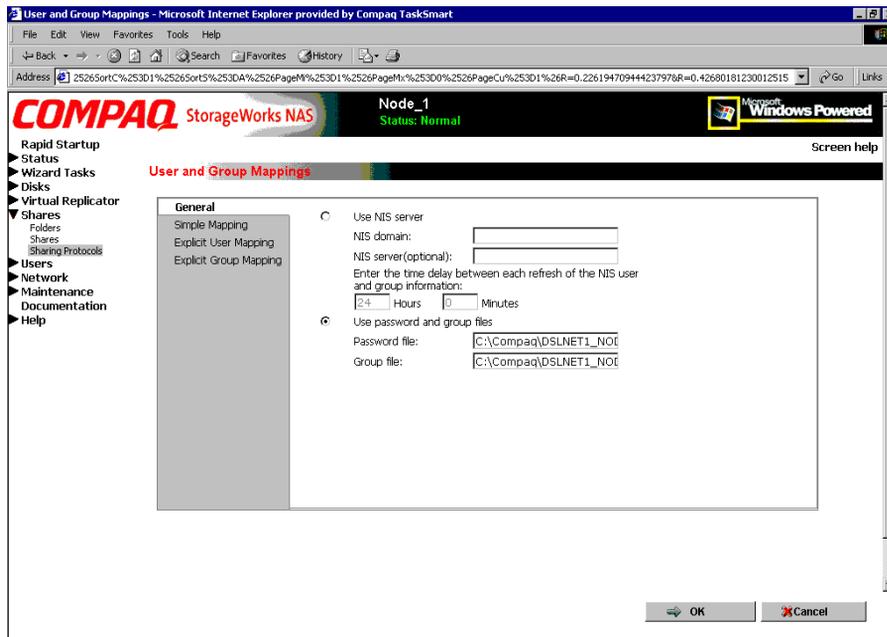


Figure 9-14: User and Group Mappings dialog box, General tab

From the **General** tab of the **User and Group Mappings** dialog box:

1. If an NIS server is being used:
 - a. Select **Use NIS server**.
 - b. Enter the NIS domain name.
 - c. Enter the NIS server name. This field is optional. In the **Hours** and **Minutes** fields, indicate how often the system will connect to the NIS domain to update the user list.
2. If custom password and group files are being used:
 - a. Select **User password and group files**.
 - b. Enter the path and name of the password file.
 - c. Enter the path and name of the group file.
3. After this basic information is entered, click **OK**.

Simple Mapping

Simple (or implicit) mapping is the first level of user name mapping. In simple mode, user and group names that match exactly in name are automatically equated.

While simple mappings are the most easily managed and are the most forthright type of map, security problems can arise. For example, if a UNIX user is coincidentally an exact match of a Windows user, the system will equate them and an inadvertent mapping will occur, granting a user inappropriate access.

To use simple mappings, the feature must be enabled. If this feature is turned off, the administrator must manually create an explicit map for each user.

To enable simple mapping, click the **Enable Simple Mapping** option and then enter the Windows domain name.

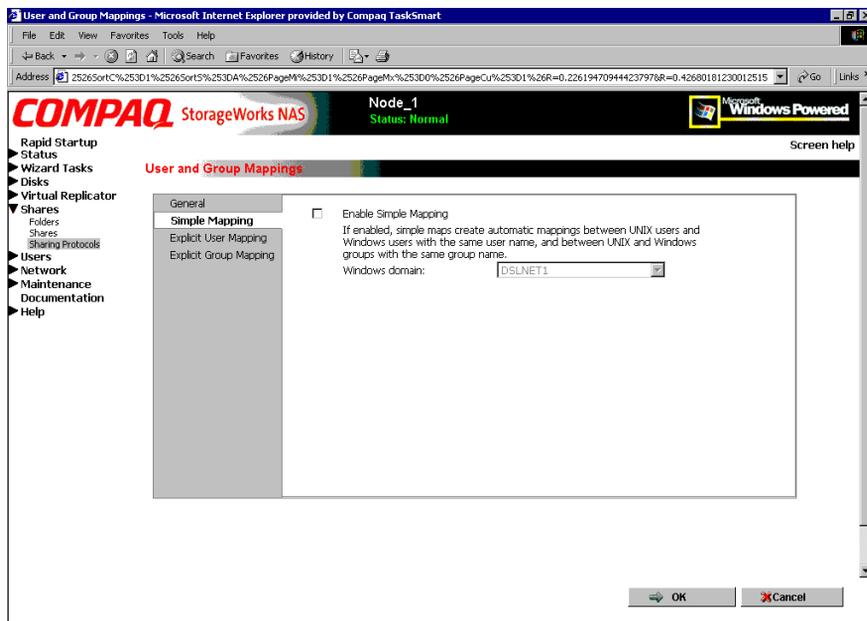


Figure 9-15: User and Group Mappings dialog box, Simple Mapping tab

IMPORTANT: Only one domain can have simple mappings. For information about how to set up simple mappings for multiple domains, see the Compaq “OEM Supplemental Help” chapter of the SFU help, found on the NAS B2000 console.

Explicit User Mapping

Explicit mappings allow the administrator to map any user or group manually to any other user and group. Advanced mappings override simple mappings, giving administrators the capability of using simple mapping for most users and then using advanced mappings for the users with unique names on the different systems. Alternatively, simple mapping can be disabled completely, relying solely on explicit mappings. Explicit mappings create the most secure mapping environment.

Security issues seen in simple mappings do not exist in explicit mappings. Explicit user mappings specifically correlate two users together, thus preventing the inadvertent mapping.

To enter explicit user mappings, select the **Explicit User Mapping** tab. Figure 9-16 is an example of the Explicit User Mapping tab.

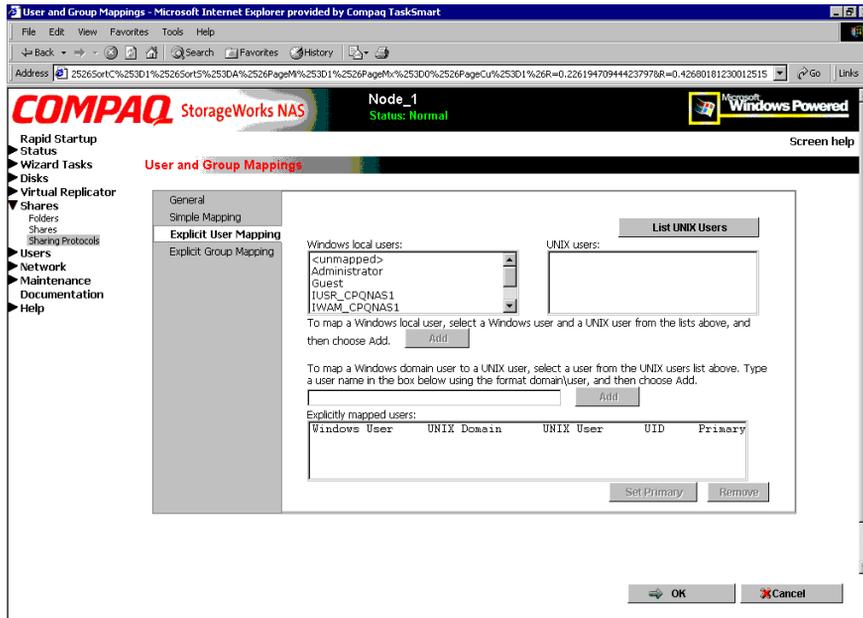


Figure 9-16: User and Group Mappings dialog box, Explicit User Mapping tab

To create explicit user mappings:

1. Click the **List UNIX Users** button to populate the **UNIX users** box.
2. To map a local Windows user to a UNIX user, highlight the Windows user in the **Windows local users** box and highlight the UNIX user that you want to map, and then click **Add**. The **Explicitly mapped users** box at the bottom of the screen is populated with the new mappings. Repeat this process until all desired users have been mapped.

3. To map a domain Windows user to a UNIX user, enter the domain and the user name in the box in the middle of the screen (use the `Domain\username` format) and highlight the UNIX user that you want to map, and then click **Add**. The map is added to the **Explicitly mapped users** box at the bottom of the screen. Repeat this process until all desired users have been mapped.
4. To map multiple Windows users to one UNIX user, one of the mapped Windows users must be set as the primary mapping. To indicate which user map is the primary mapping, highlight the desired map in the **Explicitly mapped users** box, and then click the **Set Primary** button.
5. To delete a map, highlight the map in the **Explicitly mapped users** box, and then click the **Remove** button.
6. After all entries are completed, click **OK** to activate the new entries.

Explicit Group Mapping

To enter explicit group mappings, select the **Explicit Group Mapping** tab. Figure 9-17 is an example of the Explicit Group Mapping tab.

Explicit mappings allow the administrator to map any user or group manually to any other user and group. Explicit mappings override simple mappings, giving administrators the capability of using simple mapping for most groups and then using explicit mappings to make changes to simple mappings. Simple mapping can be turned off for greater security.

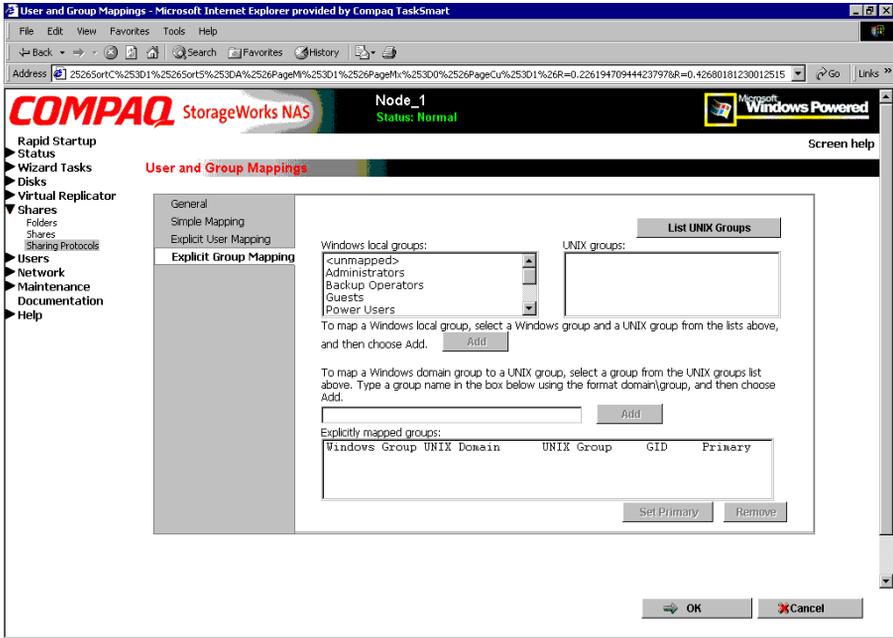


Figure 9-17: User and Group Mappings dialog box, Explicit Group Mapping tab

To create explicit group mappings:

1. Click the **List UNIX Groups** button to populate the **UNIX Groups** box.
2. To map a local Windows group to a UNIX group, highlight the Windows group in the **Windows local groups** box and highlight the UNIX group to map, and then click **Add**. The **Explicitly mapped groups** box at the bottom of the screen is populated with the new mappings. Repeat this process until all desired groups have been mapped.
3. To map a domain Windows group to a UNIX group, enter the domain and the group name in the box in the middle of the screen (use the `Domain\groupname` format) and highlight the UNIX group to map, and then click **Add**. The map is added to the **Explicitly mapped groups** box at the bottom of the screen. Repeat this process until all desired groups have been mapped.
4. To map multiple Windows groups to one UNIX group, one of the Windows groups must be set as the primary mapping. Therefore, to indicate which group map is the primary mapping, highlight the desired map in the **Explicitly mapped groups** box, and then click the **Set Primary** button.
5. To delete a map, highlight the map in the **Explicitly mapped groups** box and then click the **Remove** button.
6. After all entries are completed, click **OK** to activate the new entries.

Backing up and Restoring Mappings

The user name-mapping server has the capability to save and retrieve mappings from files. This capability is useful for backing up mapping settings prior to making changes and for exporting the mapping file from one server to others, using the same mapping information.

The user name-mapping server can save existing mappings to a file or load them from a file and populate the mapping server. This feature is found in the MMC under the Map Maintenance tab of the User Name Mapping screen, as shown in Figure 9-18.

To access the MMC, use Terminal Services. To open a Terminal Services session, from the WebUI, select **Maintenance**, **Terminal Services**.

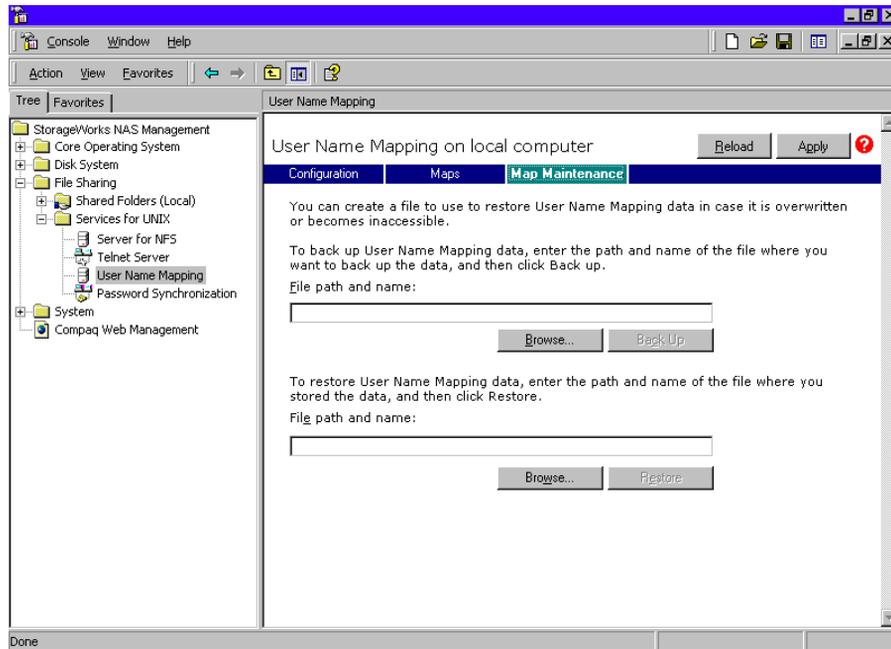


Figure 9-18: MMC User Name Mapping screen, Map Maintenance tab

Backing up User Mappings

1. Select the **Map Maintenance** tab from the **User Name Mapping** screen.
2. Type the path and name of the file to be used for backup in the **File path and name** field or click **Browse** to locate the file.

NOTE: If the file is being created for the first time, follow these steps:

1. Browse to the target directory.
 2. Right-click in the file-listing pane, select New, Text Document. Enter a name for the file and then press Enter.
 3. Double-click the new file to select it.
3. Click **Backup**.

