

administration guide

hp StorageWorks NAS b2000

Fourth Edition (February 2003)

Part Number: 292278-004

This guide provides information on performing the administrative tasks necessary to manage the HP StorageWorks NAS b2000 server. Overview information as well as procedural instructions are included in this guide.



© Hewlett-Packard Company, 2003.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Compaq Computer Corporation is a wholly-owned subsidiary of Hewlett-Packard Company.

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

Intel and Pentium are trademarks of Intel 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.

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

NAS b2000 Administration Guide
Fourth Edition (February 2003)
Part Number: 292278-004



- About this Guide. 15**
- Overview. 16
 - Intended Audience 16
 - Prerequisites 16
- Conventions 17
 - Document Conventions 17
 - Text Symbols 17
 - Equipment Symbols 17
- Rack Stability 19
- Getting Help 19
 - HP Technical Support 19
 - HP Storage Website 20
 - HP Authorized Reseller 20
- 1 System Overview 21**
- Product Definition and Information 22
 - Server Hardware Features 22
 - Optional Features. 23
 - Software Features. 23
 - Product Information 23
 - Product Manageability. 24
 - Product Redundancy 24
 - Product Scalability. 25
- Deployment Scenarios 26
- Environment Scenarios. 27
 - Workgroup 27
 - Domain. 27
- User Interfaces 28
 - NAS b2000 Web-Based User Interface 28
 - Status 28
 - Network 29
 - Disks 29
 - Users 29
 - Shares. 29
 - Maintenance. 29
 - HP Utilities 29
 - Help 29

Take a Tour	29
Rapid Startup Wizard	29
Set Server Appliance Name.	29
Set Administrator Password	29
Set Default Page	29
NAS b2000 Desktop	30
NAS Management Console	31
CPQTeam Setup	31
Install Data Copy	31
2 Setup Completion and Basic Administrative Procedures	33
Setup Completion	34
Setting up Ethernet NIC Teams (Optional)	34
Installing the CPQTeam Utility.	34
Opening the CPQTeam Utility	36
Adding and Configuring NICs in a Team	36
Fault Tolerance	37
Load Balancing	38
Configuring the NIC Team Properties.	39
Renaming the Teamed Connection.	39
Showing a Connection Icon on the Taskbar	40
Configuring the TCP/IP Protocol on the New Team	40
Checking the Status of the Team.	41
NIC Teaming Troubleshooting	43
Managing System Storage	43
Creating and Managing Users and Groups	44
Creating and Managing File Shares	44
Installing and Configuring Data Replication Software	45
Activating the iLO Port Using the License Key	46
Basic Administrative Procedures	47
Setting the System Date and Time.	48
Shutting Down or Restarting the Server	49
Viewing and Maintaining Audit Logs	50
Using Terminal Services	51
Setting up E-mail Alerts.	51
Updating the Software	52
Changing System Network Settings.	52
3 Storage Management Overview	53
Storage Management Process.	53
Storage Elements Overview	55
Physical Hard Drives	55
Arrays.	56
Logical Drives (LUNs)	57
Fault-Tolerance Methods	58
RAID 0—Data Striping	58
Advantages	58
Disadvantages	59
RAID 1—Drive Mirroring.	59
Advantages	59

Disadvantages	60
RAID 5—Distributed Data Guarding	60
Advantages	61
Disadvantages	61
RAID ADG—Advanced Data Guarding	61
Advantages	62
Disadvantage	62
Physical Storage Best Practices	63
Logical Storage Elements Overview	63
Partitions	63
Volumes	64
Utilizing LDM Storage Elements	64
Persistent Storage Management Elements Overview	65
File System Elements	66
File-Sharing Elements	66
4 Advanced Storage Management Planning	67
Fundamental Storage Configuration Planning Issues	68
System Priorities	68
Array Configuration (Striping) Methods	69
Vertical Array Configurations	69
Horizontal Array Configurations	72
NSPOF Horizontal Array Configurations	74
Recommended System Configurations	75
Storage Enclosure Configuration Options	75
Recommended Configuration Methods	76
When Fault Tolerance is Most Important	76
When Capacity Utilization Is Most Important	77
When I/O Performance Is Most Important	77
Physical Storage Planning Issues	78
Hard Drive Sizes and Types	78
Mixed Drive Sizes	78
Mixed Drive Types	78
Spare Drive Sizes	78
Use and Number of Spare Disks	78
LUN Sizing	79
LUNs Cannot be Extended	79
LUN Management under Windows Powered OS	80
Storage Sizing Considerations	80
RAID Issues	80
Spare Disk Issues	81
Snapshot Issues	81
Growth Issues	81
Allocation Unit Size Issues	82
Consolidation Issues	82
Storage Management Planning Scenarios	83
A Complete and Detailed Storage Planning Example	83
Initial Storage Needs	83
Snapshot Storage Needs	83
Total Storage Need	84

Array Configuration Requirements	85
Drives Required	87
Storage Enclosures Required	88
Conclusion	89
A Simple Sizing Comparison	89
An Example of a Storage Subsystem Using Different Array Configurations	90
Planning Worksheet	91
Migration Issues	94
Developing a Migration Plan	94
System Wide Migration	94
Departmental Migration	94
Performing the Migration	95
Backup and Restore	95
Ethernet Copy	96
Drive Migration	96
Storage Capacity Expansion Issues	97
5 Physical Storage Management	99
Hard Drive Management	100
Defining Hard Drive LED Indicators	100
Replacing Failed Hard Drives	102
Compromised Fault Tolerance	104
Moving Hard Drives	104
Moving Arrays	105
Array and LUN Management	106
ACU Overview	106
Features of the ACU	106
Accessing the ACU	107
Entering Controller Settings	109
Creating a New Array	111
Creating Logical Drives (LUNs)	114
Expanding the Capacity of an Existing Array	116
Migrating an Existing LUN to a New RAID Level or Stripe Size	119
6 Persistent Storage Manager	121
Operational Overview	121
Reading Snapshots	122
Creating Snapshots	122
PSM Snapshot Attributes	122
Read Only	122
Read/Write	122
Always Keep	123
Automated Snapshot Deletion	123
Data Recovery	123
File/Folder/Volume Recovery	123
Snapshots and Drive Defragmentation	123
PSM and Backup	124
Snapshots Performance Impact	125
Recovering Snapshots after a System Restore or System Loss	125
Granule Size Update Utility	126

Clearing the Cache File from the System	127
Re-extending Volumes from Old Snapshots	127
Volume Display in Persistent Storage Manager	127
Persistent Storage Manager Storage Limitations	128
Accessing Persistent Storage Manager	128
Global Settings	129
Maximum Persistent Images	129
Inactive Period	129
Inactive time-out	130
Image directory	130
Restore Defaults	130
Volume Settings	130
Available Volume	130
Size	131
Free Space	131
Cache Size	131
Usage	131
Volume Configuration Settings	131
Warning threshold reached when	132
Begin deleting images when	132
Cache size	132
Schedules	132
Create a New Schedule	133
Editing Persistent Image Schedule Properties	134
Deleting a Persistent Image Schedule	135
Persistent Image and Group Information	136
Image name and location on volume	136
Persistent image group name	136
Number of images in group	136
Volumes included in this image	136
Image attributes	136
Retention weight	137
Most recent image in group	137
Oldest image in group	137
Next image in group to be deleted	137
Managing Persistent Images	137
Creating a New Persistent Image	138
Deleting a Persistent Image	139
Editing Persistent Image Properties	140
Undo Persistent Image Changes	141
Restoring an Image	142
Known Issues	144
Event log error at cache full	144
Display Error on SAK	144
Always Keep error at cache file full	144
Improper display of default Cache File Size	144
Page file setting	144
No Boot - No Revert	144
Reverting of System Drive Prohibited	144
No support for mount points in UNIX, AppleTalk, or NetWare	144

7	User and Group Management	145
	Domain Compared to Workgroup Environments	146
	User and Group Name Planning	146
	Managing User Names	147
	Managing Group Names	147
	Workgroup User and Group Management	148
	Managing Local Users	148
	Adding a New User	149
	Deleting a User	149
	Modifying a User Password	150
	Modifying User Properties	150
	Managing Local Groups	151
	Adding a New Group	152
	Deleting a Group	152
	Modifying Group Properties	153
	General Tab	153
	Members Tab	153
	Drive Quotas	154
	Managing Quotas	154
	Enabling and Disabling Quota Management	156
	Creating New Quota Entries for a User or Group	157
	Deleting Quota Entries for a User or Group	158
	Modifying Quota Entries for a User or Group	158
8	Folder and Share Management	161
	Folder Management	162
	Navigating to a Specific Volume or Folder	163
	Creating a New Folder	164
	Deleting a Folder	165
	Modifying Folder Properties	165
	Creating a New Share for a Volume or Folder	166
	Managing Shares for a Volume or Folder	167
	Managing File Level Permissions	168
	Share Management	174
	Share Considerations	174
	Defining Access Control Lists	174
	Integrating Local File System Security into Windows Domain Environments	175
	Comparing Administrative (Hidden) and Standard Shares	175
	Planning for Compatibility between File Sharing Protocols	175
	NFS Compatibility Issues	175
	Managing Shares	176
	Creating a New Share	176
	Deleting a Share	177
	Modifying Share Properties	178
	CIFS Sharing	178
	NFS Sharing	180
	FTP Sharing	180
	Web Sharing (HTTP)	180
	NetWare Sharing (NCP)	181
	AFP (AppleTalk) Sharing	182

Installing Services for AppleTalk	182
Installing Windows NT Services for Macintosh	182
Protocol Parameter Settings	183
CIFS Protocol Settings	184
NFS Protocol Settings	184
FTP Protocol Settings	184
HTTP Protocol Settings	184
NCP (NetWare) Protocol Settings	184
AFP (AppleTalk) Protocol Settings	185
9 UNIX File System Management	187
Network File System	188
Server for NFS	188
Authenticating User Access	189
Indicating the Computer to Use for the NFS User Mapping Server	189
Logging Events	190
Installing NFS Authentication Software on the Domain Controllers and Active Directory Domain Controllers	191
NFS File Shares	192
Creating a New Share	192
Deleting a Share	193
Modifying Share Properties	193
Encoding Types	195
NFS Protocol Properties Settings	196
NFS Async/Sync Settings	197
NFS Locks	197
NFS Client Groups	199
Adding a New Client Group	200
Deleting a Client Group	200
Editing Client Group Information	201
NFS User and Group Mappings	202
Types of Mappings	202
Explicit Mappings	202
Simple Mappings	202
Squashed Mappings	203
User Name Mapping Best Practices	203
Creating and Managing User and Group Mappings	204
General Information	204
Simple Mapping	205
Explicit User Mapping	206
Explicit Group Mapping	207
Backing up and Restoring Mappings	208
Backing up User Mappings	209
Restoring User Mappings	209
NFS File Sharing Tests	210
Terminal Services, Telnet Service, and Remote Shell Service	211
Using Terminal Services	211
Using Telnet Server	211
Using Remote Shell Service	211
Password Synchronization	212

Password Synchronization Best Practices	212
Password Synchronization Requirements	213
Implementing Password Synchronization	213
Configuring Advanced Settings	213
Installing Password Synchronization on Domain Controllers and Active Directory Domain Controllers	214
Customizing Password Synchronization	215
10 NetWare File System Management	217
Installing Services for NetWare	218
Managing File and Print Services for NetWare	219
Creating and Managing NetWare Users	220
Adding Local NetWare Users	220
Enabling Local NetWare User Accounts	221
Managing NCP Volumes (Shares)	222
Creating and Managing NCP File Shares Using the WebUI	222
Creating a New NCP Share	222
Deleting an NCP Share	223
Modifying NCP Share Properties	224
Creating and Managing NCP Shares using the NAS Management Console	225
Creating a New NCP Share using the NAS Management Console	226
Modifying NCP Share Properties using the NAS Management Console	229
NOTES:	230
11 Remote Access Methods and Monitoring	231
Web Based User Interface	232
Terminal Services	232
Integrated Lights-Out Port	232
Features	233
Security Features	233
Manage Users Feature	233
Manage Alerts Feature	234
Integrated Lights-Out Port Configuration	234
Using the Integrated Lights-Out Port to Access the NAS b2000	234
Telnet Server	235
Enabling Telnet Server	235
Configuring Telnet Server	235
Authentication Information	235
Auditing Information	236
Server Settings	236
Sessions Information	236
Remote Shell Daemon	236
Insight Manager	237
Insight Manager Console	237
Insight Manager Agent Web Interface	237
Enterprise Management Applications	238
HP OpenView (Windows-Based Operating System)	238
Insight Manager for HP OpenView, Version 2.0	238
Tivoli NetView (AIX)	239
Insight Manager for Tivoli NetView (AIX), Version 2.0	239
Installing the Management Software on the Client Machine	239

Insight Manager for HP OpenView (Windows 2000 Operating System)	239
Insight Manager for Tivoli NetView (AIX)	240
A Backup Management	241
Backup Solutions	241
System Environments	241
SCSI Direct Connect Environments	241
Hardware Options	242
Software Options	242
Best Practices	243
Regular and Reliable Backups	243
Automated Tape Libraries	243
Multiple Backup Devices	243
Backup Schedules	244
Media Rotation	244
Offsite Storage	244
Server Setup Information Archival	245
Snapshots and Quick Online Restores	245
Readiness Testing	246
Disaster Recovery	246
B PSM Error Codes	247
Index	255
Figures	
1 Primary WebUI screen	28
2 NAS b2000 desktop	30
3 Installing Network Teaming	35
4 Network Teaming installation complete	35
5 CPQTeam utility icon	36
6 CPQTeam Properties dialog box	36
7 NIC Properties, Teaming Controls tab, Fault Tolerant option	37
8 NIC Properties, Teaming Controls tab, Load Balancing option	38
9 CPQTeam dialog box	39
10 NIC Team Properties dialog box	40
11 NIC Team TCP/IP Properties dialog box	41
12 Updated CPQTeam Properties dialog box	42
13 NAS data copy install wizard	46
14 Maintenance menu	47
15 Date and Time dialog box	48
16 Shutdown menu	49
17 Logs menu	50
18 Terminal Services session	51
19 Network menu	52
20 Storage Management process	54
21 Separate physical drive (P1, P2, P3) read/write (R/W) operations	55
22 Configuring the physical drives into an array dramatically improves read/write efficiency	56
23 RAID 0 (data striping) (S1-S4) of data blocks (B1-B12)	56
24 2 arrays (A1, A2) and 5 logical drives (L1 through L5) spread over 5 physical drives	57

25	RAID 1 (drive mirroring) of P1 onto P2	59
26	RAID 5 (distributed data guarding) showing parity information (P)	60
27	RAID ADG (advanced data guarding) with two sets of parity data	61
28	System characteristics	68
29	Vertical Array configurations	70
30	Horizontal Array configurations	72
31	NSPOF Horizontal Array configuration, RAID 1+0	74
32	Recommended Configuration methods	76
33	Hot plug hard drive LED indicators	100
34	ACU Logical Drive view	108
35	ACU Physical view	108
36	Controller Settings dialog box	110
37	ACU main configuration screen	111
38	Create Drive Array screen	112
39	Example Array B	113
40	Example Array Logical Configuration view with two arrays	113
41	Create Logical Drive dialog box	114
42	Example array - Configuration View screen with two arrays	116
43	Array expansion example - Logical Configuration View screen	117
44	Expansion wizards - Logical Drive screen	118
45	Migrate RAID/Stripe Size screen	119
46	PSM Main screen	128
47	Global settings	129
48	Volume settings	130
49	Volume configuration settings	131
50	Persistent Image Schedules	132
51	Create Persistent Image Schedule	133
52	Edit schedule properties	134
53	Delete scheduled images	135
54	Persistent Image and Group Information	136
55	Managing persistent images	137
56	Create new persistent image	138
57	Delete Verification	139
58	Edit Persistent Image Properties	140
59	Undo Image Changes	141
60	Images available to restore	142
61	Restore confirmation screen	143
62	Local Users dialog box	148
63	Create New User dialog box	149
64	User Properties dialog box	150
65	Local Groups dialog box	151
66	Create New Group dialog box, General tab	152
67	Group Properties dialog box, General tab	153
68	Group Properties dialog box, Members tab	154
69	Disk Quota dialog box	155
70	Default Quota Dialog box	156
71	Quota Entries dialog box	157
72	New Quota Entry dialog box	158
73	Quota Entry dialog box for a user	159
74	Volumes dialog box	163

75	Folders dialog box	164
76	Create a New Folder dialog box, General tab	165
77	Folder Properties dialog box, General tab	166
78	Create New Share dialog box, General tab	167
79	Security Properties dialog box for folder name NTSF Test	169
80	Access Control Settings dialog box for folder name NTSF Test, Permissions tab	170
81	User or Group Permission Entry dialog box for folder name NTSF Test	170
82	Access Control Settings, Auditing tab dialog box for folder name NTSF Test	171
83	Select User, Computer, or Group dialog box	172
84	Auditing Entry dialog box for folder name NTSF Test	172
85	Access Control Settings, Owner tab dialog box for folder name NTSF Test	173
86	Create a New Share dialog box, General tab	177
87	Share Properties dialog box, General tab	178
88	Share Properties dialog box, CIFS Sharing tab	179
89	Share Properties dialog box, NFS Sharing tab	180
90	Share Properties dialog box, NetWare Sharing tab	181
91	Local Area Connection Properties page, Install option	182
92	Sharing Protocols dialog box	183
93	NAS Management Console Server for NFS screen, User Mapping tab	190
94	NAS Management Console Server for NFS screen, Logging tab	190
95	Create a New Share dialog box, General tab	192
96	Share Properties dialog box, General tab	194
97	NFS Sharing tab	194
98	NFS Sharing Protocols menu	196
99	NFS Async/Sync Settings dialog box	197
100	NFS Locks dialog box	198
101	NFS Client Groups dialog box	199
102	New NFS Client Group dialog box	200
103	Client Groups dialog box	201
104	Edit NFS Client Groups dialog box	201
105	Mapping Server "ls -al" Command example	203
106	User and Group Mappings dialog box, General tab	205
107	User and Group Mappings dialog box, Simple Mapping tab	206
108	User and Group Mappings dialog box, Explicit User Mapping tab	207
109	User and Group Mappings dialog box, Explicit Group Mapping tab	208
110	NAS Management Console User Name Mapping screen, Map Maintenance tab	209
111	Password Synchronization screen	212
112	Password Synchronization screen, Advanced Settings dialog box	214
113	Local Area Connection Properties page, Install option	218
114	Installing File and Print Services for NetWare	219
115	File and Print Services for NetWare screen	219
116	New User dialog box	220
117	NetWare Services tab	221
118	Create a New Share dialog box, General tab	223
119	Share Properties dialog box, General tab	224
120	Share Properties dialog box, NetWare Sharing tab	225
121	Create Shared Folder dialog box	226
122	NetWare Basic Share Permissions dialog box	227
123	Customize Permissions dialog box, Share Permissions tab	228
124	Customize Permissions dialog box, Security tab	229

125 Telnet Server interface screen 235
 126 Web Enabled interface. 238

Tables

1 Document Conventions 17
 2 NIC Teaming Troubleshooting 43
 3 Summary of RAID Methods 62
 4 Vertical Carving Disk Use per RAID Level 70
 5 Horizontal Configuration Disk Use per RAID Level 72
 6 Suggested Storage Enclosure Configurations 75
 7 Example Storage Need Worksheet. 84
 8 Example Array Configuration Requirements Worksheet 87
 9 Example Drives Required Worksheet 88
 10 Example Enclosures Required Worksheet 89
 11 Example Usable Space Using Different Configurations 90
 12 Usable Storage Need Worksheet 91
 13 Array Configuration Storage Needs Worksheet 92
 14 Drive and Enclosure Requirements Worksheet 93
 15 Hard Drive LED Combinations 101
 16 Storage Enclosure Drive Bay Configuration 103
 17 Optimum Stripe Sizes for Different Environments 115
 18 Adjusting Granule Size 126
 19 Group Name Examples 147
 20 Command Line Interface Command Prompts 211
 21 PSM Error Codes 247

about this guide

This administration guide provides information to help you:

- Plan your storage configuration
- Setup physical storage
- Setup virtual storage
- Manage users and groups
- Manage folders and shares
- Manage a UNIX file system
- Manage a NetWare file system
- Remotely access the NAS b2000 server

About this Guide topics include:

- [Overview](#), page 16
- [Conventions](#), page 17
- [Rack Stability](#), page 19
- [Getting Help](#), page 19

Overview

This section covers the following topics:

- [Intended Audience](#)
- [Prerequisites](#)
- [Conventions](#)

Intended Audience

This book is intended for use by system administrators who are experienced with setting up and managing a network server.

Prerequisites

Before beginning, make sure you consider the items below.

- Knowledge of Microsoft Windows NT or 2000 operating systems.
- Knowledge of HP hardware.
- Location of all documentation shipped with your device.

Conventions

Conventions consist of the following:

- [Document Conventions](#)
- [Text Symbols](#)
- [Equipment Symbols](#)

Document Conventions

The document conventions included in [Table 1](#) apply in most cases.

Table 1: Document Conventions

Element	Convention
Cross-reference links	Figure 1
Key and field names, menu items, buttons, and dialog box titles	Bold
File names, application names, and text emphasis	<i>Italics</i>
User input, command and directory names, and system responses (output and messages)	Monospace font COMMAND NAMES are uppercase monospace font unless they are case sensitive
Variables	<i><monospace, italic font></i>
Website addresses	Underlined sans serif font text: http://www.hp.com

Text Symbols

The following 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 death.



Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

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

Equipment Symbols

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings.



Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of personal injury from electrical shock hazards, do not open this enclosure.



Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.



Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of personal injury from a hot component, allow the surface to cool before touching.



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of personal injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Rack Stability

Rack stability protects personal and equipment.



WARNING: To reduce the risk of personal injury or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
 - The full weight of the rack rests on the leveling jacks.
 - In single rack installations, the stabilizing feet are attached to the rack.
 - In multiple rack installations, the racks are coupled.
 - Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.
-

Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our website: <http://www.hp.com>.

HP Technical Support

In North America, call technical support at 1-800-652-6672, available 24 hours a day, 7 days a week.

Note: For continuous quality improvement, calls may be recorded or monitored.

Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the HP website under support: <http://www.hp.com>.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP Storage Website

The HP website has the latest information on this product, as well as the latest drivers. Access storage at: <http://www.hp.com>. From this website, select the appropriate product or solution.

HP Authorized Reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the HP website for locations and telephone numbers: <http://www.hp.com>.