Restoring User Mappings

User mappings can be restored using the following procedures.

1. Select the **Map Maintenance** tab from the **User Name Mapping** screen.
2. Type the path and name of the file in the **File path and name** field or click **Browse** to locate the file.
3. After locating the file, click **Restore**.

NFS File Sharing Tests

Compaq recommends performing the following tests to verify that the setup of the shares, user mappings, and permissions grant the desired access to the NFS shares.

1. Create an NFS share.

See “NFS File Shares” earlier in this chapter for information on creating shares.

2. Verify that the NFS share exists.

Use Terminal Services to log in to the NAS B2000 and access the command line interface:

```
nfsshare <sharename> (sharename represents the name of the share.)
```

3. Map a user.

See “User and Group Mappings” in this chapter for instructions.

4. Verify that the mappings exist.

Use Terminal Services to log in to the NAS B2000 and access the command line interface:

```
mapadmin list -all
```

5. On the UNIX system, use the mapped user to create a file.

- a. As the root user, mount the share:

```
mount -t nfs <nfs server IP address:/nfs share>  
/mount-point
```

- b. Log in as a mapped user.
- c. Change directories to the mount-point directory.
- d. Create the file as the mapped user (example: file1).

6. Verify that the same permissions are set up for the user on both the UNIX side and the Windows side.

- a. List the permissions on the UNIX side:

```
ls -l /mount-point/file1
```

(Example screen display: -r--r----- unixuser1 unixgroup1)

- b. List the permissions on the Windows side: (change to the nfs share directory)

From a command line interface accessed from Terminal Services on the NAS B2000:

```
cacls file1
```

(Example display: DOMAIN1\Windowsuser1:R

 DOMAIN1\Windowsgroup1:R

- c. Compare and verify the permissions from UNIX and Windows.

Terminal Services, Telnet Service, and Remote Shell Service

In addition to the WebUI, three services are available for remote administration of Services for UNIX. These services let users connect to machines, log on, and obtain command prompts remotely. See Table 9-1 for a list of commonly used commands.

Using Terminal Services

Microsoft Terminal Services can be used to remotely access the NAS B2000 desktop. This provides the administrator flexibility to automate setups and other tasks. SFU file-exporting tasks and other SFU administrative tasks can be accomplished using Terminal Services to access the SFU user interface from the MMC or from a command prompt.

Terminal Services is included in the WebUI of the NAS B2000. To open a Terminal Services session, from the WebUI, select **Maintenance, Terminal Services**. See the “Remote Access Methods and Monitoring” chapter for information on setting up and using Terminal Services.

Using Telnet Service

Telnet is a UNIX command-line utility. The Telnet service is included on the NAS B2000, but, by default, it is not activated. To use Telnet services, see the information in the “Remote Access Methods and Monitoring” chapter.

Using Remote Shell Service

The Remote Shell is a UNIX method for allowing UNIX users to run commands remotely. It can be used in a fashion similar to Telnet or can be used to directly invoke a remote command. Remote Shell service is not activated by default. See Chapter 11 for setup and use.

Table 9-1 describes some common SFU commands.

Table 9-1: Command Line Interface Command Prompts

Command	Function
nfsstat /?	Learn about viewing statistics by NFS operation type
showmount /?	View the format of the command to display NFS export settings on NFS servers
showmount -a	View users who are connected and what they currently have mounted
showmount -e	View exports from the server and their export permissions
rpcinfo /?	Learn how to display Remote Procedure Call (RPC) settings and statistics.
mapadmin /?	View how to add, delete, or change user name mappings
tnadmin /?	View how to change Telnet Server settings
nfsshare /?	Learn how to display, add, and remove exported shares

Password Synchronization

Password synchronization is an optional service that automatically synchronizes Windows passwords with UNIX passwords across multiple machines or environments. This service is included on the NAS B2000, but it is not activated.

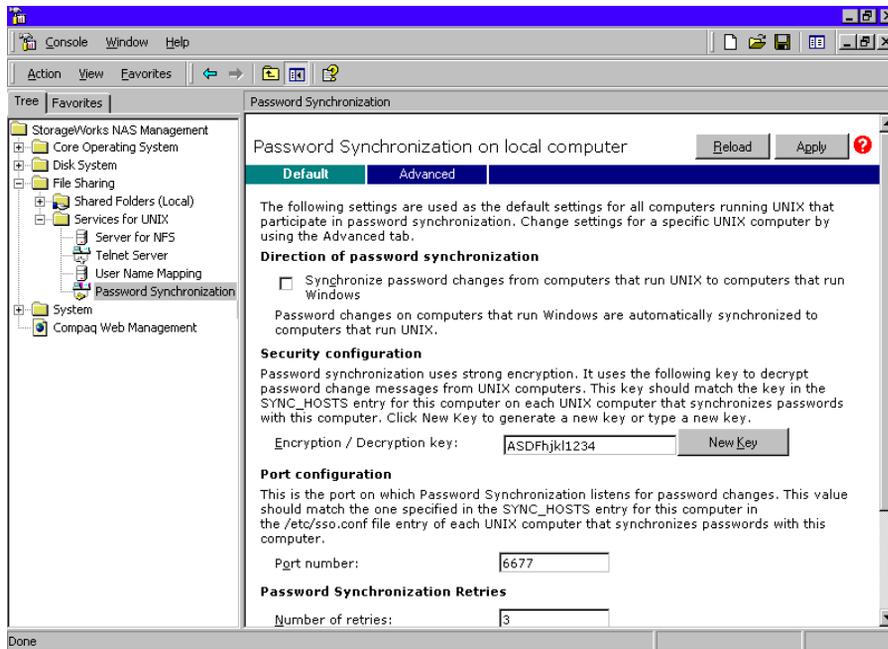


Figure 9-19: MMC Password Synchronization screen

Password synchronization ensures that the machines contain identical and most current user password database. When the user or administrator changes a password, the new password is updated across all target machines.

Without password synchronization, the user could have different passwords on different machines. If the administrator or user changed the password, the change would affect only that single machine.

Password Synchronization Best Practices

- Install Password Synchronization on all domain controllers to ensure consistent synchronization of the Domain and the UNIX passwords.

- Ensure consistent password policies.

If you are providing Windows-to-UNIX password synchronization, make sure the Windows password policy is as restrictive in all areas as the UNIX policy. Failure to ensure that password policies are consistent may result in synchronization failure when a user changes a password on the less restrictive system and the change is rejected by the more restrictive system.

- Avoid synchronizing administrator passwords.

Do not synchronize passwords for members of the Windows Administrator groups or the passwords of UNIX Superuser or Root accounts.

- When Password Synchronization is installed, members of the local Administrators or Domain Administrators group are added to the PasswordPropDeny group, which prevents their passwords from being synchronized. If you add a user to either the Administrators or Domain Admins group, be sure to add the user to the PasswordPropDeny group.
- The `sync_users` statement in the `sso.conf` file on UNIX systems prevents the passwords of Superusers from being synchronized.

Password Synchronization Requirements

For the password synchronization service to function, the work environment **must** meet the following criteria.

- The password policies must be the same on Windows NT and UNIX.
- User and group names must match exactly in spelling. No advanced mapping component exists to correct for any mistakes or differences.
- The UNIX system must be using CRYPT to encrypt its password database. If the UNIX machine is using anything else, such as MD5, the password synchronization service does not work.
- The password synchronization service must be installed on the primary and backup domain controllers. Click the **Advanced** button to select settings other than default.

Implementing Password Synchronization

The password synchronization service is a service residing on the NFS server. The service does not have to be on the same server as the NFS server, but the service is included on each NAS B2000 device. The password synchronization service detects updates on the Windows NT side and transmits the changes to the target UNIX machines, as specified in the service configuration.

To access the password synchronization module on the NAS device, use Terminal Services to access the **MMC**. From the MMC, select **File Sharing, Services for UNIX, and Password Synchronization**.

Configuring Advanced Settings

To configure advanced settings for password synchronization, use the following procedures:

1. Type the name or IP address of the UNIX computer in the Computer Name box.
2. Click **Add** and then click **Configure**. The password synchronization settings dialog box for the specific computer is displayed.

This dialog box allows the user to perform steps such as supplying new encryption keys or changing password synchronization port numbers.



Figure 9-20: MMC Password Synchronization screen, Advanced Settings dialog box

Installing Password Synchronization

The password synchronization service must be installed on all primary domain controllers (PDCs) and backup domain controllers (BDCs) located in a domain, to support the service. The PDCs contain the primary copy of the user passwords.

Password synchronization should be installed by itself. Core SFU components are not needed to install the service on a domain controller.

NOTE: This procedure does not install SFU.

IMPORTANT: Before installing password synchronization, be sure to close all applications and notify connected users that the server is rebooting.

To install Password Synchronization **without NFS Authentication Tools** on a domain controller:

1. Allow the C:\WINNT\bin\SFU directory of the NAS B2000 to be shared:

```
net share SFU=C:\WINNT\bin\SFU
```
2. On the domain controller, connect to the share:

```
net use Z: \\NAS_machine_name\SFU
```
3. Change directories from the domain controller to the root of the connected share of the NAS B2000:

```
cd /d Z:\
```
4. Run the installation program on the domain controller (case sensitive):

```
OemSetup.msi ADDLOCAL=PasswdSync SFUDIR=C:\SFU  
OEMINSTALL=TRUE SOURCELIST=Z:\ /l*v %temp%\sfusetup.log /q
```
5. Restart the domain controller. The domain controller must be restarted manually after installing the password synchronization. If the domain controller is not restarted, password synchronization will not run correctly.
6. Run the Administration User Interface on the domain controller and set up password synchronization:

Click **Start, Programs, Windows Services for UNIX, Services for UNIX Administration**.

To install Password Synchronization and **NFS Authentication Tools** on the domain controller:

1. Allow the C:\WINNT\bin\SFU directory of the NAS B2000 to be shared:

```
net share SFU=C:\WINNT\bin\SFU
```

2. On the domain controller, connect to the share:

```
net use Z: \\NAS_machine_name\SFU
```

3. Change directories from the domain controller to the root of the connected share of the NAS B2000:

```
cd /d Z:\
```

4. Run the installation program on the domain controller in the following order (case sensitive):

```
OemSetup.msi ADDLOCAL=NFSServerAuth SFUDIR=C:\SFU  
OEMINSTALL=TRUE SOURCELIST=Z:\ /l*v %temp%\sfusetup.log /q
```

```
OemSetup.msi ADDLOCAL=PasswdSync SFUDIR=C:\SFU  
OEMINSTALL=TRUE SOURCELIST=Z:\ /l*v %temp%\sfusetup.log /q
```

5. Restart the domain controller. The domain controller must be restarted manually after installing the password synchronization. If the domain controller is not restarted, password synchronization will not run correctly.

Customizing Password Synchronization

Use **Default** to select password synchronization settings. Select different settings for each UNIX host in the Hosts tab.

- **Direction of Password Synchronization**—This option must remain unchecked. Password changes on Windows NT/2000 are always propagated to UNIX computers. Synchronize password changes from UNIX machines to Windows NT/2000.
- **Security configuration**—Password synchronization uses strong encryption for propagating passwords.
- **Encryption key**—Password synchronization comes with a default Encryption Key (displayed). Enter an encryption key of your own, regenerate the key, or do both.
- **Port configuration**—This port is where the password synchronization service checks for password changes. UNIX machines must be configured to use the defined port number.
- **Password Sync Retries**—Select Password Sync Retries to determine how Password Synchronization failures are handled.
- **Logging**—Significant password synchronization events are logged to the event log. Select the option to allow or deny extensive logging.

NetWare File System Management

File and Print Services for NetWare (FPNW) is one part of the Microsoft software package called Services for NetWare. The most common use of the NetWare network operating system is as a file and print server. Customers using NetWare as the platform to host their file and print services have become accustomed to its interface from both a user and an administrator point of view and have built up an investment in NetWare file and print services. File and Print Services for NetWare helps customers preserve their NetWare skill set while consolidating the number of platforms. This reduces hardware costs and simplifies file and print server administration by making the NAS B2000 emulate a NetWare file and print server. FPNW eases the addition of the NAS B2000 into a mixed infrastructure by providing a NetWare user interface (UI) to a Windows 2000-based server; administrators and users see their same, familiar NetWare UI. Additionally, the same logon for clients is maintained without a need for any client configuration changes.

This service also provides the ability to create Novell volumes, which are actually NTFS shares, from which users can map drives and access resources. Novell Login scripts are supported on the NAS B2000 or through an existing NDS (Novell Directory Services) account.

IMPORTANT: IPX\SPX protocol is required on the Novell servers.

Topics discussed in this chapter include:

- Installing Services for NetWare
- Managing File and Print Services for NetWare
- Creating and Managing NetWare Users
- Creating and Managing NetWare Volumes (Shares)

Installing Services for NetWare

The installation of FPNW on the NAS B2000 allows for a smooth integration with existing Novell servers. FPNW allows a Windows 2000-based server to emulate a NetWare file and print server to users, clients, and administrators. This emulation allows authentication from Novell clients, the use of Novell logon scripts, the creation of Novell volumes (shares), the use of Novell file attributes, and many other Novell features.

Additional information on Microsoft Directory Synchronization Services and the File Migration Utility can be found at:

<http://www.microsoft.com/WINDOWS2000/guide/server/solutions/NetWare.asp>

IMPORTANT: The printing capabilities of File and Print Services for NetWare are not supported on the NAS B2000.

To install Services for NetWare:

1. From the desktop of the NAS B2000, click **Start**, navigate to **Settings-Network and Dial-up Connections**, click **Local Area Connection**, and then click **Properties**.
2. Click **Install**. The Select Network Component Type dialog box is displayed. Figure 10-1 is an example of the Select Network Component Type dialog box.

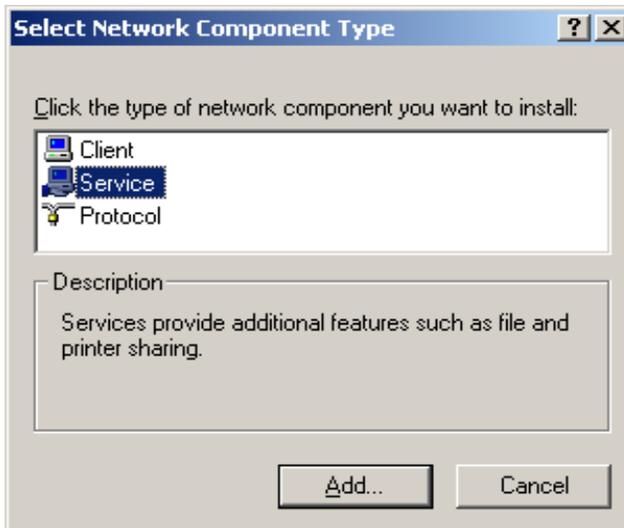


Figure 10-1: Local Area Connection Properties page, Install option

3. Select **Service** and click **Add**.

4. Click the **Have Disk** icon and navigate to the location of Services for NetWare. Services for NetWare is located in the path: c:\compaq\SFN\FPNW\.
5. Select the NETSFNTSRV file and click **OK**.
File and Print Services for NetWare should now appear as an option to install.
6. Select **File and Print Services for NetWare** and click **OK**.