System Overview

1

The HP StorageWorks NAS b2000 server 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 Desktop

Note: The NAS b2000 Desktop can be accessed via a directly connected keyboard and mouse, through Terminal Services, or by using an integrated Lights-Out port.

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 Model 1 server:

- Intel Pentium Xeon 2.80 GHz processor, with 512-KB L2 ECC cache
- 1-GB 200 MHz PC2100 DDR SDRAM memory
- 64-bit PCI-X I/O technology—two hot-pluggable 100MHz slots, one 133MHz slot
- Two 36GB 10K U320 Universal non hot-pluggable hard drives
- Three 146GB 15K U320 Universal hot-plug hard drives for data
- Two embedded 10/100/1000 WOL (Wake on LAN) network interface controllers (NICs)
- Smart Array 5i Plus controller and Battery Backed Write Cache Enabler Module
- IDE DVD-ROM drive
- Redundant hot-plug power supplies and fans
- Two open PCI hot-plug controller slots
- One non hot-plug PCI controller slot
- Embedded iLO board with license key

The following features are included in the NAS b2000 Model 2 server:

- Intel Pentium Xeon 2.80 GHz processor, with 512-KB L2 ECC cache
- 1-GB 200 MHz PC2100 DDR SDRAM memory
- 64-bit PCI-X I/O technology:
 - Two hot-pluggable 100MHz slots
 - One 133MHz slot with Smart Array 5304/256 controller installed
- Two 36.4-GB 10K U320 Universal non hot-pluggable hard drives
- External storage:
 - Four 146GB 15K U320 Universal hot-plug hard drives for data
 - 4314 Storage Cabinet with redundant power supplies and fans
- Two embedded 10/100/1000 WOL (Wake on LAN) network interface controllers (NICs)
- Smart Array 5i Plus controller and Battery Backed Write Cache Enabler Module
- IDE DVD-ROM drive
- Redundant hot-plug power supplies and fans
- Two open PCI hot-plug controller slots
- One non hot-plug PCI controller slot
- Embedded iLO board with license key

Optional Features

The following features are optional for the NAS b2000 server:

- Additional memory
- Smart Array 5300 controller
- StorageWorks 4300 Family storage enclosures
- Network interface cards (NICs)
- Tape drive
- 36.4-GB, 72.8-GB and 146-GB hard drives
- Processor
- SAN Fibre Channel Adapter for tape backup

Software Features

Advanced features included and supported by the NAS b2000 include:

- Array Configuration Utility (ACU)
- 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+0, 5 and ADG
- StorageWorks Data Copy (Trial Version)
- Windows Powered OS
- Columbia Data Products Persistent Storage Manager
- Optional third party supported software (not included):
 - Backup software
 - Management software
 - Quota management
 - Virus protection

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

<http://h18000.www1.hp.com/products/storageworks/nas/supportedsoftware.html>

Product Information

The NAS b2000 provides 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.
- Insight Manager is a comprehensive tool designed to be a key component in the systems management environment. It monitors the operations of HP servers, workstations, and clients. Insight Manager provides system administrators more control through visual interface, comprehensive fault and configuration management, and industry leading remote management.
- The integrated Lights-Out feature 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.

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 36GB hard drives preconfigured with the NAS operating system so that the active system volume is mirrored (RAID 1+0) 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 Model 1 also contains three 146GB 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+0, and 5. The NAS b2000 Model 2 ships with four 146GB hard drives in an external storage cabinet connected to the Smart Array 5304 controller.

Note: RAID 1+0 requires an even number of drives.

A power supply 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 integrated Lights-Out port 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 port 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 146GB disk drives storage capacity, it can grow up to 584GB of raw storage. Externally the NAS b2000 can support up to 27 terabytes of raw storage capacity when utilizing three Smart Array 5304 controllers, spanning 186 146GB hard drives spread over 13 StorageWorks 4300 Family storage enclosures and the internal drives.

Note: Each fully populated StorageWorks 4300 Family storage enclosure supports 14 hard drives.

Deployment Scenarios

The default shipping configuration contains two 10/100/1000 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. The integrated Lights-Out management port is also available. HP 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, HP 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 planned, the ability of the NAS b2000 to support most 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 confidence that it can support both CIFS and NFS simultaneously, 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 7 of 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)
- NetWare 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:

<http://www.microsoft.com/windows2000/technologies/directory/ad/default.asp>

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 Desktop

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.

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 shown in [Figure 1](#).

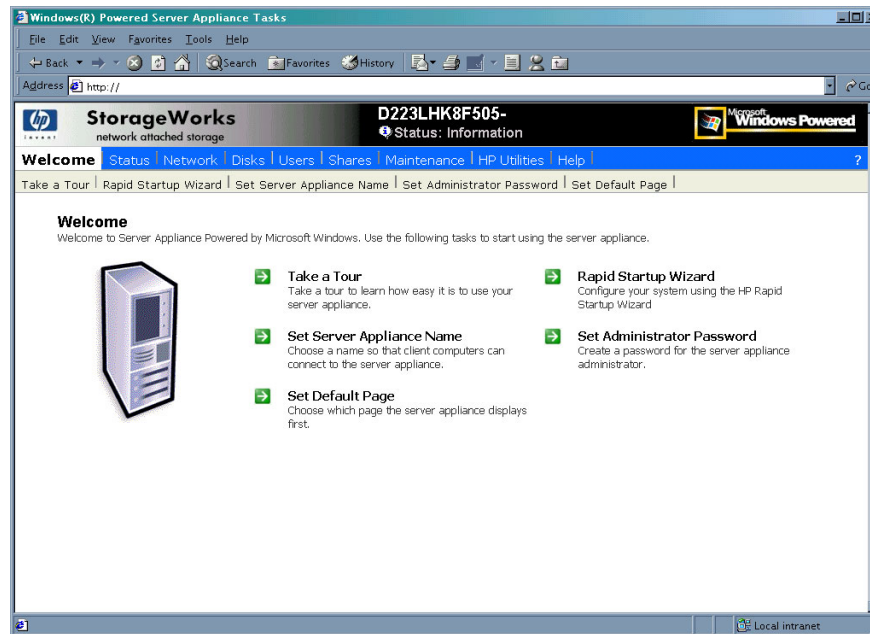


Figure 1: Primary WebUI screen

As shown in [Figure 1](#), the following areas are administered through this interface:

Status

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

Network

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

Disks

Use this option to manage disks, volumes, and disk quotas, and snapshots.

Users

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

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 Chapter 8 for additional information.

Maintenance

Maintenance tasks include setting date and time, performing system restarts and shutdowns, viewing audit logs, accessing Terminal Services, setting up Email alerts, linking to remote management, and HP System Management.

HP Utilities

Access HP system management utilities such as NAS Data Copy, remote management, enable floppy boot, and the HP System Management WebUI.

Help

This option contains help information for the WebUI.

Take a Tour

Take a tour and learn how to use your server appliance.

Rapid Startup Wizard

Use this utility to enter system setup and configuration information.

Set Server Appliance Name

Choose a name so that client computers can connect to the server appliance.

Set Administrator Password

Create a password for the server appliance administrator.

Set Default Page

Choose which page the server appliance displays first.

NAS b2000 Desktop

The NAS b2000 desktop can be accessed by:

- Directly connecting a keyboard and mouse
- Using the WebUI Maintenance tab and selecting **Terminal Services**
- Using the integrated Lights-Out port

Note: When using Terminal Services to connect to the NAS b2000 desktop do not use the window close feature (✕). Click on **Start/Log Off Administrator** to exit Terminal Services.




Figure 2: NAS b2000 desktop

The following icons are available from the Desktop:

- NAS Management Console
- CPQTeam Setup
- Install Data Copy

NAS Management Console

Click this icon to access the following folders:

- **Core Operating System** is used to manage local users and groups, access performance logs and alerts, and manage the event viewer.
- **Disk System** contains access to the Compaq Array Configuration Utility and local disk management, including a volume list and a graphical view of the disks.
- **File Sharing** contains modules for the configuration of file sharing exports. CIFS (Windows) and NFS (UNIX) file shares are managed through this folder.
- **System** contains system summary information.

CPQTeam Setup

Click this icon to install the Compaq Network Teaming and Configuration utility. See Chapter 2 for additional information on this feature.

Install Data Copy

Click this icon to install the trial version of the NAS Data Copy data replication software. See Chapter 2 for additional information on this feature.

Setup Completion and Basic Administrative Procedures

2

This chapter continues the process of setting up the system that was started using the *HP StorageWorks 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
 - Installing and configuring data replication software
 - Activating the iLO port using the license key
- 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
 - Updating the software
 - 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
- Installing and 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.

Note: 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
- NIC teaming troubleshooting

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 after selecting Start/Run:

```
c:\winnt\bin\nicteam\EN\cpqsetup.exe (English)
```

```
c:\winnt\bin\nicteam\JP\cpqsetup.exe (Japanese)
```

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

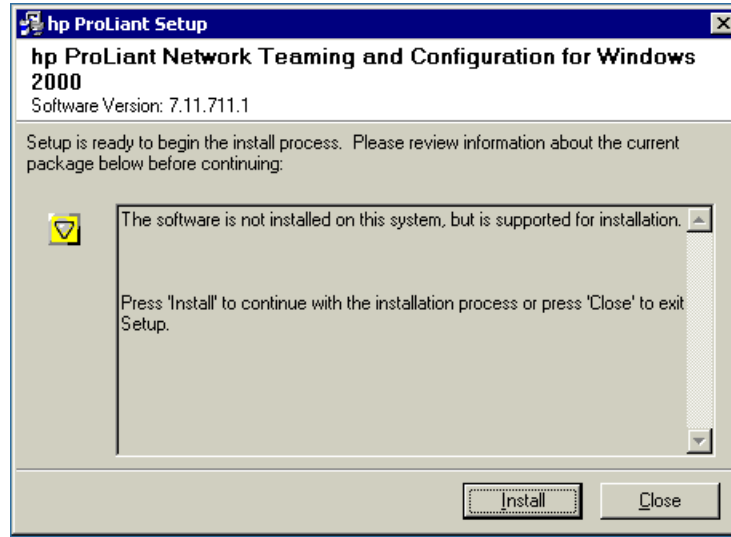


Figure 3: Installing Network Teaming

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

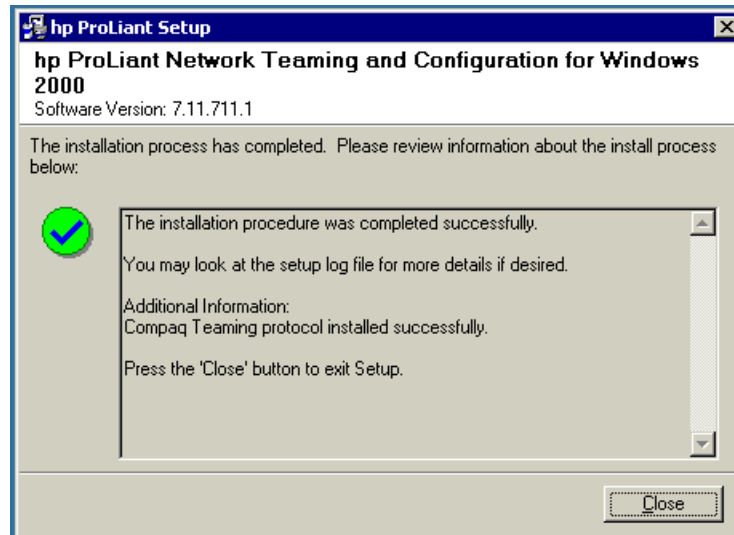


Figure 4: Network Teaming installation complete

4. Restart the system.



Caution: To ensure proper functioning of the software you must restart the server at this time.

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 **CPQTeam utility** icon.




Figure 5: 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. The type of NIC and the slot and port used is shown.

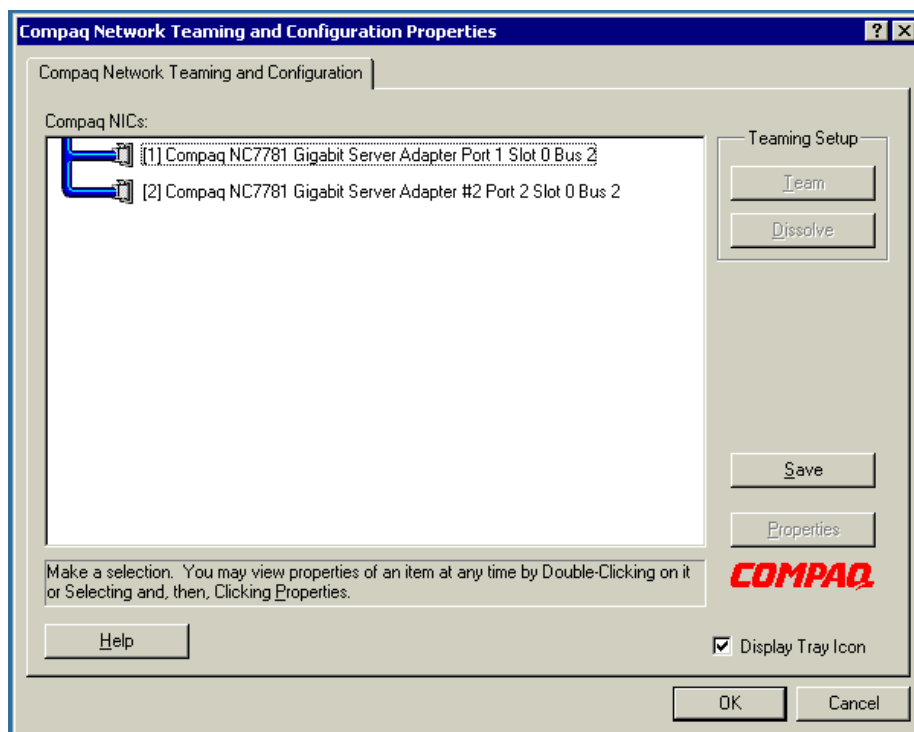


Figure 6: CPQTeam Properties dialog box

2. Highlight the NICs to team.

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

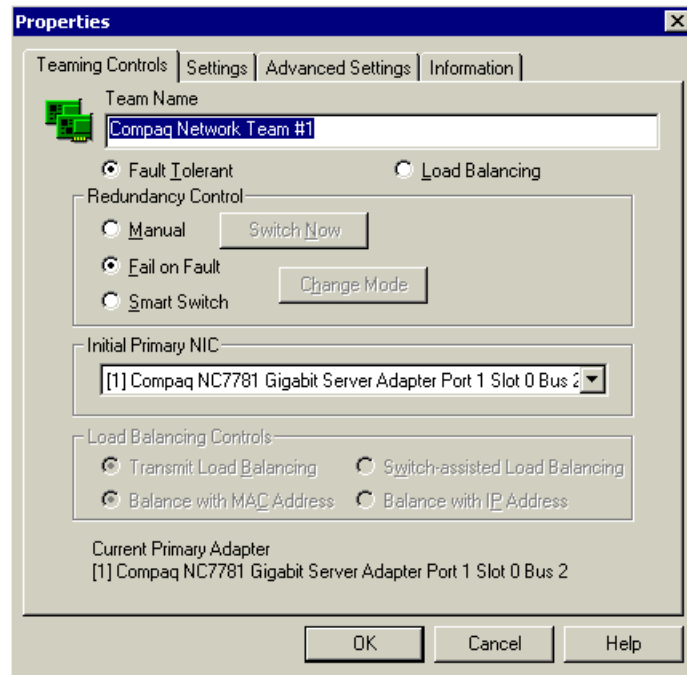


Figure 7: NIC Properties, Teaming Controls tab, Fault Tolerant option

- Configure the team by choosing either **Fault Tolerant** 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 you change from a Primary NIC to a Secondary NIC only when you click **Switch Now**.

Note: The **Switch Now** option is disabled until you select **Manual** and then click **OK**.

- **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 8: NIC Properties, Teaming Controls tab, Load Balancing option

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

- **Transmit Load Balancing**—All transmit IP frames are load balanced among all team members based on a Load Balancing algorithm in the teaming device driver. The Current Primary adapter transmits all other frames, and receives all frames for the team. If a failover event occurs, one of the non-Primary adapters assumes the role of Current Primary adapter, and transmit IP packets are load balanced among all remaining team members. If a failure occurs in any of the non-Primary adapters, the packets are load balanced among all remaining team members.
- **Switch-assisted Load Balancing**—All transmit packets are load balanced among all team members based on a Load Balancing algorithm in the teaming device driver. The receive packets are load balanced among all team members by the switch. If a failure of any team member occurs, the packets are load balanced among the remaining adapters. There is no primary adapter in a Switch-assisted Load Balancing team.
- **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 following 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 following 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.

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

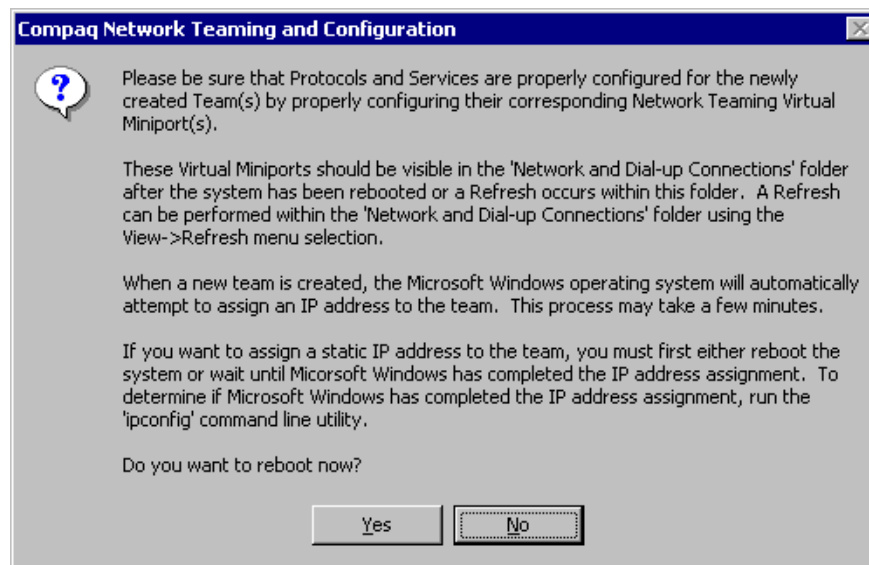


Figure 9: CPQTeam dialog box

6. 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. HP 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."

Showing 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, by default the new adapter is set to DHCP. To manually configure the IP address, perform the following steps.

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.

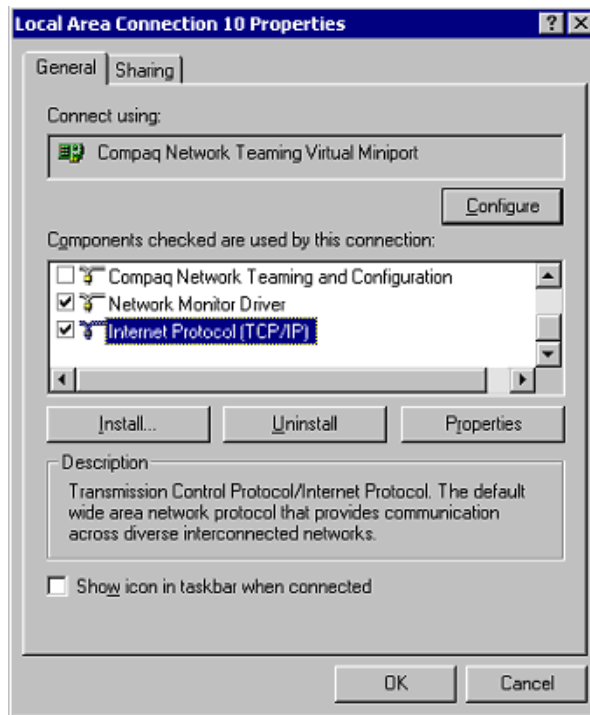


Figure 10: NIC Team Properties dialog box

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:

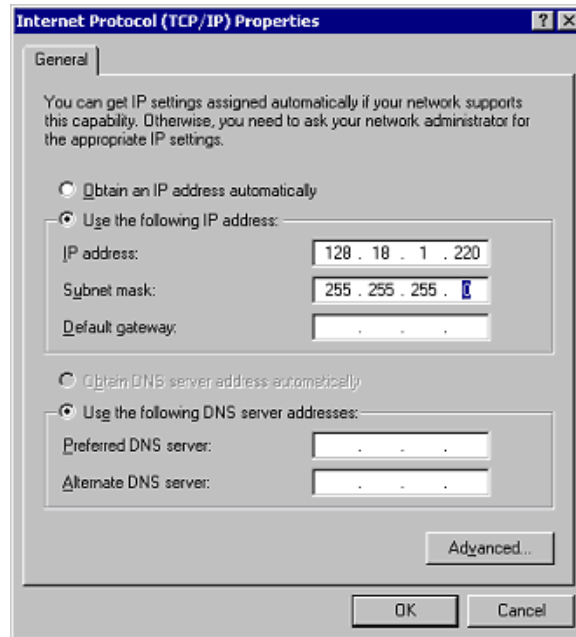


Figure 11: NIC Team TCP\IP Properties dialog box

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

4. 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.

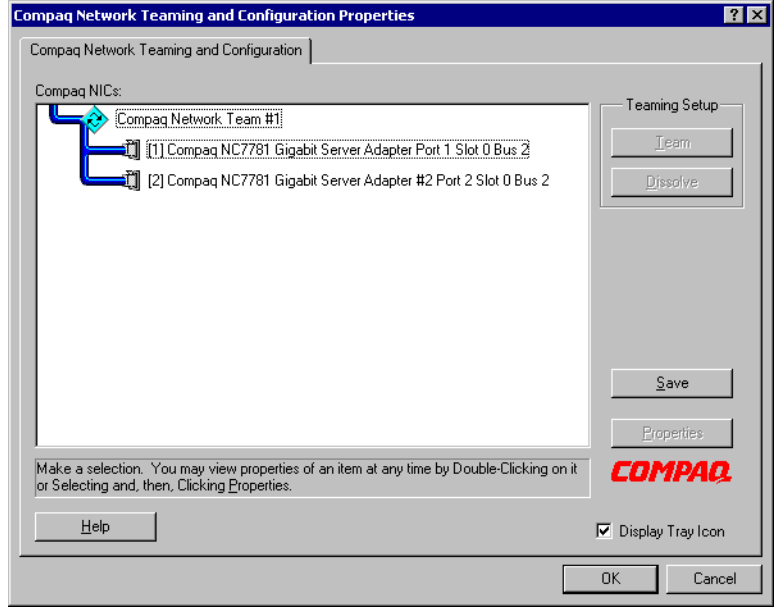



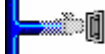





Figure 12: Updated CPQTeam Properties dialog box

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.