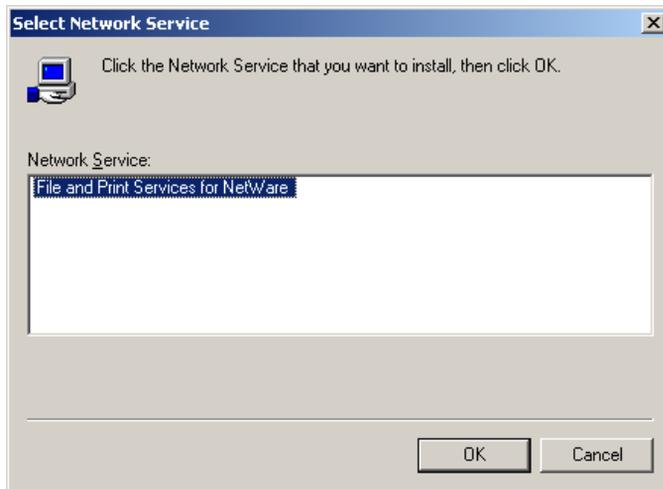


Figure 10-2: Installing File and Print Services for NetWare

Managing File and Print Services for NetWare

To access FPNW:

1. From the desktop of the NAS B2000, click **Start, Settings, Control Panel**, and then double-click **FPNW**.

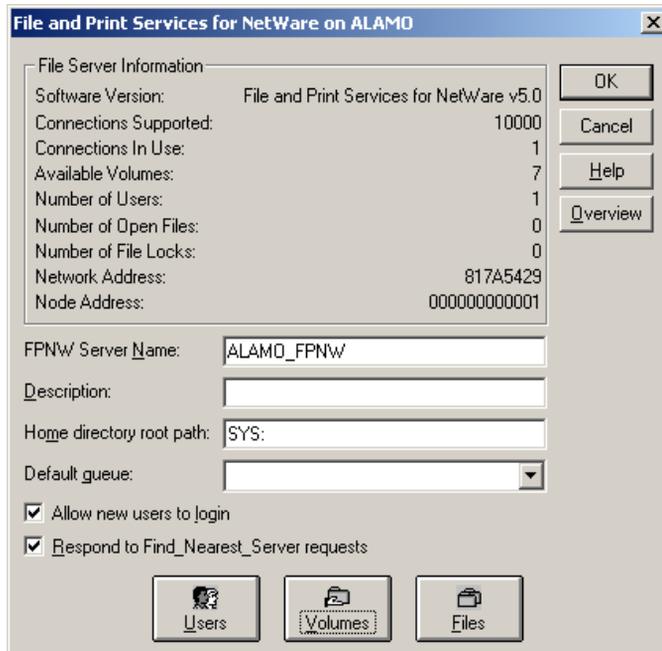


Figure 10-3: File and Print Services for NetWare screen

2. Enter an **FPNW Server Name** and **Description**.

This name must be different from the server name used by Windows or LAN Manager-based clients to refer to the server. If you are changing an existing name, the new name will not be effective until you stop and restart File and Print Services for NetWare. For example, in Figure 10-3 the Windows server name is Alamo and the FPNW server name is Alamo_FPNW.

3. Indicate a **Home directory root path**.

This path is relative to where the Sysvol volume has been installed. This will be the root location for the individual home directories. If the directory specified does not already exist, it must first be created.

4. Click the **Users** button to:

See connected users, disconnect users, send broadcast messages to all users connected to the server, and to send a message to a specific user.

5. Click the **Volumes** button to:

See users connected to specific volume and to disconnect users from a specific volume.

6. Click the **Files** button to:

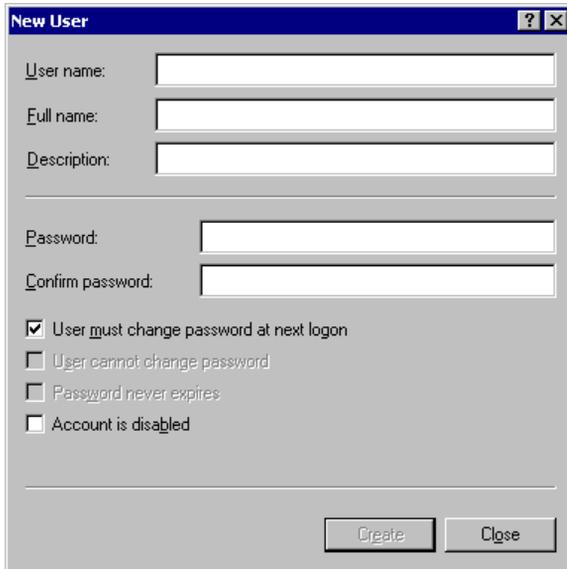
See open files and close open files.

Creating and Managing NetWare Users

To use Services for NetWare, the Novell clients must be entered as local users on the NAS B2000.

Adding Local NetWare Users

1. From the NAS B2000 desktop, click the **NAS Management Console** icon, click **Core Operating System**, and then click **Local Users and Groups**.
2. Right-click the **Users** folder and then click **New User**.



The image shows a 'New User' dialog box with the following fields and options:

- User name: [Text Input]
- Full name: [Text Input]
- Description: [Text Input]
- Password: [Text Input]
- Confirm password: [Text Input]
- User must change password at next logon
- User cannot change password
- Password never expires
- Account is disabled

Buttons: Create, Close

Figure 10-4: MMC New User dialog box

3. Enter the user information, including the user's User name, Full name, Description, and Password. Then, click **Create**.
4. Repeat these steps until all NetWare users have been entered.

Enabling Local NetWare User Accounts

1. In the Users folder (MMC, Core Operating System, Local Users and Groups), right-click an NCP client listed in the right-pane of the screen and then click **Properties**.
2. Select the **NetWare Services** tab.

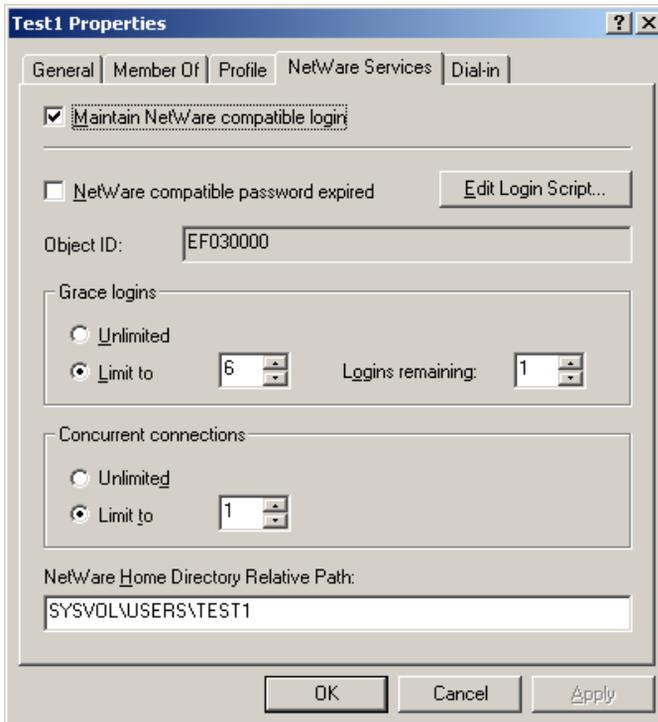


Figure 10-5: NetWare Services tab

3. Select **Maintain NetWare compatible login**.
4. Set other NetWare options for the user and click **OK**.

NOTE: The installation of *File and Print Services for NetWare* will also create a supervisor account, which is used to manage FPNW. The supervisor account is required if the NAS B2000 was added as a bindery object into NDS.

Managing NCP Volumes (Shares)

As with all protocols, before the file shares can be set up, the pools and virtual disks that will contain the data to be shared must be established. Creating storage pools and virtual disks is discussed in Chapter 6.

NCP file shares are created in the same manner as other file shares; however, there are some unique settings. NCP shares can be created and managed through two user interfaces:

- WebUI
- MMC

Procedural instructions for using each of these interfaces are included in the following sections.

Creating and Managing NCP File Shares Using the WebUI

Complete information on managing all types of file shares is documented in the “Shares Management” chapter of this guide. The following information is specific to NCP share management and is extracted from the “Shares Management” chapter and duplicated below.

NOTE: NCP shares can be created only after Microsoft Services for NetWare is installed. See the previous section “Installing Services for NetWare” for instructions on installing SFN.

In addition to the Share Management Wizard of the WebUI, shares can be managed through the **Shares** menu option of the WebUI. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties

Each of these tasks is discussed in this section.

Creating a New NCP Share

To create a new share:

1. From the WebUI main menu, select the **Shares** directory and then select the **Shares** option. The Shares dialog box is displayed. From the **Shares** dialog box, click **New**. The **General** tab of the **Create a New Share** dialog box is displayed.

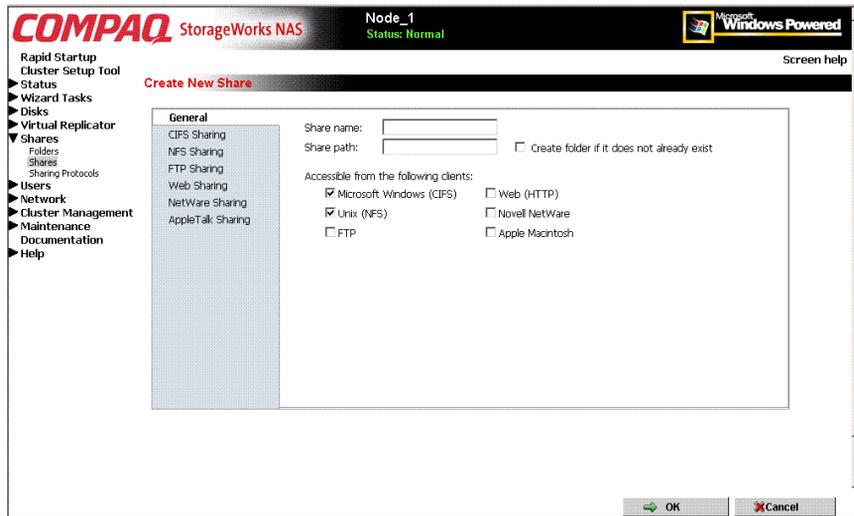


Figure 10-6: Create a New Share dialog box, General tab

2. In the **General** tab, enter the share name and path. Check the Novell NetWare client protocol checkbox.

To create a folder for the share, check the indicated box and the system will create the folder when it creates the share.

3. Select the **NetWare Sharing** tab to enter NCP-specific information. See “Modifying Share Properties” for information on this tab.
4. After all share information is entered, click **OK**.

Deleting an NCP Share

IMPORTANT: Before deleting a share, warn all users to exit that share. Then confirm that no one is using the share.

To delete a share:

1. From the **Shares** menu, click **Delete**.
2. Verify that this is the correct share and click **OK**.

Modifying NCP Share Properties

To change share settings:

1. From the **Shares** menu, select the share to modify and then click **Properties**. The **General** tab of the **Share Properties** dialog box is displayed.

The name and path of the selected share is displayed.

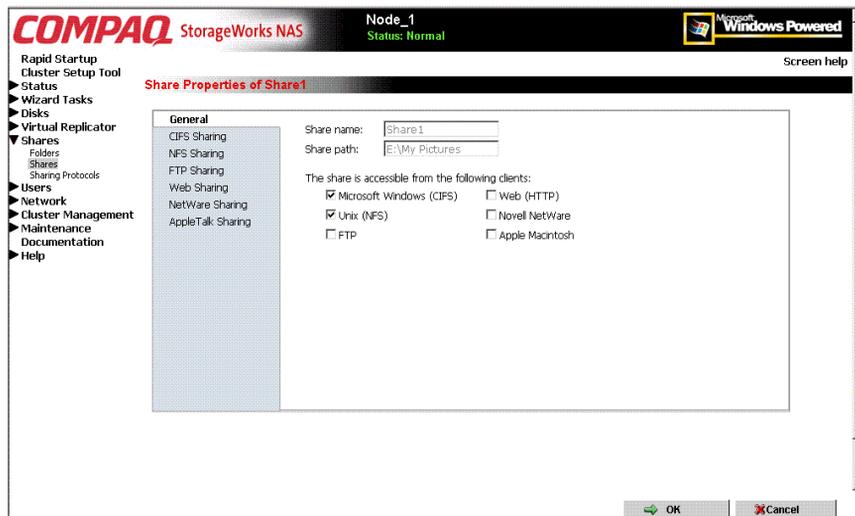


Figure 10-7: Share Properties dialog box, General tab

2. To enter or change client protocol information, check the **Novell NetWare** client-type box and then click the **NetWare Sharing (NCP)** tab.

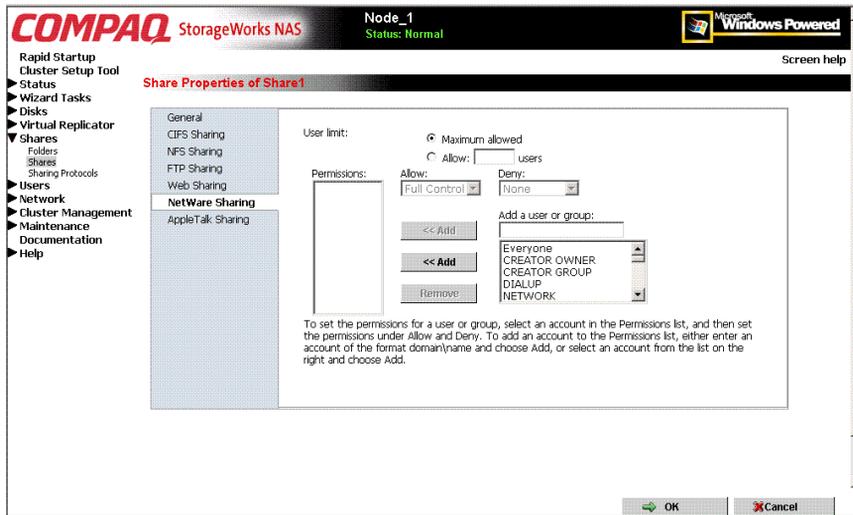


Figure 10-8: Share Properties dialog box, NetWare Sharing tab

3. From the **NetWare Sharing** tab of the **Share Properties** dialog box:
 - a. Enter a user limit.
 - b. Enter Permissions information.