Table 2: NIC Teaming Troubleshooting

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.
	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 HP authorized service provider.
	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. Note: Only NICs assigned as members of a team are displayed as Unknown. If a teamed NIC is turned off, it displays as Unknown.
	Disabled—The NIC has been disabled through the Device Manager or NCPA.

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

Managing System Storage

The NAS administrator uses the Compaq Array Configuration Utility (ACU) to manage the storage hardware, Logical Disk Manager to manage volumes, and Persistent Storage Manager to manage snapshots. See the following chapters for more detailed information on managing system storage:

- Chapter 4 discusses system storage planning decisions in detail.
- Chapter 5 discusses hard drive, array, and LUN management procedures.
- Chapter 6 discusses snapshot management procedures.
- Chapter 8 discusses folder and share management procedures.

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 Chapter 7:

- Domain compared to workgroup environments
- User and group name planning
 - Managing user names
 - Managing group names
- Workgroup user and group management
 - Managing local users
 - Managing local groups
- Drive quotas
 - Managing quotas
 - Enabling and disabling quota management
 - Creating new quota entries for a user or group
 - Deleting new quota entries for a user or group
 - Modifying new quota entries for a user or group

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 Chapter 8:

- 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
- Share Management
 - Share considerations
 - Defining Access Control Lists
 - Integrating local file system security into Windows domain environments
 - Comparing administrative (hidden) and standard shares

- Planning for compatibility between file sharing protocols
- Managing shares
- Protocol parameter settings

UNIX specific information is discussed in the "UNIX File System Management" chapter.

Installing and 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.

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 HP. Further information can be found at the HP website.

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.

To install the trial version of NAS Data Copy:

- Select **Data Copy** from the **HP Utilities** tab or,
- Double-click on the **Install Data Copy** icon on the NAS b2000 desktop

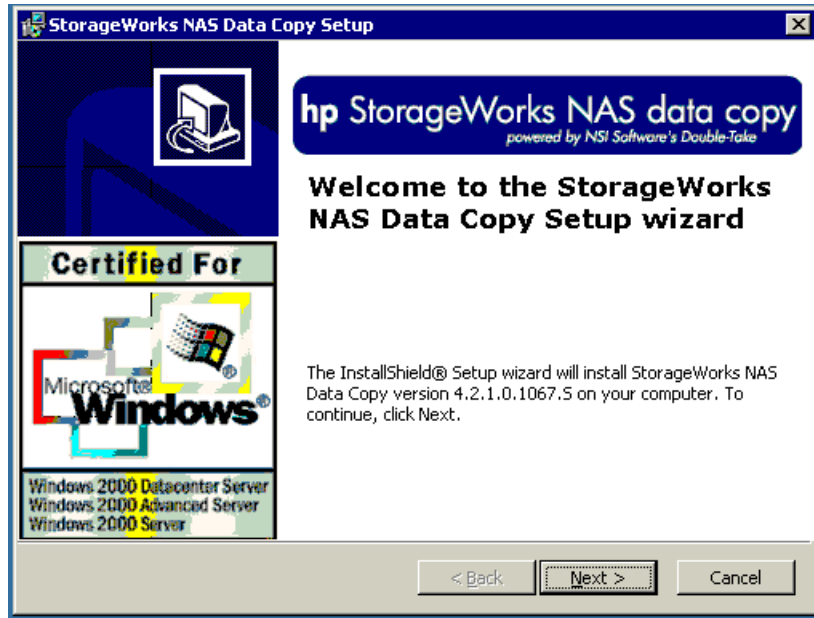


Figure 13: NAS data copy install wizard

Follow the onscreen instructions to complete the installation.

Activating the iLO Port Using the License Key

To activate the iLO port, locate the Integrated Lights-Out Advanced Pack License Installation Card and follow the enclosed instructions.

To configure the iLO port, click on **HP Utilities**, then click **Remote Management**.

Basic Administrative Procedures

Basic administrative procedures include:

- Setting the system date and time
- Shutting down or restarting the server
- Viewing and maintaining audit logs
- Using Terminal Services
- Setting up email alerts
- Updating the software
- Changing system network settings

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

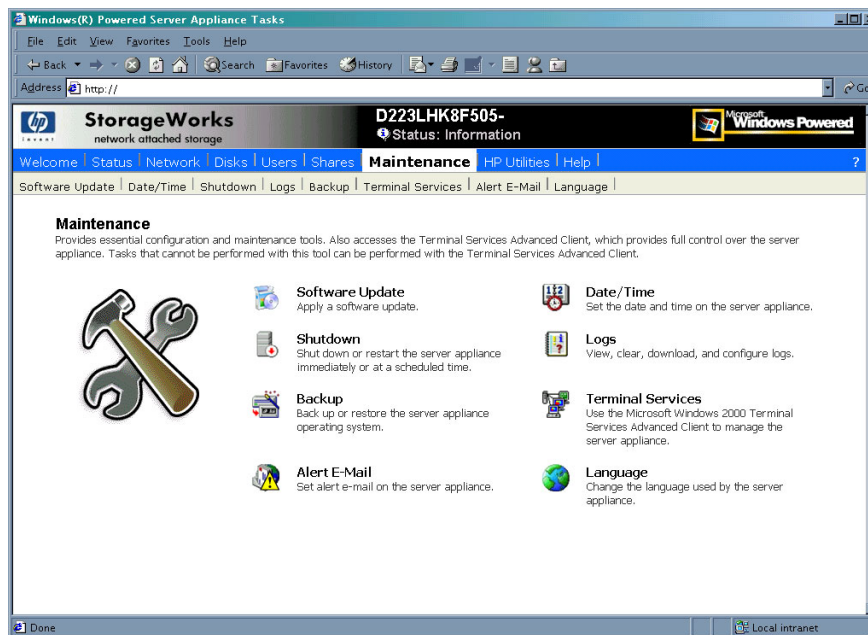


Figure 14: Maintenance menu

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.

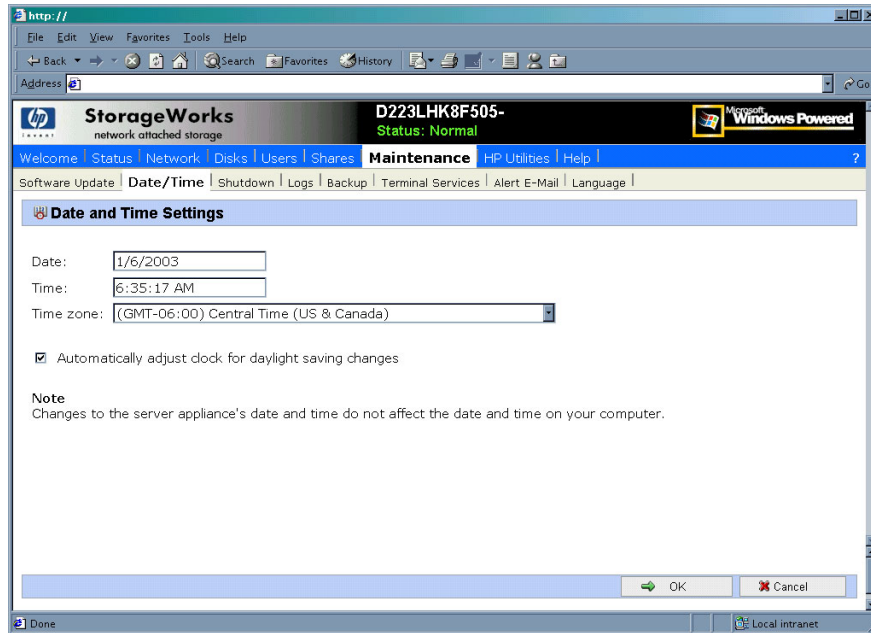


Figure 15: Date and Time dialog box

Shutting Down or Restarting the Server



Caution: 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: **Restart**, **Shut Down**, and **Scheduled Shutdown**.

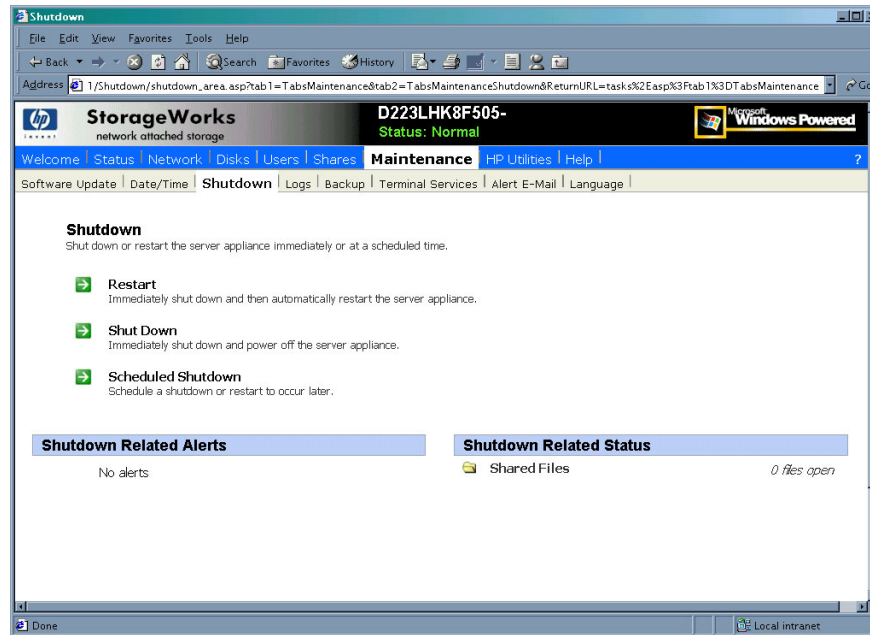


Figure 16: Shutdown menu

- a. To shut down and automatically restart the server, click **Restart**.
 - b. To shut down and power off the server, click **Shut Down**.
 - 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 seven different logs.

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

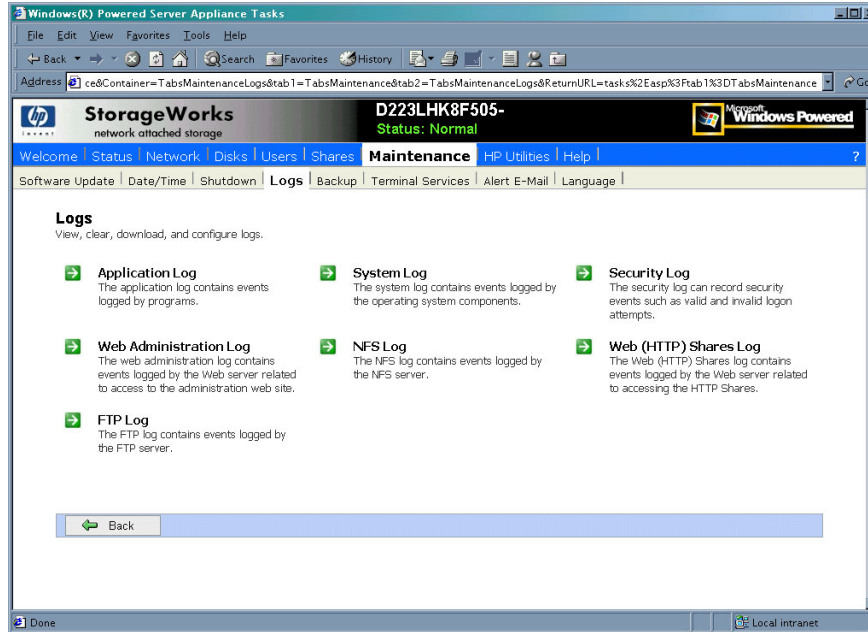


Figure 17: Logs menu

A variety of logs are available and are listed in [Figure 17](#).

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. Backup software and antivirus programs are examples of approved applications.

In addition, Terminal Services is used to access the NAS Management Console 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.

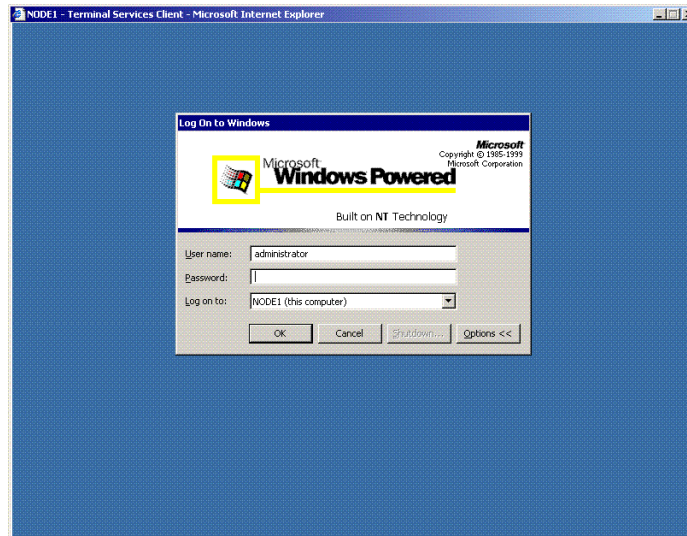


Figure 18: Terminal Services session



Caution: Two open sessions of Terminal Services are allowed to operate at the same time. After completing an application do not use the window close feature (X) to close that session of Terminal Services. Click on **Start/Log Off Administrator** to exit Terminal Services.

Setting up E-mail Alerts

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

To activate this option:

1. From the WebUI, select **Maintenance, Alert E-mail**. The **Set Alert E-Mail** dialog box is displayed.
2. Select **Enable Alert E-mail**.
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. After all settings have been entered, click **OK**.

Updating the Software

To update the software, click on **Software Update** from the **Maintenance** menu. The Software Update Wizard will guide you through selecting, verifying, and updating the desired software.

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.

Online help is available for these settings. [Figure 19](#) is an illustration of the Network settings menu.

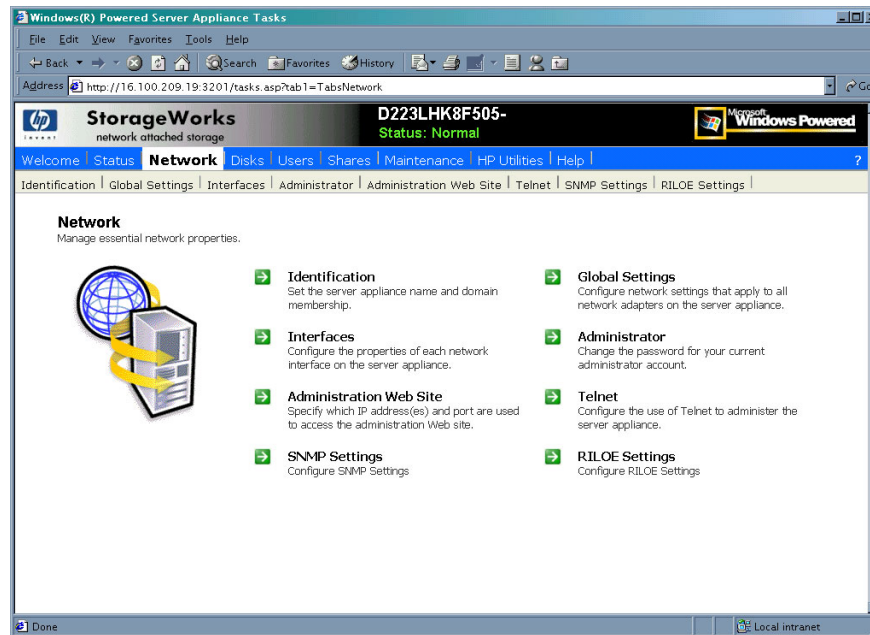


Figure 19: Network menu

Storage Management Overview

3

With the HP 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, Logical Disk Manager to manage the volume level, and Persistent Storage Manager to manage the snapshots.

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.

This chapter defines and discusses physical, logical, and snapshot storage concepts on the StorageWorks NAS b2000, including:

- Storage Management Process
- Storage Elements Overview
- Logical Storage Elements Overview
- Persistent Storage Elements Overview
- File System Elements Overview
- File Shares Elements Overview

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 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, Logical Disk Manager is used to take the LUNs and create logical volumes that can be basic or dynamic, which can then be broken down into partitions or volumes. Folders, subfolders, and file shares are created on the resulting volumes or partitions to organize, store, and give access to the system data. Persistent Storage Manager is used to create snapshots of the data at specific times.

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

See [Figure 20](#) for an illustration of these storage management elements.

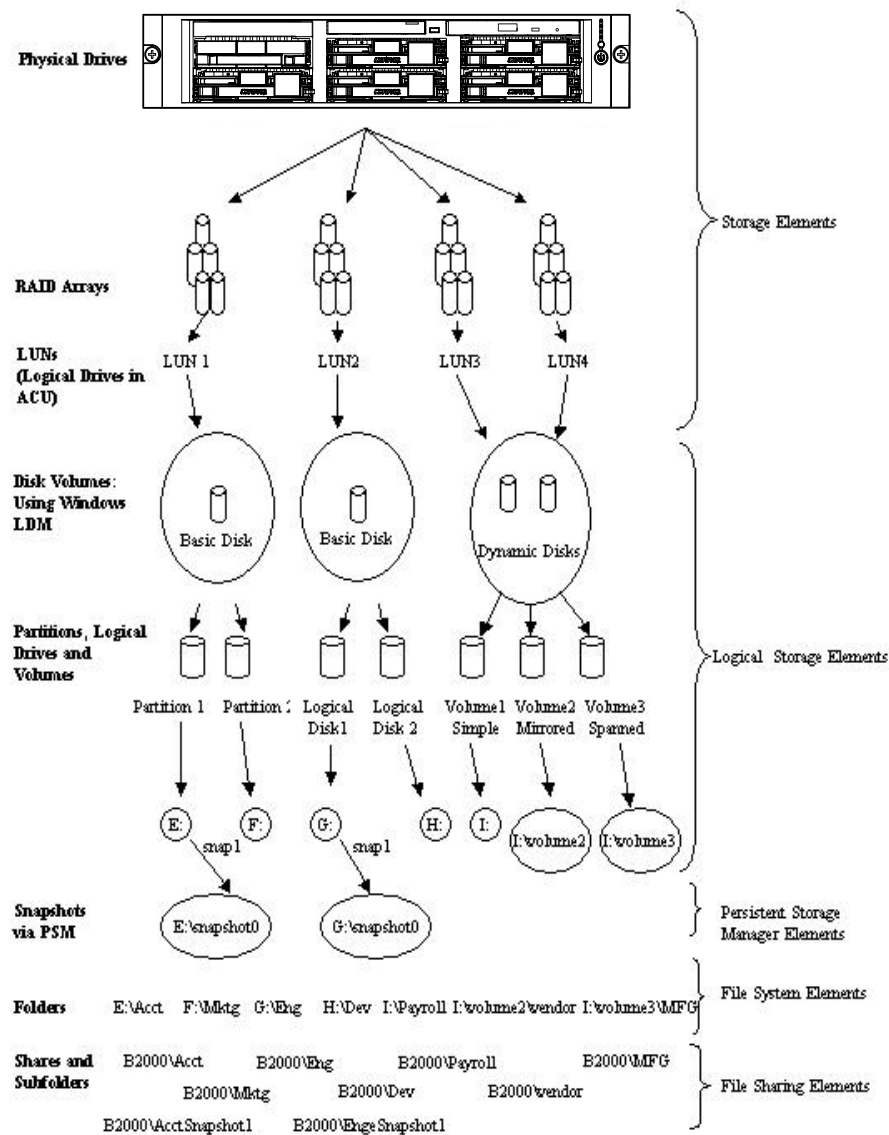


Figure 20: Storage Management process

Storage Elements Overview

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 146GB disk drives storage capacity, it can grow up to 291.2 GB of raw storage. Externally the NAS b2000 can support up to 27 terabytes of raw storage capacity when utilizing three Smart Array 5304 controllers, spanning 186 146GB hard drives spread over 13 StorageWorks 4300 Family storage enclosures and the internal drives.

Note: Each fully populated StorageWorks 4300 Family storage enclosure supports 14 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 21 illustrates the read/write process with separate physical hard drives.




Figure 21: 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 called 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 22: Configuring the physical drives into an array 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 23](#).




Figure 23: 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.

Note: 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 dynamic disk is 32 and the largest single LUN that can be utilized by the operating system is 2 TB. It should also be noted that the largest basic disk that can exist is 2 TB and the largest volume that can exist is 64 TB. Format of the partition or volume impacts the largest file system that can exist as well. A single NTFS partition is limited in size based on the allocation size used when formatting the disk ranging from a maximum size of between 2 TB (2^{32} allocation units x 512 bytes/allocation unit) and 256 TB (2^{32} allocation units x 65536 bytes/allocation unit).

Note: LUNs should not be expanded after they are created because Windows 2000 Advanced Server does not support the altering of 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 be designated as dynamic disks and then volumes can be expanded.

After the physical drives are grouped into 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 that can be of different RAID levels. HP 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 24: 2 arrays (A1, A2) and 5 logical drives (L1 through L5) spread over 5 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.

Note: The ADG feature is available only with the optional Smart Array 5300 Controller installed. RAID 1 support requires an even number of drives. Refer to [Table 3](#) 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.

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, HP recommends including no more than 7 drives in a RAID 0 array.

See [Figure 23](#) 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 25](#). 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 25: 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.

Note: HP supports a configuration that uses RAID 1 on the system drives in a two drive RAID array.



Caution: 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

Using 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 26: 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. HP 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—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 27: 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 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 Smart Array controllers. The decision chart in the following table may help determine which option is best for different situations.

Table 3: Summary of RAID Methods

	RAID 0 Striping (no fault tolerance)	RAID 1 / RAID 1+0 Mirroring	RAID 5 Distributed Data Guarding	RAID ADG 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
Read performance	High	High	High	High
Write performance	High	Medium	Low	Lowest
RAID overhead	Low	High	Medium	Medium

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 56 drives in each array.
 - For RAID 1+0, include between 2 and 56 drives in each array.
 - For RAID 5, include between 3 and 14 drives in each array.
 - For RAID ADG, include between 4 and 56 drives in an array.

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.

Logical Storage Elements Overview

Logical Storage elements consist of those components that translate the physical storage elements to the file system elements as presented in [Figure 20](#). The b2000 utilizes the Logical Disk Manager (LDM) for managing the various types of disk presentation to the file system. LDM has two types of LUN presentation, basic disk and dynamic disk. Each of these types of disk has special features that enable different types of management. Through the use of basic disks, partitions or extended partitions may be created. Partitions can only encompass one LUN. Through the use of dynamic disks, volumes can be created that span multiple LUNs. The sections below discuss in brief each of these types of representations and the considerations that need to be observed. More detailed information regarding LDM use can be obtained through the online help of the tool and the Microsoft website.