The **Permissions** box lists the currently approved users for this share.

- *To add a new user or group*, either select a user or group from the box at the bottom right of the screen or manually enter the user or group name in the **Add a user or group** box. Then click **Add**. That user or group is added to the Permissions box.
 - *To remove access to a currently approved user or group*, select the user or group from the Permissions box, and then click **Remove**.
 - *To indicate the allowed access for each user*, select the user and then expand the **Allow** and **Deny** drop-down boxes. Then, select the appropriate option.
4. After all NetWare Sharing information has been entered, click **OK**. The **Share** menu is redisplayed.

Creating and Managing NCP Shares using the MMC

In addition to the WebUI available on the NAS B2000, shares can be managed through the Microsoft Management Console (MMC). Tasks include:

- Creating a new share
- Modifying share properties

Each of these tasks is discussed in this section.

Creating a New NCP Share using the MMC

To create a new file share:

1. From the NAS B2000 desktop, click the **Microsoft Management Console** icon, click **File Sharing, Shared Folders**, and then **Shares**.
2. Right-click **Shares**, and then click **New File Share**. The Create Shared Folder dialog box is displayed.

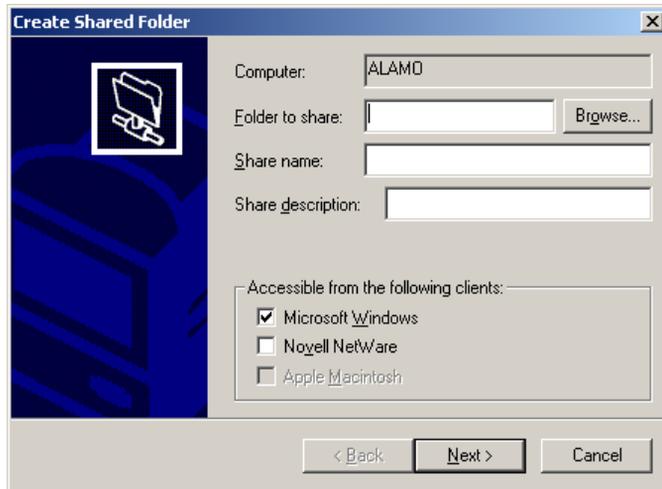


Figure 10-9: Create Shared Folder dialog box

3. In **Folder to Share**, type the path of the directory to be shared.
4. In **Share Name**, type the name of the share. Users will see this name.

5. In **Share Description**, type a description for the share.
6. Select the **Novell NetWare** checkbox and then click **Next**. The dialog box illustrated in Figure 10-10 is displayed.



Figure 10-10: NetWare Basic Share Permissions dialog box

7. Select the appropriate permissions level.

If a custom permissions level is desired, select the **Customize share and folder permissions** radio button and then click **Custom**. The Customize Permissions dialog box is displayed. Figure 10-11 is an illustration of the Customize Permissions dialog box.

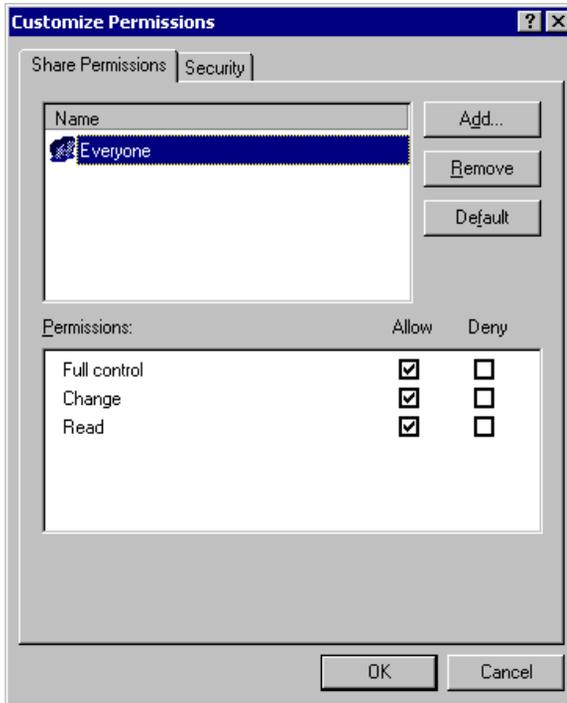


Figure 10-11: Customize Permissions dialog box, Share Permissions tab

8. In the Share Permissions tab, enter choose the appropriate permissions level for each user or group that is configured to have access to that share.
9. To enter file system permissions, select the **Security** tab. The following dialog box is displayed.

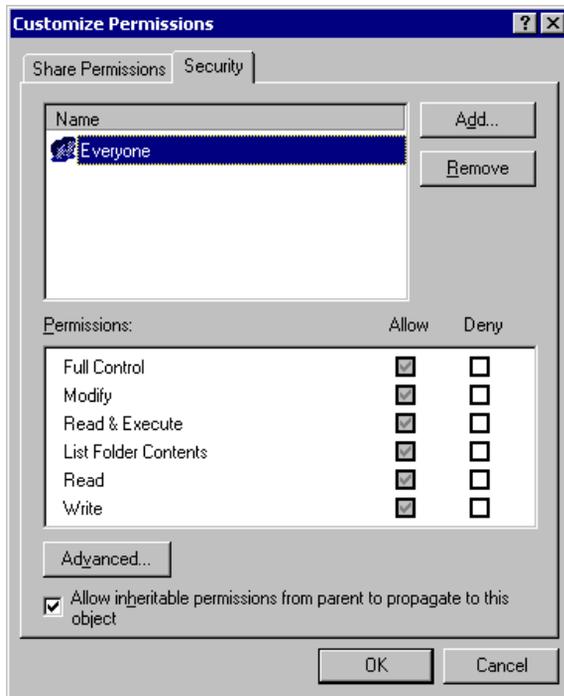


Figure 10-12: Customize Permissions dialog box, Security tab

10. In the Security tab of the Permissions dialog box, enter the file system security properties that apply to the share folder on the server.
11. After the permissions have been entered, click **OK** to return to the Create Shared Folder screens. Click **Finish** to create the share.
12. To create additional shares, click **Yes** at the “Create another shared folder” prompt. Otherwise, click **No** to exit.

Modifying NCP Share Properties using the MMC

To change share settings through the MMC:

1. From the NAS B2000 desktop, select the **Microsoft Management Console** icon and then select **File Sharing, Shared Folders, and Shares**.
2. In the details pane, right-click the desired share and then click **Properties**.
3. Click the **Share Permissions** tab.
4. *To grant permissions to an additional group or user*, click **Add**, select the group or user, and then click **Add**. After any additional groups or users have been added, click **OK**.
5. *To change the permissions granted to the group or user*, select the desired group or user and then select **Allow** or **Deny** for each item.
6. *To remove permissions for the group or user*, select the desired group or user and then click **Remove**.

NOTES:

1. Permissions can be set on a shared volume regardless of its type of file system.
2. Share permissions are effective only when the share is accessed over the network.
3. The group of permissions you set for the share applies equally to all files and subdirectories in the volume.
4. Permissions on an NTFS share operate in addition to NTFS permissions set on the directory itself. Share permissions specify the maximum access allowed.

Remote Access Methods and Monitoring

The *StorageWorks* NAS B2000 comes from the factory with full remote manageability. Several methods of remote access are provided:

- Web-based user interface
- Terminal services
- Remote Insight Lights-Out Edition board
 - Features
 - Remote Insight Lights-Out Edition Board Configuration
 - Using the Remote Insight Lights-Out Edition Board to Access the NAS B2000
- Telnet Server
 - Enabling Telnet Server
 - Configuring Telnet Server
- Remote Shell Daemon
- Compaq Insight Manager
 - Compaq Insight Manager Console
 - Compaq Insight Manager Agent Web Interface

- Enterprise management applications
 - HP OpenView
 - Tivoli NetView (AIX)

These options let administrators use interfaces with which they are already familiar.

Web-Based User Interface

The NAS B2000 includes a Web-based user interface (WebUI) for the administrator to remotely manage the machine. Of all of the remote access methods, the WebUI is the most intuitive and easiest to learn and use.

The WebUI permits complete system management, including system configuration, user and group management, shares management, UNIX file system management, and storage management.

In addition, the WebUI includes wizards to guide the administrator while performing repetitive tasks, such as creating a share.

To access the WebUI:

1. Launch a Web browser.
2. In the URL field, enter:
`http://<your NAS B2000 machine name or IP address>:3201/`

Extensive procedural online help is included in the WebUI.

Terminal Services

The NAS B2000 supports Terminal Services, with a license for two concurrently running open sessions. Terminal Services provides the same capabilities as being physically present at the server console.

Use Terminal Services to access:

- The NAS B2000 desktop
- The NAS B2000 Microsoft Management Console (MMC)
- A command line interface
- Backup software
- Antivirus programs
- Telnet Server
- Remote Shell

To access Terminal Services from the WebUI, select **Maintenance, Terminal Services**. For additional procedural information on Terminal Services, see the “Setup Completion and Basic Administrative Procedures” chapter.

Remote Insight Lights-Out Edition Board

The following information provides an overview of the Remote Insight Lights-Out Edition board capabilities. For further information, refer to the *Compaq Remote Insight Lights-Out Edition Installation and Users Guide* on the Documentation CD.

The Remote Insight Lights-Out Edition board is a PCI-based, single-board computer that uses a Web interface to provide remote management for the server.

Regardless of the state of the host operating system or the host CPU, complete capability for the server is available. A built-in processor, combined with a standard external power supply, makes the Remote Insight Lights-Out Edition board independent of the host server and its operating system. The Remote Insight Lights-Out Edition board provides remote access, sends alerts, and performs other management functions, even when the host server operating system is not responding or the server has lost power.

Features

The Remote Insight Lights-Out Edition board provides the following features:

- Hardware-based graphical remote console access
 - IMPORTANT:** The remote client console must have a direct browser connection to the Remote Insight Lights-Out Edition board without passing through a proxy server or firewall.
- Remote restart
- Server failure alerting
- Integration with Compaq Insight Manager
- Local Area Network (LAN) access through onboard NIC
- Browser support for Internet Explorer 5.50 or later

- Reset and failure sequence replay
- Auto configuration of IP address through domain name system (DNS) or Dynamic Host Configuration Protocol (DHCP)
- Virtual power button

Security Features

- SSL encryption for login and network traffic
- User administration allows capability to define 12 user profiles
- Event generation for invalid login attempts
- Logging of user action in the Event Log

Manage Users Feature

The Manage Users feature allows those with supervisory access to add and delete users or to modify an existing user's configuration. Manage Users also lets the administrator modify:

- User name
- Logon name
- Password
- Simple network management protocol (SNMP) trap IP address
- Receive host OS-generated SNMP traps
- Supervisor access
- Logon access
- Remote console access
- Remote server reset access

Manage Alerts feature

The Manage Alerts feature allows the user to:

- Select alert types received
- Generate a global test alert
- Generate an individual test alert
- Clear pending alerts
- Enable alerts

Refer to the *Remote Insight Lights-Out Edition Board User Guide* for more information about the Remote Insight Lights-Out Edition board features and functionality.

Remote Insight Lights-Out Edition Board Configuration

The Remote Insight Lights-Out Edition board on the NAS B2000 is initially configured through the Rapid Startup Utility. SNMP is enabled and the Compaq Insight Management Agents are preinstalled.

The Remote Insight Lights-Out Edition board comes with factory default settings, which the administrator can change. Administrators may want to add users, change SNMP trap destinations, or change networking settings. Refer to the *Remote Insight Lights-Out Edition Installation and User Guide* for information about changing these settings.

There are several methods for performing Remote Insight Lights-Out Edition board configuration changes:

- Web interface
- Remote Insight Lights-Out Edition board configuration utility access by pressing **F8** during a system restart.

NOTE: You must connect locally with a monitor, keyboard, and mouse.

- Initial Remote Insight Lights-Out Edition board access using the default DNS name

The Remote Insight Lights-Out Edition board is preconfigured by the Rapid Startup Utility, using the following default settings:

- **User Name:** Administrator
- **Password:** (last four digits of the serial number)
- **DNS Name:** RIBXXXXXXXXXXXXX (The 12 Xs are the MAC address of the Remote Insight Lights-Out Edition board.)
- **IP Address:** The IP address entered during system setup

Using the Remote Insight Lights-Out Edition Board to Access the NAS B2000

Using the Web interface of a client machine is the recommended procedure for remotely accessing the server:

1. In the URL field of the Web browser, enter the IP address of the Remote Insight Lights-Out Edition board.
2. Supply an administrator-level user name and password.
The NAS B2000 desktop is displayed.

Telnet Server

Telnet Server is a utility that lets users connect to machines, log on, and obtain a command prompt remotely. Telnet Server is preinstalled on the NAS B2000, but must be activated prior to use.

IMPORTANT: For security reasons, the Telnet Server must be restarted each time the server is restarted.

Enabling Telnet Server

To enable Telnet Server, use Terminal Services to access a command line interface and enter the following command:

```
net start tlntsvr
```

Configuring Telnet Server

To enter Telnet parameter settings, access the Telnet Server user interface. Use Terminal Services to go to the MMC. Then select **File Sharing, Services for UNIX, Telnet Server**.

In the Telnet Server UI, indicate the following:

- Authentication information
- Auditing information
- Server Settings
- Sessions information

Each of these topics is discussed in the following paragraphs.

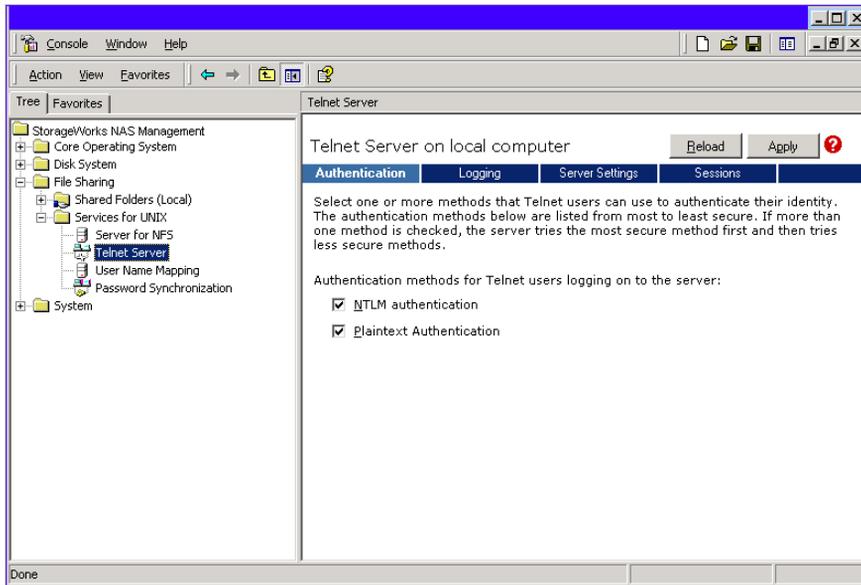


Figure 11-1: Telnet Server interface screen

Authentication Information

The Authentication tab is used to select user authentication methods allowed by the Telnet Server. The administrator determines what method of authentication is appropriate based on work environment.