Partitions

Partitions exist as either Primary Partitions or Extended Partitions and can be composed of only one Basic disk no larger than 2 TB. Basic disks can also only contain up to 4 primary partitions and 1 extended partition. In addition, the partitions on them cannot be extended beyond the limits of a single LUN nor can they be altered once they are created. Extended partitions allow the user to create multiple logical drives, but they cannot be altered once they are created. These partitions or logical disks can be assigned drive letters or be mounted as mount points on existing disks. If mount points are utilized, it should be noted that Services for Unix does not support mount points at this time. When creating mount points, meaningful volume labels should be utilized to identify them in Persistent Storage Manager since PSM utilizes the volume label for managing the snapshots.

Volumes

When planning dynamic disks and volumes there is a limit to the amount of growth a single volume can undergo. Volumes are limited in size and are limited to no more than 32 separate LUNs with each LUN not exceeding 2 terabytes (TB). Volumes also cannot exceed 64 TB of disk space. Additionally, a single NTFS partition is limited in size based on the allocation size used when formatting the disk, ranging from a maximum size of between 2 TB (2^{32} allocation units x 512 bytes/allocation unit) and 256 TB (2^{32} allocation units x 65536 bytes/allocation unit).

The RAID level of the LUNs included in a volume must be considered. All of the units that make up a volume 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+0 and a RAID 5 array in the same volume set. By keeping all the units the same, the entire volume retains the same performance and high-availability characteristics, making managing and maintaining the volume much easier. It should be noted that if a dynamic disk goes offline, then the entire volume dependent on the one or more dynamic disks is unavailable. There could be a potential for data loss depending on the nature of the failed LUN.

Volumes are created out of the dynamic disks and can be expanded on the fly to extend over multiple dynamic disks if they are spanned volumes. However, once a type of volume is selected it cannot be altered, i.e. a spanning volume cannot be altered to a mirrored volume without deleting and recreating the volume, unless it is a simple volume. Simple volumes can be mirrored or converted to spanned volumes. Fault tolerant disks cannot be extended either. Therefore, selection of the volume type is important. Please note that the same performance characteristics on numbers of reads and writes apply when using fault tolerant configurations as is the case with controller based RAID. These volumes can also be assigned drive letters or be mounted as mount points off existing drive letters. In general, HP recommends utilizing the Array controller for the management of fault tolerance over the use of LDM since LDM places an additional level of operating system overhead on volumes. If mount points are utilized, it should be noted that Services for Unix does not support mount points at this time. When creating mount points, meaningful volume labels should be utilized to identify them in Persistent Storage Manager since PSM utilizes the volume label for managing the snapshots.

The administrator should carefully consider how the volumes 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 dynamic disk set would not be efficient. These applications or groups would be better served by being divided up into separate dynamic disks, which could then grow as their space requirements increased, following the allowable growth limits.

Utilizing LDM Storage Elements

No matter which type of storage element is created in LDM the last step in creating the element is determining its drive letter or mount point and formatting the element. Each element created can exist as a drive letter(s), assuming one is available and/or as mount points off of an existing folder of a drive letter. Either method is supported. However, mount points can not be utilized for shares that will be shared using Microsoft Services for Unix (NFS). They can be setup with both but the use of the mount point in conjunction with NFS shares causes instability with the NFS shares.

Formats consist of NTFS, FAT32, and FAT and all three types can be used on the NAS device. However, Persistent Storage Manager can only utilize volumes that are NTFS formatted.

Persistent Storage Management Elements Overview

Persistent Storage Manager lets the administrator make replicas, called snapshots, of disks in a matter of seconds. Snapshots enable the 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.

Snapshots use existing space from the volume, partition, or logical drive to maintain the data required to present the original data. This space is called the cache file. By default the cache file consumes 10 percent of the available space of a Logical Storage element. Snapshots can be read only, read write or always keep, and if they are shared, users can access a snapshot and edit the data. If snapshots are shared with write access enabled, the snapshot will revert if changed back to read only using PSM.

Snapshot Facts:

- Snapshots are created on a per volume, partition or logical drive basis.
- Snapshots can be read only, read write, or always keep.
- Snapshots are mounted as a mount point on the root of the volume, partition, or logical drive.
- Snapshots can be shared in the same manner as any other folder, drive, or mount point.
- Snapshots are meant to be temporary.
- Snapshots are automatically deleted if disk space becomes critical and they are not set to always keep.
- Persistent Storage Manager only writes to the cache file on the first change of the underlying data.

Detailed information on Persistent Storage Manager can be found in Chapter 6 of this guide.

File System Elements

File system elements are composed of the folders and subfolders that are created under each Logical Storage element (partitions, logical disks, and volumes). Folders are used to further subdivide the available file system providing another level of granularity for management of the information space. Each of these folders can contain separate permissions and share names that can be used for network access. Folders can be created for individual users, groups, projects, and so on.

Detailed information on file system elements can be found in Chapter 8 of this guide.

File-Sharing Elements

The NAS b2000 supports several file sharing protocols, including CIFS, NFS, FTP, HTTP, NCP, and AppleTalk. On each folder or Logical Storage element, different file sharing protocols can be enabled using specific network names for access across a network to a variety of clients. Permissions can then be granted to those shares based on users or groups of users in each of the file sharing protocols.

Detailed information on file-sharing elements can be found in Chapter 8 of this guide.

Advanced Storage Management Planning

4

This chapter details issues surrounding storage configuration, storage sizing, and performance planning for the HP 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 (referred to as Logical Drives in ACU)
 - Storage Sizing 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
- Storage Capacity Expansion 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, disks, and partitions extensive analysis of desired system performance must be completed. 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 28](#), 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.

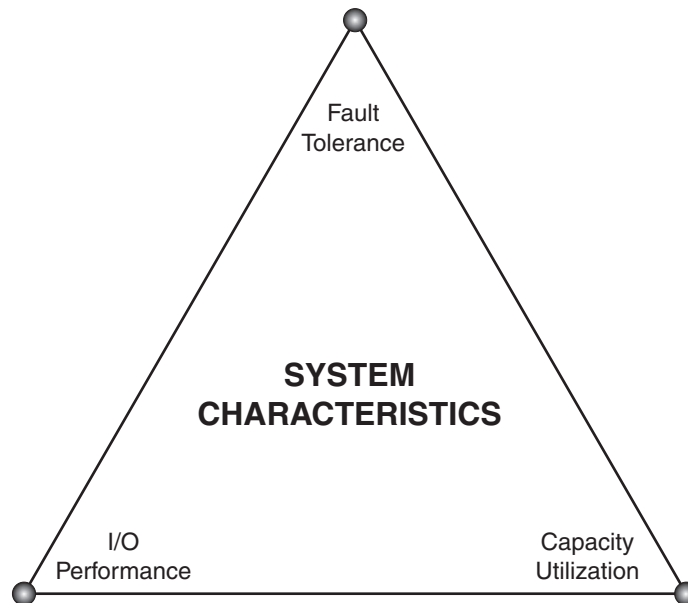


Figure 28: System characteristics

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.

Note: Smart Array controllers support up to 32 logical drives per controller.

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 (only applicable when using Smart Array 5300 controller and storage enclosures)
- 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 29](#) and [Figure 30](#) 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). To implement vertical striping of RAID arrays, HP recommends using a fully populated Smart Array 5304 controller with four disk storage enclosures. This configuration allows for greater choices of RAID configurations, as well as for the creation of larger arrays.

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 consists of at least one drive from each storage enclosure. See [Figure 29](#) and [Table 4](#) for RAID specific disk use information of vertical carving configurations.




Figure 29: Vertical Array configurations

Table 4: 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.

Note: 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. The b2000 does support mount points in CIFS environments and disks may be concatenated using dynamic disks. Both approaches help to address the exhaustion of disk 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, i.e. with only a maximum of four cabinets to span and retaining redundancy the arrays are limited to four physical disks. Using small arrays means that the same amount of storage would also occupy more LUNs. An additional disadvantage of vertical configurations is that the creation of many small arrays, LUNs, operating system disks, 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 30](#) and [Table 5](#) for information about horizontal array configurations and disk use.




Figure 30: Horizontal Array configurations

Table 5: Horizontal Configuration Disk Use per RAID Level

RAID Level	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 logical storage elements.

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 31](#) for an illustration of this technique.




Figure 31: NSPOF Horizontal Array configuration, RAID 1+0

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 31](#) 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.) Four 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. [Figure 6](#) provides a list of the striping and array configuration methods available for the different hardware configurations.

Table 6: Suggested Storage Enclosure Configurations

Striping Method	One 4300 Family Storage Enclosure	Two 4300 Family Storage Enclosures	Three 4300 Family Storage Enclosures	Four 4300 Family Storage Enclosures
Vertical (NSPOF)	None	RAID 1+0	RAID 1+0 RAID 5	RAID 1+0 RAID 5 RAID ADG
Horizontal	RAID 0 RAID 1 +0 RAID 5 RAID ADG	RAID 0 RAID 1+0 RAID 5 RAID ADG	RAID 0 RAID 1+0 RAID 5 RAID ADG	RAID 0 RAID 1+0 RAID 5 RAID ADG
NSPOF Horizontal	None	RAID 1+0	None	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 32 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.

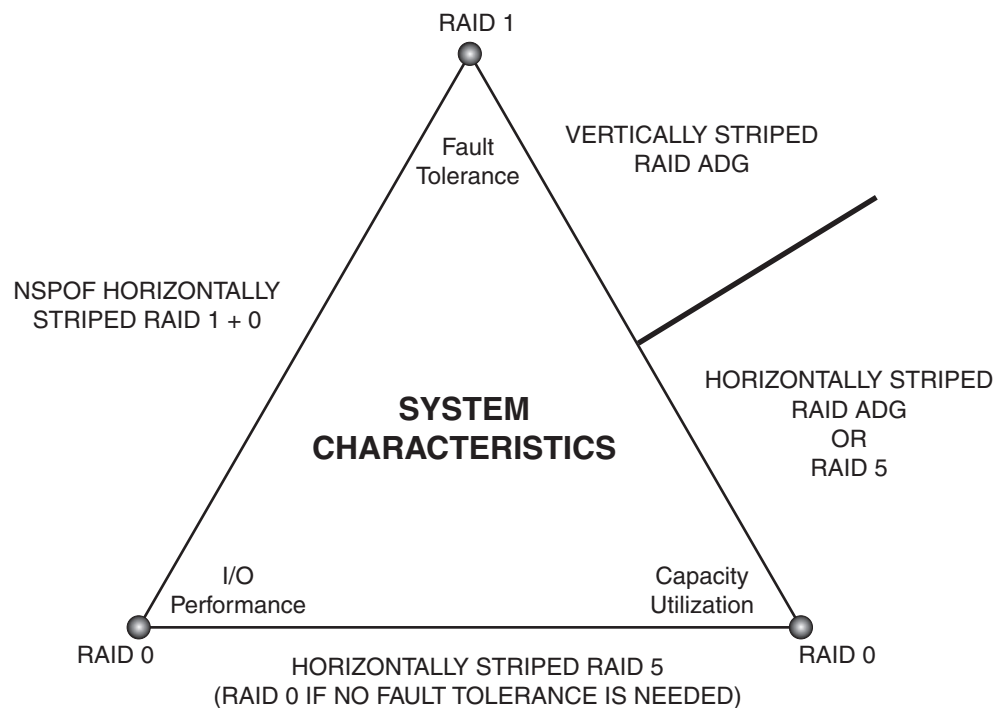


Figure 32: Recommended Configuration methods

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 32, 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 32](#) 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.