Auditing Information

Telnet Server can log various events. The Logging tab allows the administrator to enable logging and select the events that should be logged. Note that errors and significant events are always logged to the Windows event list as well.

Server Settings

Use the Server Settings tab to change Telnet Server parameters. These parameters determine how the NAS B2000 Telnet Server operates. For example, one parameter is the number of simultaneous Telnet Server connections that the server allows.

Sessions Information

The sessions screen provides the ability to view or terminate active sessions.

Remote Shell Daemon

The remote shell, commonly referred to as *rsh* in UNIX, is a method for allowing users to access a command prompt or to run a command on another machine. It can be used in a fashion similar to Telnet Server or can be used to directly invoke a remote command.

By default, the Remote Shell is not automatically started on the NAS B2000. The administrator will need to start this service by entering the following command:

```
net start rshsvc
```

NOTE: For security reasons, each time the B2000 is restarted, the Remote Shell service will have to be restarted.

In the following example, the remote shell runs the `ls -al` command on *<server name>* and returns the results to the screen:

```
rsh <server name> ls -al
```

NOTE: An *.RHOSTS* file must be created to allow client access to the server. See the SFU help topic "Rshsvc" on how to create the *.RHOSTS* file.

Currently, SFU implements only the remote command functionality of *rsh*. If a command line is needed, use Telnet Server.

For more information regarding the setup and use of Remote Shell or the Remote Shell service, refer to the online help documentation.

Compaq Insight Manager

The NAS B2000 is equipped with the latest Compaq Insight Management Agents for Servers, allowing easy manageability of the server through Compaq Insight Manager, HP OpenView, and Tivoli NetView.

Compaq Insight Manager is a comprehensive management tool that monitors and controls the operation of Compaq servers and clients. Compaq Insight Manager Version 5.40 or later is needed to successfully manage the NAS B2000. Compaq Insight Manager consists of two components:

- Windows-based console application
- Server- or client-based management data collection agents

Management agents monitor over 1,000 management parameters. Key subsystems make health, configuration, and performance data available to the agent software. The agents act upon that data by initiating alarms in the event of faults. The agents also provide updated management information, such as network interface or storage subsystem performance statistics.

NOTE: The NAS B2000 also supports Compaq Insight Manager XE.

Compaq Insight Manager Console

System monitoring applications such as Compaq Insight Manager allow the administrator to accomplish normal administrative tasks from any remote location with a Web browser. To manage the NAS B2000 using the Compaq Insight Manager console:

1. From the **Setup** menu, access **Discover IP devices**. Then click **New**.
2. Enter the IP address range of the device and then click **Add**.
3. Click **Close** when finished.
4. Click **Find Devices** and select the device to view.
5. Click **Add/Update All Devices**.
6. Double-click the server to show a device information window. The window allows the user to view management data collected by the Compaq agents.

Figure 11-2 is an example of the device information window.

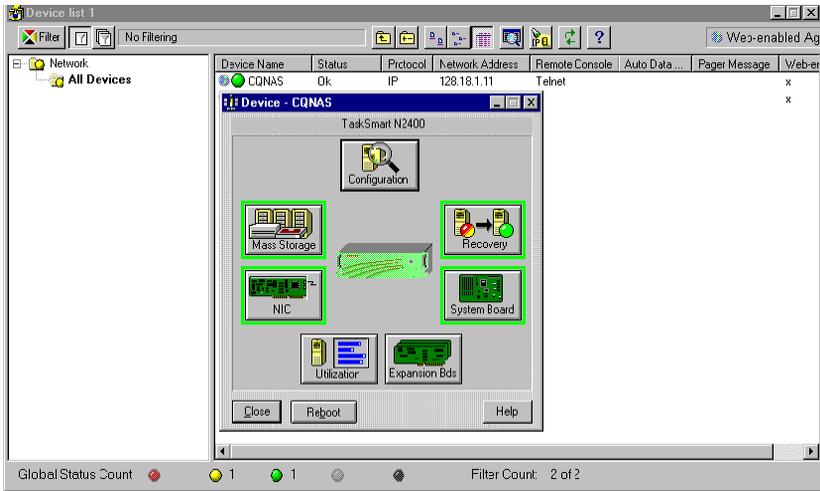


Figure 11-2: Device Information window

For more information about using Compaq Insight Manager, refer to the Compaq Management CD.

Compaq Insight Manager Agent Web Interface

There are two options for accessing the Compaq Insight Manager Agent Web interface.

- From the Compaq Insight Manager console, right-click the device name and select **View Web Data**. The agent Web interface of the server launches in a browser within Compaq Insight Manager.
- Open a Web browser. Enter the server's IP address, using port 2301. An example IP address is `http://122.18.1.14:2301`. The default logon account is "anonymous." Click the account name to log on as an administrator. The user name and password are both `administrator`, lowercase. After the user is logged on as an administrator, the user can change the password.

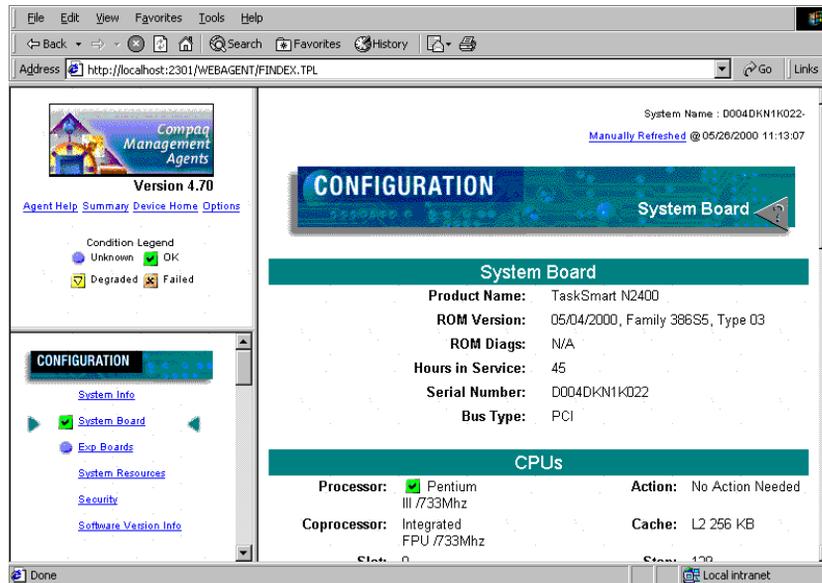


Figure 11-3: Compaq Insight Manager Agent Web interface

For more information about Compaq Insight Manager, see the Compaq Management CD.

Enterprise Management Applications

The following enterprise management applications are installed onto the client machine:

- HP OpenView
- Tivoli NetView

HP OpenView (Windows-Based Operating System)

The NAS B2000 can be managed using HP OpenView by following these steps:

1. Install Compaq Insight Manager for HP OpenView Version 2.0 or later onto the client machine.
2. Modify *CPQCONFIG.DAT* in HP OpenView.

Compaq Insight Manager for HP OpenView, Version 2.0

Compaq Insight Manager for HP OpenView integrates Compaq hardware management and event notification into the HP OpenView Network Node Manager (NNM) network management console.

Depending on the version of Compaq Insight Manager for HP OpenView being used, an addition to the *CPQCONFIG.DAT* may be needed to manage the NAS B2000 from HP OpenView. Because HP OpenView works with different platforms, the location of the *CPQCONFIG.DAT* file varies. Locate the file, add the following line to the file, and save it:

```
Server : ntsrvr : 1.3.6.1.4.1.232.11.2.7.2.1.4.0 string  
"StorageWorks NAS B2000";
```

This command string allows HP OpenView to identify the NAS B2000.

Compaq Insight Manager for HP OpenView introduces an integrated browser launch from the NNM console to the home page of the Web-enabled Compaq Management Agents, as shown in Figure 11-4. The interface is used to collect in-depth information about installed Compaq hardware.

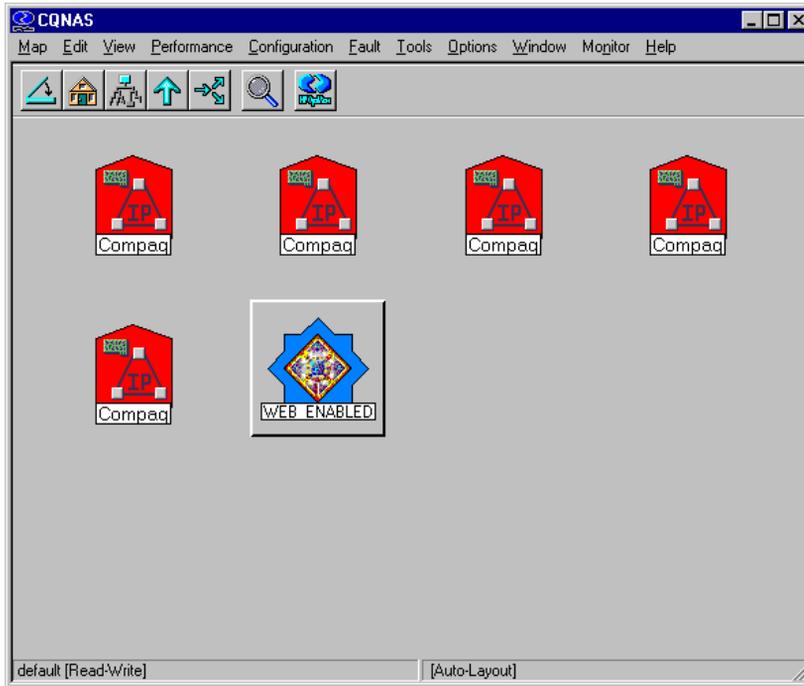


Figure 11-4: Web Enabled interface

The information collected includes system status, system health, prefailure monitors, performance and environmental data, event alarms, and Windows-based OS statistics. Compaq Insight Manager for HP OpenView can be obtained from the Compaq website:

<http://www.compaq.com/products/servers>

Tivoli NetView (AIX)

The NAS B2000 can be managed using Tivoli NetView (AIX). Two steps are required to be able to manage the NAS B2000 through Tivoli NetView.

1. Install Compaq Insight Manager for Tivoli NetView Version 2.0 or later onto the client machine.
2. Modify *CPQCONFIG.DAT* in Tivoli NetView.

Compaq Insight Manager for Tivoli NetView (AIX), Version 2.0

Compaq Insight Manager for Tivoli NetView integrates Compaq hardware management and event notification into the Tivoli NetView network management console. This release introduces an integrated browser launch from the NetView console to the home page of the Web-enabled Compaq Management Agents, which is used to collect in-depth information about Compaq hardware. The information collected includes system status, system health, prefailure monitors, performance and environmental data, event alarms, and Windows-based OS statistics. Compaq Insight Manager for Tivoli NetView can be obtained from the Compaq website:

<http://www.compaq.com/products/servers>

Depending on the version of Compaq Insight Manager for Tivoli NetView being used, an addition to the *CPQCONFIG.DAT* may be needed to manage the NAS B2000 from Tivoli NetView. Often the location of the *CPQCONFIG.DAT* file varies. Locate the file and add the following line to the file and save it:

```
Server : ntsrvr : 1.3.6.1.4.1.232.11.2.7.2.1.4.0 string  
"StorageWorks NAS B2000";
```

This command string allows Tivoli NetView to identify the NAS B2000.

Installing the Management Software on the Client Machine

If a remote machine needs to have remote access to the NAS B2000, it must first have the management software loaded onto it. The following directions provide instruction on loading the Compaq Insight manager console onto the client machines.

Compaq Insight Manager:

1. Insert the Compaq Management CD Version 5.40 or later into the management console of the client machine.
2. From the Compaq Management CD window, click **Compaq Insight Manager**.
3. Select **Compaq Insight Manager** and follow the installation directions.

Compaq Insight Manager for HP OpenView (Windows 2000 Operating System)

For detailed instructions on downloading and installing Compaq Insight Manager for HP OpenView (Windows 2000 operating system), visit the following website:

<http://www.compaq.com/products/servers/management/nt-ov-dl.html>

For detailed instructions on downloading and installing Compaq Insight Manager for HP OpenView (HPUX) visit the management area website:

<http://www.compaq.com/products/servers/>

Compaq Insight Manager for Tivoli NetView

For detailed instructions on downloading and installing Compaq Insight Manager for Tivoli NetView (AIX) visit the management area of the Compaq website:

<http://www.compaq.com/products/servers/>

Backup Utility Management

This appendix is a backup guide for *StorageWorks* NAS devices. This document guides the reader through the process of determining which backup and restore solution is best suited to the NAS device and their business environment.

As a source and a destination for departmental, workgroup, and enterprise data, the NAS B2000 becomes an integral part of company computing environments. Therefore, efficient backup and reliable restore capabilities are a priority.

This appendix provides pointers for setting up and maintaining reliable backups, including:

- Backup solutions
- Best practices

Backup Solutions

There are three main considerations when developing a backup solution:

- System environments
- Hardware options
- Software options

System Environments

In many departmental and workgroup situations, it is common to connect a tape backup device directly to the NAS device, using a SCSI connection. In this scenario, the server has exclusive use of the tape device. Compaq has several tape solutions with wide industry acceptance available for use with the NAS B2000.

The NAS device is deployed into a SCSI direct-connect environment.

SCSI Direct-Connect Environments

The NAS device may be directly connected to a large tape library using an optional SCSI tape controller. The optional High Voltage Differential (HVD) or Low Voltage Differential (LVD) controllers have two SCSI busses, each capable of supporting up to two DLT 7000 (35/70-GB) devices, for a total of four tape drive devices.

Hardware Options

Selecting the correct type device and connection type ensures a reliable backup of data that is well suited to the particular computing environment. Compaq recommends several tape solutions for use with the NAS B2000.

For a full list of qualified tape solutions, refer to the Compaq website:

www.Compaq.com

Additional backup recommendations and information is available in the Backup whitepapers, also available at the Compaq website.

Before purchasing a tape device, ensure that the backup software supports the preferred device. Most backup software supports a wide range of backup devices, and Compaq has done extensive testing and certification on many popular backup software packages. The administrator should confirm specific choices by consulting the software vendor's website. Vendors usually post a hardware compatibility guide for each version of the backup software application.

Software Options

After choosing the tape hardware devices, the next step is to select the backup software. If backup software is already being used on other servers, the same software may be used to reduce the complexity and setup time of the backup solution.

Before purchasing backup software, verify that it is supported on the chosen backup device. Most backup software supports all types of backup devices, and Compaq has done extensive testing and certification on many popular backup packages. The administrator must confirm the specific choice by consulting the software vendor's website. Vendors usually post a hardware compatibility guide for each version of the backup software application.

Important capabilities to look for in backup software include the following:

- Autochanger support
- Tape media management database
- File history database with extensive search capabilities
- Ability to define backup groups and schedules
- Ability to take advantage of multiple tape devices concurrently, to reduce backup window
- Capabilities to analyze, summarize, and report status automatically
- Options for sharing tape drives in a shared library environment
- Options to enable backup of open and locked files

- Options to back up system state and system databases
- Options to interact with software from a remote console application
- Options for disaster recovery

Best Practices

After deciding on a backup solution, establish procedures that will enhance the reliability and effectiveness of the backups. The following sections describe general recommendations for performing a backup. Keep company-specific needs and environment in mind when implementing these suggestions.

Regular and Reliable Backups

The NAS B2000 has a range of high-availability features, including:

- RAID 1 (mirroring) for the operating system drives
- RAID 3/5 for the data drives
- Redundant power supplies and fans
- Redundant HBAs for multiple paths to data
- Snapshots

Despite these features, the only way to reliably safeguard data against accidental loss, intentional tampering, or hardware failures is with regularly scheduled backup and offsite storage of backup media.

Compaq recommends that data disks be configured in RAID arrays. This configuration makes data loss due to disk failure unlikely, because two drives in the same array must fail at the same time for data loss to occur. Although unlikely, such a failure can occur. Backups prevent an inconvenience from becoming a tragedy.

Automated Tape Libraries

Automated tape libraries improve performance, capacity, and reliability of tape backup operations and should be used whenever possible. Libraries must be enabled by additional licensing, installation of library control modules, and configuration steps. Benefits of tape libraries include:

- Enhanced performance by the automated, instantaneous handling of tapes, requiring no lag time for an administrator to arrive and manually change the tape.
- Improved capacity because tape libraries include storage slots for additional tape cartridges. Enough media can be loaded so that operations can continue overnight, over the weekend, or all week, without intervention or tape changes.
- Increased reliability because tapes are handled less and the human element of forgetfulness in changing tapes is eliminated.

Multiple Backup Devices

To take advantage of multiple backup devices, the server must be configured correctly. Generally, backing up multiple disks requires multiple tape drives. If the NAS device has 500 GB of disk space and this space is arranged as a single virtual disk, it is not possible to directly take advantage of multiple tape drives. If possible, make multiple, smaller virtual disks. This procedure lets the administrator back up the multiple devices in parallel, sending the data from one or two disks to each tape in parallel. This type of configuration greatly reduces the time required for backup and makes the most efficient use of the tape backup device.

If it is necessary to use a single, large virtual disk, the administrator configures several backup groups to contain the various directory trees at the root of the virtual disk, so that more than one tape device can work in parallel.

Also, note how the virtual disks are constructed when setting up backup jobs. To increase the performance of the backups, schedule the back up of virtual disks so that disks that share a common set of physical drives are scheduled at different times. The underlying physical disks can devote more time to each of the backup jobs, rather than having two backup jobs competing for disk I/O.

Backup Schedules

An automatic, periodic backup is much more reliable than occasional backups that occur only when someone remembers to execute them. The specific needs of the organization will determine what type of schedule to implement.

A weekly or biweekly full backup is the basis of any good backup schedule. Add to that baseline daily incremental or differential backups to capture any daily changes that occur between full backups. Depending on the rate of data change, and the capacity and performance of the backup devices, adjust the backup schedule to fit the environment of the organization. Incremental backups capture changes to the data that have occurred since the last backup. Differential backups capture all the changes that have occurred since the last full backup.

If the backup devices do not have sufficient capacity for a complete, full backup, distribute the backups so they occur throughout the backup cycle. This strategy can meet the backup needs of the organization until a larger tape backup device or library can be installed. For example, instead of doing a full backup of disks C:, X:, Y:, and Z: on Friday, back up C: on Monday, X: on Tuesday, Y: on Wednesday, and Z: on Thursday. Schedule incremental or differential backups on the same distributed schedule.

NOTE: The suggested scenarios for backup times are based on a hypothetical company situation.

Media Rotation

Most backup software solutions are equipped to label and track media usage accurately. Take advantage of these capabilities to maintain different media pools for full backups and incremental/differential backups, as well as archive media. The retention time on each of these types of backup is different. For example, using differential backups on the same tape as for full backups causes the tape space to be wasted after the retention time for the differential data has passed. Keep separate pools to avoid this problem.

Offsite Storage

Set up a regular process for moving important long-term media, such as backups and archives, offsite for safekeeping. This ensures that the administrator can recover the data in the event of a complete facility destruction where the NAS device resides. As an alternative to a commercial offsite storage facility, if the company has multiple buildings, the offsite media can be stored in another building. This alternative provides some protection in the event of a building fire where the NAS device is located.

When employing offsite storage, strike a balance between safety and convenience by deciding how long to keep the media onsite. After the media has been moved offsite, restores will take much longer because the media is not readily available.

A periodic audit of the offsite facility ensures that the media is being stored in secure, environmentally acceptable conditions, and that it can be located and returned to the facility in a timely manner.

Server Setup Information Archival

After the administrator has established a regular backup schedule, it is necessary to document the setup attributes of the NAS device. There is always the possibility of the complete loss of the server in cases of fire, flood, or weather disasters. To maximize the ability to recover from server disasters and to minimize the time required for recovery, keep current copies of the following information in a safe location:

- Server name
- IP addresses
- Gateways
- DNS servers
- NIS servers
- User mapping database
- Storage setup
 - Virtual Replicator pool names, sizes and member storage units (LUNs)
 - Virtual disk names and sizes
 - Snapshot names and schedules
 - Share names, paths, and access permission settings

This information greatly increases the ability to quickly and accurately recover from catastrophic failures such as fires, weather disasters, theft, and complete hardware failure.

Snapshots and Quick Online Restores

The Virtual Replicator aggregates disk RAID arrays into large pools of storage from which virtual disks are created. VR also enables the snapshot capability for the NAS device. Snapshots are temporary, online copies of the virtual disks. When a snapshot is made, an identical copy of the original source disk is created without any additional disk space being utilized. Snapshots can be completed in a few seconds, because disk space is used only when the original files change.

Snapshots can be used as a convenient source of data for a backup. There are some applications that must be stopped before backups are made. A backup requires that the file system is recorded in a consistent state, where no changes occur during the backup. Because snapshots are created in a matter of seconds and maintain a consistent view of the file system from that point on, snapshots can drastically reduce the amount of time applications must be paused or shut down during backup operations. The NAS device facilitates automatically creating snapshots at any given time, and can even be set up to create a snapshot, execute a backup, and delete the snapshot upon successful completion of the backup job.

IMPORTANT: Snapshots should be considered an additional convenience for restores, not a replacement for tape backup. In the event of disk failures, snapshots can be lost along with the original virtual disk data. Snapshots will be automatically deleted without warning by the Virtual Replicator to regain space when disk space is low.

Although snapshots should never be considered a replacement for regular data backup to removable media, they can be a highly convenient feature for immediate, tapeless recoveries. If a file is accidentally deleted or corrupted, it can be recovered quickly by accessing the snapshot, selecting the file or directory, and copying it back to its original location on the virtual disk.

To use the snapshot capability for a quick online restore, take a snapshot on a regular basis or before the source disk is altered. This ensures a backup of all the original files, applications, and configurations.

Readiness Testing

Completing regular backups is important, but it is only the first step in the backup process. To verify the integrity of those backups, the administrator must conduct periodic testing to confirm the ability to recover files and directories. Regularly testing the recoverability of random files or directories ensures that the backup solution is working as planned.

Disaster Recovery

Disasters that cause the loss of an entire server or server operating system drives require a complete restoration of the server. The specific procedure for recovering from a disaster depends on your environment, the backup software, and the optional disaster recovery modules that may have been installed.

In general, it is necessary to complete the following steps to fully recover from a disaster:

- Use the QuickRestore process to reinstall the NAS B2000 system image.
- Configure the arrays, logical drives, pools, and virtual disks as they were at the time of the disaster.
- Reinstall the backup application.
- Add the NAS B2000 into the appropriate domain.
- Re-establish user accounts if the NAS B2000 is a part of a workgroup.
If the files were retained, re-establish user rights to drives, files, and directories.

NOTE: Being a part of a domain negates this requirement.

- Recover the backup application file and media history databases.
- Recover data from backup applications.
- Recover the system state.
- Re-create file shares.

A

- accelerator ratio
 - ACU (Array Configuration Utility) 5-18
- access control list See ACL
- Access Control Settings
 - Auditing tab dialog box for folder name
 - NTSF Test, illustrated 8-16
 - dialog box for folder name NTSF Test
 - Permissions tab, illustrated 8-14
 - Owner tab dialog box for folder name
 - NTSF Test, illustrated 8-18
- Access Control Settings, advanced
 - adding user or group 8-15
 - auditing tab 8-16
 - Owner tab 8-18
 - permissions tab 8-14
 - removing user or group 8-15
- ACL (Access Control List)
 - described 8-20
 - permissions 8-20
 - security, integrating 8-20
 - user and group access 8-20
- Active Directory, domain environment 7-3
- ActiveAnswers
 - website 1-9, 7-3
- ACU (Array Configuration Utility)
 - accelerator ratio 5-18
 - controller
 - settings screen 5-16
 - controller settings, ACU 5-16
 - expand priority 5-18
 - main configuration screen 5-15
- ACU (Array Configuration Utility)
 - configuring the array controller 5-12
 - creating logical drives 3-7
 - features 5-12
 - graphical configuration utility 5-12
 - managing drive arrays 5-12
 - rebuild priority 5-17
 - storage management 5-12
 - view screen with two arrays,
 - illustrated 5-25
 - wizards 5-12
- ACU (Array Configuration Utility)
 - Controller Settings dialog box,
 - illustrated 5-17
- ACU (Array Configuration Utility) Physical
 - Drive View, illustrated 5-15
- adaptive load balancing See ALB
- ADG See RAID ADG
- administrative procedures
 - listed 2-1
- ADU (Array Diagnostic Utility), drive
 - failure 5-6
- advanced data guarding See RAID methods
- AFP shares See shares, AFP

ALB (adaptive load balancing)
 defined 2-10
 teaming option, Cisco Fast
 EtherChannel 2-10
 using 2-10
AppleTalk shares See shares, AFP
archival and recovery, data A-8
array
 benefits of 3-4
 limitations 3-7
Array Configuration Utility See ACU
array controller
 configuring 5-12
 firmware 5-9
 moving drives 5-9
Array Diagnostic Utility See ADU
arrays
 capacity expansion 4-51
 moving 5-10
 moving drives
 conditions 5-9
 steps 5-9
auditing
 viewing and maintaining logs 2-24
Auditing Entry dialog box for folder NTFS
 Test, illustrated 8-17
auditing, logging events 9-7

B

backing up and restoring, mappings 9-37
Backup Operation Information screen,
 illustrated 6-47
backup solutions
 catastrophic failures, listed A-8
 disaster recovery A-10
 media rotation A-7
 tape devices
 described A-2
 hardware compatibility A-3

backups
 considerations A-1, A-2
 hardware options A-2
 incremental 6-47
 recommendations A-4
 recommended schedules A-6
 restoring from a snapshot 6-43, A-8
 snapshots, creating 6-25
 software options A-3
best practices
 physical storage 3-15
 virtual storage 3-24
block of data 3-6

C

caching, CIFS file shares 8-29
capacity expansion
 configuration 4-51
 defined, ACU (Array Configuration
 Utility) 5-26
 options 4-51
 procedure 4-51
 procedure, ACU (Array Configuration
 Utility) 5-26
 process information, power loss, ACU
 (Array Configuration Utility) 5-28
cautions
 data backup 5-9, 5-10
 defined xv
 drive bays 5-8
 drive defragmentation 3-20
 embedded smart array controller
 configuration 5-16, 5-19
 grounding xvi
 hard drives 5-3
 replacing 5-10
 Lifeguard service 3-22
 logical drive
 extending 5-12
 LUNs 6-24
 overheating xvi
 setting pool policies 6-11, 6-49

- smart array controller
 - configuration 5-13
- snapshot recovery 6-43
- snapshots 3-21
- ventilation clearances xvi
- virtual disk management 6-16
- CIFS (Common Internet File System)
 - shares See shares, CIFS shares
- Cisco Fast EtherChannel See FEC
- client groups See NFS, client groups
- Command Line Interface Command
 - Prompts, table 9-41
- command prompt, commands, NFS 9-41
- Compaq authorized resellers, telephone numbers xvii
- Compaq Insight Manager
 - Agent Web interface 11-13
 - Agent Web interface, illustrated 11-13
 - components 11-11
 - console Web browser 11-11
 - detecting drive failure 5-6
 - discover IP devices 11-12
 - installing management software 11-17
 - management agents
 - network interface 11-11
 - storage subsystem
 - performance 11-11
 - remote administration 11-10
- Compaq Insight Manager for HP OpenView, remote administration 11-14
- Compaq Insight Manager for Tivoli NetView (AIX), Version 2.0, remote administration 11-16
- Compaq Network Teaming and Configuration utility See CPQTeam
- Compaq website xvii
- Compaq website, telephone numbers xvii
- component-level repairs, warning xvi
- configuration
 - capacity expansion, ACU (Array Configuration Utility) 5-26
 - creating a new array, ACU (Array Configuration Utility) 5-19
- controller
 - settings screen
 - accelerator ratio, ACU (Array Configuration Utility) 5-18
 - expand priority, ACU (Array Configuration Utility) 5-18
 - rebuild priority, ACU (Array Configuration Utility) 5-17
 - settings, ACU (Array Configuration Utility) 5-16
- CPQTeam
 - installing the utility 2-4
 - opening the utility 2-5
- CPQTeam dialog box, illustrated 2-11
- CPQTeam installation, complete, illustrated 2-5
- CPQTeam Properties dialog box, illustrated 2-6
- CPQTeam utility icon, illustrated 2-5
- create
 - a new array
 - creating logical drives, ACU (Array Configuration Utility) 5-23
 - manually, ACU (Array Configuration Utility) 5-19
 - logical drive screen
 - logical drive size, ACU (Array Configuration Utility) 5-24
 - logical drives, ACU (Array Configuration Utility) 5-23
- Create a New Folder dialog box, General tab, illustrated 8-6
- Create a New Share dialog box, General tab, illustrated 8-26, 9-10, 10-10
- Create a New Virtual Disk dialog box, illustrated 6-17
- Create Logical Drive dialog box, illustrated 5-23
- Create New Group dialog box, General tab, illustrated 7-16
- Create New Pool dialog box, illustrated 6-9
- Create New Share dialog box, General tab, illustrated 8-10

- Create New User dialog box,
 - illustrated 7-12
- Create Shared Folder dialog box,
 - illustrated 10-13
- Customize Permissions dialog box, Security tab, illustrated 10-16
- Customize Permissions dialog box, Share Permissions tab, illustrated 10-15

D

- data
 - archival and recovery A-8
 - block 3-6
 - organization of folders 8-3
 - stripes 3-6
- data backup, caution 5-9, 5-10
- data guarding See RAID methods
- data replication, defined 2-20
- data striping See RAID 0
- Date and Time dialog box, illustrated 2-22
- Default Quota dialog box, illustrated 6-55
- deployments
 - domain 7-2
 - StorageWorks NAS B2000, single device 1-6
 - typical 1-6
 - workgroup 7-2
- Device Information Window,
 - illustrated 11-12
- DHCP (dynamic host configuration protocol) 11-5
- disaster recovery
 - backup solutions A-10
 - steps A-10
- Disk Quota dialog box, illustrated 6-54
- Display/Delete Snapshot Schedule dialog box, illustrated 6-39
- Display/Delete the Snapshot Schedule Wizard screen, illustrated 6-40
- distributed data guarding 3-11. See RAID methods. See also RAID methods
- DNS (domain name system) 11-5

- domain controller
 - domain environment 7-3
 - functions 1-9
- domain environment
 - compared to workgroup 7-2
 - described 1-9, 7-3
 - domain controller 7-3
- drive See hard drives
- drive arrays
 - creating a new array, ACU (Array Configuration Utility) 5-19
 - creating logical drives, ACU (Array Configuration Utility) 5-23
 - expanding capacity, ACU (Array Configuration Utility) 5-26
- drive arrays
 - managing, ACU (Array Configuration Utility) 5-12
- drive bays, caution 5-8
- drive defragmentation, caution 3-20
- drive failure
 - backups 5-8
 - fault tolerance 5-8
 - moving arrays 5-10
- drive letters
 - setting for a snapshot 6-34
 - setting for a virtual disk 6-19
- drive mirroring See RAID methods
- drive quotas See quotas
- drive striping See RAID, 0

E

- Edit NFS Client Groups dialog box,
 - illustrated 9-24
- electric shock
 - warning xvi
- email alerts, setting up 2-26
- embedded smart array controller
 - configuration, caution 5-16, 5-19

enterprise management

- Compaq Insight Manager for HP OpenView 11-14
- Compaq Insight Manager for Tivoli NetView (AIX), Version 2.0 11-16
- HP OpenView 11-14
- remote administration 11-14
- Tivoli NetView (Linux) 11-16

environments

- domain 1-9, 7-1
- domain and workgroup, compared 7-2
- users and groups 7-1
- workgroup 1-8, 7-1

environments, defined 1-8

equipment damage, warning xvi

Ethernet teaming

- adding NICs to a team 2-6
- checking the status of 2-15
- configuration procedures 2-11
- CPQTeam 2-3
- CPQTeam utility
 - installing 2-4
 - opening 2-5
- features 2-3
- procedures, listed 2-3
- renaming the team 2-11
- taskbar icon, displaying 2-12
- TCP/IP protocol configuration 2-13
- teaming options
 - fault tolerance 2-8
 - load balancing 2-9
- troubleshooting 2-16

events, logging, auditing 9-7

expanding

- array, setting priority, ACU (Array Configuration Utility) 5-16
- priority, ACU (Array Configuration Utility) 5-18

explicit mapping

- groups 9-34

explicit mappings

- defined 9-26
- user 9-32

F

failures, catastrophic A-8

Fast EtherChannel See FEC

fault tolerance

- compromised 5-8
- configuring teams, help 2-8
- data loss 5-8
- increasing 2-3
- methods of 3-8
- nondrive problems 5-8
- RAID 0 3-9
- RAID 1 3-9
- RAID 4 3-11
- RAID 5 3-11
- RAID ADG 3-12
- teaming option
 - fail on fault 2-8
 - manual 2-8
 - smart switch 2-8

FEC (Fast EtherChannel)

- defined 2-10
- teaming 2-10

figures

- ACU (Array Configuration Utility)
 - example Array B 5-21

File and Print Services for NetWare, illustrated 10-5

file caching

- CIFS shares 8-29

file sharing See shares

file system security, managing with the Remote Insight Lights-Out Edition board 8-12

Folder Properties dialog box, General tab, illustrated 8-8

folders
 creating 8-6
 creating a share for 8-9
 deleting 8-7
 managing shares for 8-11
 modifying 8-8
 navigating to a specific folder 8-4
 organization of data 8-3
 security of 8-19
 tasks 8-3

Folders dialog box, illustrated 8-5

Format Virtual Disk dialog box,
 illustrated 6-20

formatting, virtual disks 6-20

FTP shares See shares, FTP

G

GID (group ID), authentication 9-5

grounding plug, warning xvi

grounding, caution xvi

group ID See GID

group identification, user name
 mapping 9-25

Group Name Examples, table 7-5

Group Properties dialog box Members tab,
 illustrated 7-19

Group Properties dialog box, General tab,
 illustrated 7-17

groups, local

 creating 7-16

 creating, with wizard 7-8

 deleting 7-17

 deleting, with wizard 7-8

 management 7-15

 member list 7-18

 modifying properties 7-17

 modifying, with wizard 7-8

 name examples 7-4

 naming rules and guidelines 7-4

 tasks 7-15

 using tags 7-5

groups, NFS client See NFS (Network File
 System)

H

Hard Drive LED Combinations, table 5-4

hard drives

 array, limitations 3-7

 caution 5-3

 drive failure

 ADU 5-6

 LED status indicators 5-6

 POST (Power-on Self-test)

 messages 5-6

 recognizing 5-6

 replacing failed drive 5-5

 failure, lost data 5-10

 fault tolerance 5-8

 hot-plug 1-4

 LED indicators 5-3

 LED indicators, illustrated 5-3

 mirrored pairs 3-9

 moving, steps 5-9

 replacing 5-5, 5-7, 5-8

 replacing, caution 5-10

 striped 3-9

hazardous energy levels, warning xvi

help

 additional sources xvii

 Compaq authorized resellers, telephone

 numbers xvii

 Compaq website xvii

 Insight Manager for Tivoli (AIX),

 website 11-17

 technical support, telephone

 numbers xvii

 Windows file system security

 website 8-1

High Voltage Differential See HVD

hot-pluggable drives See hard drives

HP OpenView
 Network Node Manager 11-15
 remote administration 11-10
 Windows 2000 operating system 11-14
HTTP shares See shares, HTTP
HVD (High Voltage Differential)
 defined A-2

I

ID positions, altering 5-9
illustrations
 example Array B, ACU (Array
 Configuration Utility) 5-21
IML (Integrated Management Log),
 described xvii
Important Safety Information document xv
Installing CPQTeam, illustrated 2-4
installing File and Print Services for
 NetWare, illustrated 10-4

L

LAN (Local Area Network)
 access through NIC 11-4
LED indicators
 hard drives 5-3
libraries, tape A-5
lifeguard service
 caution 3-22
 facts 3-23
 overview 3-22
limitations
 drive array 3-7
Linux
 accessing other platforms 9-3
 Compaq Insight Manager for 11-16
 installing Insight Manager for
 Tivoli 11-17
 managing with 11-16

load balancing
 increasing 2-3
 teaming option
 adaptive load balancing 2-10
 with IP address 2-10
 with MAC address 2-10
Local Area Connection Properties page,
 illustrated 10-3
local area connection, show icons 2-12
local area network See LAN
Local Groups dialog box, illustrated 7-15
Local Users dialog box, illustrated 7-10
Log File Information screen,
 illustrated 6-48
logical drive
 advantages of 3-4
 size, defined, ACU (Array Configuration
 Utility) 5-24
logical drive, caution 5-12
logical storage units See LUNs
Logs menu, illustrated 2-24
Low Voltage Differential See LVD
LUNs
 adding to a pool 6-14
 capacity expansion 4-51
 caution 6-24
 defined 3-6
 large 3-7
 management of 5-11
 size 3-7
LVD (Low Voltage Differential)
 defined A-2

M

main configuration screen, ACU (Array
 Configuration Utility) 5-15
Maintenance menu, illustrated 2-21

- maintenance, system
 - audit logs 2-24
 - basic tasks 2-21
 - date and time 2-22
 - email alerts 2-26
 - restart 2-23
 - shutdown 2-23
 - shutdown, scheduled 2-23
 - Terminal Services 2-25
 - Manage Policies dialog box, illustrated 6-12, 6-50
 - Manage Pools dialog box, illustrated 6-7
 - Manage Snapshots dialog box, illustrated 6-29
 - Manage Virtual Disks dialog box, illustrated 6-15
 - management software
 - Compaq Insight Manager for HP OpenView (Windows 2000 operating system), installing 11-17
 - Compaq Insight Manager for Tivoli NetView (AIX) installing 11-17
 - Compaq Insight Manager, installing 11-17
 - installing 11-17
 - Mapping Server ls -al command example, illustrated 9-27
 - mappings
 - backing up and restoring 9-37
 - best practices 9-28
 - creating and managing 9-29
 - creating, with wizard 7-9
 - explicit 9-26
 - explicit group 9-34
 - explicit user 9-32
 - group identification 9-25
 - NIS server 9-31
 - password and group files 9-31
 - server computer name 9-5
 - simple 9-26, 9-31
 - squashing 9-27
 - types 9-26
 - user name 9-25
 - members
 - of local user groups 7-18
 - of NFS client groups 9-21
 - Microsoft Management Console
 - overview 1-14
 - tasks, listed 1-14
 - Microsoft Management Console, illustrated 1-14
 - Microsoft Windows file system security, website 8-1
 - migration
 - RAID level, ACU (Array Configuration Utility) 5-29
 - stripe size, ACU (Array Configuration Utility) 5-29
 - MMC New User dialog box, illustrated 10-7
 - MMC Password Synchronization screen, Advanced Settings dialog box, illustration 9-45
 - MMC Password Synchronization screen, illustrated 9-42
 - MMC Server for NFS screen, Logging tab, illustrated 9-7
 - MMC Server for NFS screen, User Mapping tab, illustrated 9-6
 - MMC User Name Mapping screen, Map Maintenance tab, illustrated 9-37
 - moving arrays
 - array controller 5-10
 - restrictions 5-10
- N**
- Namespace Recovery screen, illustrated 6-52
 - NAS (network attached storage), backup capabilities A-1

- NAS B2000
 - continuous data availability 1-4
 - defined 1-2
 - environments
 - defined 1-8
 - included software 1-3
 - manageability 1-4
 - overview 1-1
 - product information 1-4
 - supported software 1-3
- NCP (Novell Core Protocol)
 - shares See shares, NCP
- NetWare Basic Share Permissions dialog
 - box, illustrated 10-14
- NetWare Services tab, illustrated 10-8
- NetWare shares See shares, NCP
- Network menu, illustrated 2-27
- Network Node Manager See NNM
- network settings 2-27
- New NFS Client Group dialog box,
 - illustrated 9-22
- New Quota Entries dialog box,
 - illustrated 6-57
- New Snapshots Information screen,
 - illustrated 6-32
- New Virtual Disk Step 1 dialog box,
 - illustrated 6-18
- NFS (Network File System)
 - audit logs 9-7
 - authentication
 - client level 9-4
 - per-export basis 9-4
 - authentication, installing software 9-8
 - client groups
 - creating 9-21
 - deleting 9-23
 - managing 9-20
 - members of 9-21
 - modifying 9-24
 - permissions 9-20
 - commands 9-41
 - component functionality 9-4
 - design goals 9-3
 - mapping
 - types of 9-26
 - mappings
 - backing up and restoring 9-37
 - best practices 9-28
 - creating and managing 9-29
 - discussed 9-25
 - explicit group 9-34
 - explicit user 9-32
 - NIS server 9-31
 - password and group files 9-31
 - simple 9-31
 - protocol parameters
 - async/sync 9-16
 - locks 9-17
 - protocol parameters, settings 9-15
 - protocol version combinations 9-3
 - Remote Shell service 9-41
 - restrictions former 9-4
 - SFU abilities 9-1
 - SFU integration 9-4
 - shares See shares, NFS
 - shares, protocol parameters 9-15
 - Telnet Service 9-41
 - versions of 9-3
 - Windows NT access 9-4
 - Windows NT filesharing 9-4
- NFS Async/Sync Settings dialog box,
 - illustrated 9-17
- NFS Client Groups dialog box,
 - illustrated 9-21, 9-23
- NFS Locks dialog box, illustrated 9-19
- NFS Sharing Protocols menu,
 - illustrated 9-15
- NFS Sharing tab, illustrated 9-13
- NIC Properties, Teaming Controls tab, fault tolerant option, illustrated 2-7
- NIC Properties, Teaming Controls tab, Load Balancing option, illustrated 2-9
- NIC Team Properties dialog box,
 - illustrated 2-13
- NIC Team TCP/IP Properties dialog box,
 - illustrated 2-14

NIS server, mapping 9-31
NNM (Network Node Manager),
 defined 11-15
no fault tolerance 3-9
NSPOF Horizontal Array Configuration,
 illustrated 4-11

O

OCU (Offline Configuration Utility),
 defined 11-6
Offline Configuration Utility See OCU
offsite, data storage A-7
online volume growth
 adding storage to pools 6-24
 description 6-23
 preparing disks for 6-23
overheating, caution xvi
ownership of files 8-18

P

parity data
 RAID 5 3-11
 RAID ADG 3-12
password and group files, for NFS
 mapping 9-31
password synchronization 9-42
 advanced settings 9-45
 best practices 9-43
 customizing settings 9-48
 implementing 9-44
 installing 9-46
 requirements 9-44
passwords, local user 7-13
permissions
 ACLs 8-20
 NFS client groups 9-20
personal injury, warning xvi
physical storage, best practices 3-15
pool policies, caution 6-49

Pool Policy Default Settings, table 6-13,
 6-51
Pool Properties summary display 6-10
pools
 adding storage to 6-24
 adding storage units 6-14
 bringing online and offline 6-14
 capacity expansion 4-51
 creating 6-9
 creating, with wizard 6-4
 deleting 6-21
 deleting, with wizard 6-4
 facts 3-17
 large 3-7, 4-51
 management 6-7
 moving arrays 5-10
 online volume growth 6-23
 overview 3-16
 policies 6-11, 6-49
 policies, caution 6-11
 properties 6-10
 tasks 6-8
POST (Power-on Self-test) messages
 drive failure 5-6
 drive positions changed 5-9
 fault tolerance 5-8
 re-enable logical drive 5-8
powering down the server 2-23
Primary WebUI screen, illustrated 1-10
priority settings, ACU (Array Configuration
 Utility) 5-16
properties
 folders 8-8
 group, local 7-17
 NFS shares 9-12
 shares 8-27, 10-11
 user, local 7-13

- protocols
 - parameters 8-34
 - AFP 8-37
 - CIFS 8-36
 - FTP 8-36
 - HTTP 8-37
 - NCP 8-37
 - NFS 8-36, 9-15
 - async/sync 9-16
 - locks 9-17
 - supported 8-1
- Q**
- Quota Entries dialog box, illustrated 6-56
- Quota Entry User dialog box, illustrated 6-59
- quotas
 - creating 6-56
 - deleting 6-58
 - enabling 6-55
 - modifying 6-58
 - space usage, managing 6-53
- R**
- RAID
 - migration, ACU (Array Configuration Utility) 5-29
 - RAID 0 3-9
- RAID methods 3-8
 - RAID 1 3-9
 - RAID 4 3-11
 - RAID 5 3-11
 - RAID ADG 3-12
- rebuild
 - priority setting, ACU (Array Configuration Utility) 5-16
 - priority, ACU (Array Configuration Utility) 5-17
- recovery and archival, data A-8
- remote administration
 - Compaq Insight Manager
 - Agent Web interface 11-13
 - components 11-11
 - console 11-11
 - management agents 11-11
 - enterprise management
 - Compaq Insight Manager for HP OpenView 11-14
 - Compaq Insight Manager for Tivoli NetView (AIX), Version 2 11-16
 - HP OpenView 11-14
 - installing software 11-17
 - Tivoli NetView (Linux) 11-16
 - management software, Compaq Insight Manager for HP OpenView 11-17
 - methods of 11-1
 - Remote Insight Lights-Out Edition board
 - administration tasks 11-4
 - capabilities 11-4
 - configuration 11-6
 - configuration changes 11-6
 - default DNS name 11-7
 - factory default settings 11-6
 - features 11-4
 - installation and users guide 11-4
 - manage alerts feature 11-6
 - manage users feature 11-5
 - management functions 11-4
 - OCU 11-7
 - remote access 11-7
 - Web interface 11-4
 - Remote Shell Daemon 11-10
 - Remote Shell Service 11-10
 - Telnet Server 11-8
 - auditing 11-9
 - authentication 11-9
 - sessions 11-10
 - settings 11-9
 - Terminal Services 9-40
 - WebUI 11-2

- Remote Insight Lights-Out Edition board
 - capabilities 11-4
 - configuration 11-6
- Remote Shell Service
 - remote administration 11-10
 - using 9-41
- repairs, warning xvi
- replicating data 2-20
- restarting the server 2-23
- Restore Virtual Disk screen, illustrated 6-46
- Restore Virtual Disks from Snapshot screen, illustrated 6-45
- restoring, mappings 9-37
- restrictions
 - drive array 3-7

S

- safeguarding data A-4
- safekeeping data, offsite A-7
- safety information xv
- SAN (storage area network)
 - backup solutions A-2
 - defined A-2
- Schedule a Snapshot screen, illustrated 6-25
- Schedule Information screen, illustrated 6-27, 6-37
- Schedule Snapshot Deletion Task dialog box, illustrated 6-41
- SCSI tape devices
 - HVD A-2
 - LVD A-2
- security
 - ACLs 8-20
 - file level 8-19
 - file shares 8-19
 - integrating
 - described 8-20
 - local file system into domain environments 8-20
 - managing 8-12
 - permissions 8-20
 - RILOE, features 11-5

- share permissions and file-level permissions, integration 8-21
- Security Properties dialog box
 - for folder name NTSF Test, illustrated 8-13
 - options 8-13
- Select User, Computer, or Group dialog box, illustrated 8-17
- Services for UNIX See SFU
- Set Drive Letter dialog box, illustrated 6-34
- setting expand or rebuild priority, ACU (Array Configuration Utility) 5-16
- SFU (Services for UNIX)
 - defined 9-1
 - for NFS 9-4
 - NFS audit logs 9-7
- share management wizard 8-22
- Share Management Wizard, illustrated 8-23
- Share Properties dialog box, CIFS Sharing tab, illustrated 8-30
- Share Properties dialog box, General tab, illustrated 8-28, 9-12, 10-11
- Share Properties dialog box, NetWare Sharing tab, illustrated 8-33, 10-12
- Share Properties dialog box, NFS Sharing tab, illustrated 8-31
- shares
 - ACLs 8-20
 - administrative 8-21
 - AFP shares
 - creating 8-26, 10-10
 - creating, with wizard 8-24, 8-25
 - deleting 8-27, 10-11
 - modifying 8-33
 - modifying, with wizard 8-25
 - protocol parameters 8-37
 - AppleTalk shares See shares, AFP

- CIFS shares
 - administrative 8-21
 - creating 8-26, 10-10
 - creating, with wizard 8-24
 - deleting 8-27, 10-11
 - deleting, with wizard 8-25
 - file caching 8-29
 - modifying 8-29
 - modifying, with wizard 8-25
 - protocol parameters 8-36
 - standard 8-21
- creating 8-26, 10-10
- creating, for a folder 8-9
- creating, with wizard 8-24
- deleting 8-27, 10-11
- deleting, with wizard 8-25
- file caching 8-29
- FTP shares
 - creating 8-26, 10-10
 - creating, with wizard 8-24
 - deleting 8-27, 10-11
 - deleting, with wizard 8-25
 - modifying 8-32
 - modifying, with wizard 8-25
 - protocol parameters 8-36
- HTTP shares
 - creating 8-26, 10-10
 - creating, with wizard 8-24
 - deleting 8-27, 10-11
 - deleting, with wizard 8-25
 - modifying 8-32
 - modifying, with wizard 8-25
 - protocol parameters 8-37
- management methods 8-19
- managing 8-19
- managing with the wizard 8-22
- modifying 8-27, 10-11
- modifying, with wizard 8-25
- NCP shares
 - creating 8-26, 10-10
 - creating, with wizard 8-24
 - deleting 8-27, 10-11
 - deleting, with wizard 8-25
 - modifying 8-32
 - modifying, with wizard 8-25
 - protocol parameters 8-37
- NetWare shares See shares, NCP
- NFS shares
 - creating 8-26, 9-10, 10-10
 - creating, with wizard 8-24
 - deleting 8-27, 9-11, 10-11
 - deleting, with wizard 8-25
 - modifying 8-31, 9-12
 - modifying, with wizard 8-25
 - protocol parameters 8-36, 9-15
 - testing 9-39
- overview 2-19, 8-19
- protocol parameters 8-34
- standard 8-21
- tasks, listed 8-25, 10-9, 10-13
- types of 8-28
- Web shares See shares, HTTP
- sharing See shares
- Sharing Protocols dialog box,
 - illustrated 8-35
- Shutdown menu, illustrated 2-23
- shutting down the server 2-23
- simple mappings
 - creating 9-31
 - defined 9-26
- smart array controller configuration,
 - caution 5-13
- snapback, defined 6-43
- Snapshot and Task Information screen,
 - illustrated 6-36
- Snapshot Properties summary display,
 - illustrated 6-33
- Snapshot Wizard Welcome screen,
 - illustrated 6-35
- snapshots
 - access 3-19
 - backup 3-21
 - benefits 3-18
 - caution 3-21
 - creating 6-30
 - creating, with wizard 6-5

- defragmenting 3-20
 - deleting 6-33
 - effects on virtual disk 3-18
 - facts 3-19
 - initial implementation 3-19
 - management 6-29
 - naming of 3-21
 - overview 3-18
 - pool sizing 3-19
 - properties, viewing 6-33
 - read performance 3-18
 - recovery
 - caution 6-43
 - restoring a virtual disk 6-43
 - restoring files and directories 3-22
 - restoring from A-8
 - schedule for a virtual disk 6-35
 - scheduled snapshots 6-25
 - schedules
 - displaying and deleting 6-39
 - setting the drive letter of 6-34
 - Snapshot Planner utility 3-19
 - snapshots of 3-19
 - tasks 6-30
 - virtual disks 3-20
 - Windows 2000 defragmenter 3-20
 - write locality 3-18
 - write performance 3-18
- software
- included 1-3
 - supported 1-3
- squashing mappings, defined 9-27
- storage
- offsite storage A-7
 - overview
 - best practices, virtual 3-24
 - virtual 3-16
 - physical 3-4
 - arrays 3-5
 - best practices 3-15
 - hard drives 3-4
 - LUNs 3-6
 - pools See pools
 - quotas See quotas
 - snapshots See snapshots
 - virtual, overview 3-16
 - storage area network See SAN
 - storage enclosure, power down
 - precaution 5-6
 - storage management wizard 6-3
 - Storage Management Wizard screen,
 - illustrated 6-3
 - storage management, Smart Array Controller
 - described 5-12
 - StorageWorks NAS B2000, single-device
 - deployment 1-6
 - stripe size 3-6
 - fault tolerance, ACU 5-24
 - migration, ACU (Array Configuration Utility) 5-29
 - symbols, in text xv
 - system date and time, setting 2-22
 - system setup, completing 2-2

T

- tables
- Command Line Interface Command Prompts 9-41
 - Group Name Examples 7-5
 - Hard Drive LED Combinations 5-4
 - Optimum Stripe Sizes for Different Environments 5-24
 - Pool Policy Default Settings 6-13, 6-51
 - Storage Enclosure Drive Bay Configuration for Replacing Failed Hard Drives 5-7
- tags, for managing user and group names 7-5
- tape devices, backup solutions A-2, A-3
- tape libraries, automated A-5
- TCP (Transport Control Protocol), with NFS 9-3

- TCP/IP protocol, configuring the Ethernet team 2-13
 - technical support, telephone numbers xvii
 - technician notes, warning xvi
 - telephone numbers
 - Compaq authorized resellers xvii
 - technical support xvii
 - Telnet Server
 - configuring 11-8
 - defined 11-8
 - interface screen, illustrated 11-9
 - remote administration 11-8
 - auditing 11-9
 - authentication 11-9
 - sessions 11-10
 - settings 11-9
 - using 9-41
 - Terminal Services
 - remote administration 9-40, 11-3
 - using 2-25
 - Terminal Services Session, illustrated 2-25
 - Tivoli NetView
 - (AIX) remote administration 11-16
 - remote administration 11-10
 - Transport Control Protocol See TCP
 - troubleshooting, NIC teaming 2-16
 - Tru64 UNIX, accessing from Linux 9-3
- U**
- UDP (User Datagram Protocol), with NFS 9-3
 - UID (user ID)
 - NFS authentication 9-5
 - Updated CPQTeam Properties dialog box, illustrated 2-15
 - User and Group Mappings dialog box
 - Explicit Group Mapping tab, illustrated 9-35
 - Explicit User Mapping tab, illustrated 9-33
 - illustrated 9-30
 - Simple Mapping tab, illustrated 9-32
 - User Datagram Protocol See UDP
 - user ID See UID
 - User Information and Schedule Information screen, illustrated 6-42
 - user interfaces, listed 1-9
 - user management wizard 7-6
 - User Management Wizard, illustrated 7-7
 - user name mapping See mappings
 - user names
 - management 7-4
 - rules and guidelines 7-4
 - User or Group Permission Entry dialog box, for folder name NTSF Test, illustrated 8-15
 - User Properties dialog box, illustrated 7-14
 - users and groups
 - planning 7-3
 - users, local
 - creating 7-11
 - creating, with wizard 7-7
 - deleting 7-12
 - deleting, with wizard 7-7
 - modifying 7-13
 - modifying password 7-13
 - modifying, with wizard 7-7
 - naming rules and guidelines 7-4
 - setting password, with wizard 7-7
 - tasks 7-10
 - utilities
 - backup solutions
 - developing A-2
 - disaster recovery A-10
 - media rotation A-7
 - tape devices A-2
 - through SAN A-2
 - CPQTeam 2-3

V

- ventilation clearances, caution xvi
- Virtual Disk and Snapshots screen, illustrated 6-31
- Virtual Disk Backup Summary screen, illustrated 6-28
- Virtual Disk Properties screen, illustrated 6-22
- virtual disks
 - backing up 6-25
 - caution 6-16
 - creating 6-17
 - creating a scheduled snapshot of 6-25
 - creating a snapshot schedule for 6-35
 - creating, with wizard 6-4
 - defragmenting 3-20
 - deleting 6-21
 - deleting, with wizard 6-5
 - facts 3-18
 - formatting 6-20
 - growing 6-23, 6-24
 - incremental backup 6-47
 - overview 3-17
 - properties 6-22
 - restoring from a snapshot 6-43
 - setting the drive letter of 6-19
 - tasks 6-16
- Virtual Replicator See VR
- virtual storage, best practices 3-24
- Volumes dialog box, illustrated 8-4
- VR (virtual replicator)
 - adding storage units 6-14
 - best practices 3-24
 - described 3-16
 - lifeguard service
 - overview 3-22
 - policies 6-11, 6-49
 - pools See pools
 - snapshots See snapshots

- storage tools
 - listed 3-16
 - online volume growth 6-23
 - pools See pools
 - snapshots See snapshots
 - virtual disks See virtual disks
- virtual disks See virtual disks

W

- warnings
 - component-level repairs xvi
 - defined xv
 - electric shock xvi
 - equipment damage xvi
 - grounding plug xvi
 - hazardous energy levels xvi
 - improper repairs xvi
 - personal injury xvi
 - technician notes xvi
- warranty, voiding xvi
- Web browser, management 1-10
- Web Enabled interface, illustrated 11-15
- Web shares See shares, HTTP
- Web-based user interface See WebUI
- websites
 - ActiveAnswers 1-9, 7-3
 - Windows file system security 8-1
- websites, Compaq xvii
- WebUI
 - folder management 8-4
 - maintenance tasks 2-21
 - NFS file sharing management 9-10
 - overview 1-10
 - remote administration 11-2
 - share management 8-22
 - storage management 6-3
 - tasks, listed 1-11
 - users and group management 7-5

- Windows file system security
 - website 8-1
- wizards
 - share management 8-22
 - storage management 6-3
 - user management 7-6
- wizards, ACU (Array Configuration Utility) 5-12
- workgroup environment
 - compared to domain 7-2
 - described 1-8
- workgroups
 - deployment 7-2
 - group management 7-15
 - group, properties 7-17
 - groups, deleting 7-17
 - user and group administration 7-5
 - user management 7-11
 - user management, tasks 7-10
 - user, deleting 7-12
 - user, management 7-10