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.

Although it is an extreme example, an array with one 18.2 GB drive and two 72.8 GB drives results in a waste of 109.2 GB ($72.8 - 18.2 = 54.6$ GB per disk x 2 disks = 109.2 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.

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. HP 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

HP 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 because LUNs cannot be extended.

These limitations lead to the recommendation that system administrators create large LUNs. Each of these points is discussed in further detail in the following paragraphs.

LUNs Cannot be Extended

While the array controller will allow the user to grow LUNs (called logical drives in the Array Configuration Utility) after they have been created, it is not supported by Windows 2000. In addition, only the most newly created LUN on an array can be extended. In order to safely grow a LUN, the LUN would have to be deleted and then recreated including more space for the LUN. Doing this results in the loss of the data in the LUN. In some cases the array might require additional disks. After the decision to recreate the arrays has been made, but prior to actually deleting and recreating the LUN, all data on the underlying folders and 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 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 dynamic disk or used to create a new disk. 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).

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

LUN Management under Windows Powered OS

LUN management under Windows Powered OS is accomplished through the Logical Disk Manager native to Windows Powered OS. Complete details can be found in the online help for this tool. LUNs appear under LDM as disks that then need to have signature written to them, need the type of disk assigned, and need to be formatted. Depending on the type of disk the user selects impacts the way Windows Powered OS manages the LUNS. If the disk is configured as basic then each basic disk is added as a separate unit that can contain multiple partitions up to 4 primary partitions and 1 extended partition. If the disk is configured as dynamic, then volumes can be configured in the basic RAID configurations of simple, spanned, striped (RAID 0), mirrored (RAID 1), and RAID 5 and treated as a large disk. Volumes that are spanned can also be expanded by simply adding additional LUNS to them or encompassing free space from an existing unused dynamic disk.

In general, HP recommends utilizing the array controller for the management of fault tolerance over the use of LDM since LDM places an additional level of OS overhead on volumes when maintaining RAID sets.

Storage Sizing Considerations

Estimating the number and configuration of storage enclosures, disks, and arrays 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 issues

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

HP 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 choose to have multiple spare disks, so that multiple arrays can successfully recover from possible failure 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

As explained previously in Chapter 3, snapshots are a feature of Persistent Storage Manager and are a point in time view of an NT Volume. Even though snapshots do not use up any space initially, a cache file is created with the first snapshot on a volume. Snapshots create a copy of original data before allowing any changes to be made to the parent volume. Snapshot space usage tends to climb over time within the cache file, and the rate of usage is related to the rate of change of the data on a volume, the total lifetime of the snapshot, and the number of snapshots maintained for a volume.

Note: Detailed information on Persistent Storage Manager can be found in Chapter 6.

By default, the NAS b2000 default setting for snapshot cache reserve is 10 percent of the volume space exclusively for the use of snapshots. This percentage may be altered for those expecting to generate more or less snapshot data. Note the cache file is not established until the first snapshot is taken. Once it is taken cache file reserve cannot be altered.

In storage planning the administrator must include the increase in raw storage capacity needed by either 10 percent or the estimated amount of space that will be reserved for snapshot use per snapshot cache contained on each volume. This setting is per volume, hence one volume can be set to 10 percent and another set to 30 percent.

Growth Issues

It is also important to allow for significant changes in the total size of all files on a particular volume. 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 dynamic disks, used to create new volumes mounted on drive letters or mount points, or used to enlarge existing volumes if they are spanned volumes, and put into use.

Note: Storage enclosures are not hot-pluggable.

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 are 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, HP recommends always formatting the disk with an allocation unit size of 16 KB.

Note: Only disks with a 4KB allocation size and smaller can use common defrag tools so choose the allocation size wisely. Also as noted in the previous section, allocation unit size will impact the maximum size of a particular volume.

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 volume exactly matches that of the original volume, 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 volume, the allocation unit size of the source volume, 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 volume 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 volume is easily selectable when creating volumes 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 fileservers 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 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 7](#) 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. When the data on the disk changes, the original blocks are copied to the cache file on the same volume. 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, HP recommends that an additional 10 percent of disk space beyond the storage requirements be allocated for snapshots.

- Determine the percentage of the volume 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 cache file is deducted:

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

This scenario uses snapshots and uses the default snapshot cache file of 10 percent. It should be noted that the cache file is not established until the first snapshot is taken on a volume. The size can be reduced at a later time if no snapshots for that cache file on that volume exist.

Total Storage Need

By applying the Snapshot Cache File 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 Cache File Factor, the needed space is increased by a sufficient amount to allow for the snapshot reserve.

For example, this Total Storage Need is $1500 \text{ GB} / .80 = 1875 \text{ GB}$.

Table 7: 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 Cache File 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 Cache File 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 8](#) 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 HP Ultra 2 and Ultra 3 drives in the following capacities: 18 GB, 36.4 GB, 72.8 GB, and 146 GB.

In this scenario, HP 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 HP 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. Data from fourteen drives can 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 subsequent LUNs are used to build the volumes or partitions, the number of arrays in a system directly affects the size of the volumes or partitions 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:

1875 GB Total Storage Need / 436 GB Usable Space per Array = 4.3

Total Number of Arrays Required = 5 arrays required

Table 8: 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)	1875 GB / 436 = 4.3
	Total Number of Arrays Required	5 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 4-6 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:

5 Arrays x 7 Drives per Array = 35 Drives Required for all Arrays

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

HP 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.

35 Drives Required for the Arrays + 1 Spare Drive = 36 Total Drives Required

Table 9: Example Drives Required Worksheet

	Formula	Value
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 from the previous worksheet)	5 x 7 = 35 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	35 + 1 = 36 drives
	Total Drives Required	36 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.

Number of Storage Enclosures = Total Number of Drives Required / 14

For this scenario, 3 storage enclosures are required:

36 Total Number of Drives Required / 14 = 2.57 = 3 storage enclosures

Table 10: Example Enclosures Required Worksheet

	Formula	Value
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

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, LUNS (logical disks) and volumes.

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.

(56 drives - 2 drives for parity = 54 drives; 54 drives x 72.8 GB = 3931 GB)

Table 4 8 consolidates these calculations.

Table 11: 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%
56 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:

1. Analyze current data storage demands and determine the amount of data that will be migrated to the NAS b2000.
2. Determine the amount of space to allow for future growth.
3. Determine whether snapshots will be used, and if so, if a 10 percent reserve is adequate.
4. Determine the total amount of space needed.
5. Determine the sizes and types of hard drives that will be used.
6. 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.
7. Determine the number of arrays that will be required.
8. Determine the number of drives required to build these arrays.
9. Determine the number of spare drives to use to support these arrays.
10. Determine the number of drives required to build this system configuration.
11. Determine the number of storage enclosures necessary to hold all of the drives.
12. Adjust, rework, and finalize and the sizing and configuration plan.

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

Table 12: 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 Cache File 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 13: 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 14: Drive and Enclosure Requirements Worksheet

	Formula	Value
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.
- **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 three methods of actually moving the data from the old file server to the NAS b2000:

- Backup and restore
- Ethernet copy
- Drive Migration



Caution: 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 three methods are different, some of the preparatory and completion procedures of each method are the same. The 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.
2. Complete the Rapid Startup procedures.
3. Configure the NAS b2000 storage, excluding file shares.
4. Configure the arrays, LUNs, volumes, or partitions.

Note: When creating the volumes or partitions, 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 volumes or partitions 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.

5. Create any local users or groups.
6. Install the backup software onto the NAS device.
7. Create the backups of the data being migrated.
8. Restore the data from the backup onto the target volumes or partitions of the NAS device.

9. 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.

Drive Migration

Assuming that the file consolidation is occurring among ProLiant Smart Array based servers and storage. Drives may be physically relocated from the target systems to the b2000 in whole units. The Smart Array controllers supported on the b2000 will recognize these drives and present to the NAS OS. Retaining the drive positions is desirable during this migration but not necessary. For example if one has a ProLiant 380R with a Smart Array 4200 controller running Windows 2000 Advanced Server, one could migrate the entire cabinet from the older unit to the new one provided they are both using the new 1" carriers, either 42x4 or 43x4 cabinets and the base OS is Windows 2000 Advanced Server.

1. Power off both systems and the accompanying storage components to be moved.
2. Migrate the drive cabinets from the target system to the b2000, attaching it to the Smart Array 5300 that is optional.
3. Power up the system. The drives will be recognized and presented to the b2000. Shares and permissions will need to be reestablished.

Storage Capacity Expansion Issues

The NAS device is designed for expansion. Up to thirteen additional storage enclosures can be added, for a total of 186 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 and create new dynamic disks or basic disks
- Create a new volume(s) or partition(s) using the new logical drive or drives
- Use new dynamic disks to expand existing volumes

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

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 volume is thirty two. To have large dynamic disk from which to create volumes, make logical drives or LUNS as large as possible up to 2 TB.

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 as a dynamic disk or a basic disk.
5. Create new volumes or partitions, as needed.
6. Create the necessary folders and file shares.

Physical Storage Management

5

Physical storage includes the physical hard drives, as well as the initial configuration of the hard drives into arrays and logical units (LUNs).

Complete conceptual information on managing physical storage is discussed in the "Storage Management Overview" and "Advanced Storage Management Planning" chapters.

This chapter discusses the procedural aspects of managing the physical storage.

The storage topics discussed in this chapter include:

- **Hard Drive Management**
 - Defining Hard Drive LED Indicators
 - Replacing Failed Hard Drives
 - Moving Hard Drives
 - Moving Arrays
- **Array and LUN Management**
 - ACU Overview
 - Accessing the ACU
 - Entering Controller Settings
 - Creating a New Array
 - Creating Logical Drives (LUNs)
 - 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:

- Defining Hard Drive LED Indicators
- Replacing Failed Hard Drives
- Moving Hard Drives
- Moving Arrays

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 *HP 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 *HP Servers Troubleshooting Guide*.

Use the following illustration in conjunction with [Table 15](#) to analyze the status of hot plug hard drives.




Figure 33: Hot plug hard drive LED indicators

Table 15: Hard Drive LED Combinations

Activity	Online	Drive Failure	Indication
On	Off	Off	Do not remove the drive. Removing a drive during this process causes data loss. The drive is being accessed and is not configured as part of an array.
On	Flashing	Off	Do not remove the drive. Removing a drive during this process causes data loss. The drive is rebuilding or undergoing capacity expansion.
Flashing	Flashing	Flashing	Do not remove the drive. Removing a drive during this process causes data loss. The drive is part of an array being selected by the ACU. -Or- The drive is being upgraded.
Off	Off	Off	OK to replace the drive online if a predictive failure alert is received (see the "Predictive Failure Alert" section in <i>HP Servers Troubleshooting Guide</i> for details) and the drive is connected to an array controller. The drive is not configured as part of an array. -Or- If this drive is part of an array, then a powered up controller is not accessing the drive. -Or- The drive is configured as an online spare.
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>HP 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>HP 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.
- 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 HP utilities report drive information in terms of SCSI ID. Only the physical drive enclosure is labeled by bay number. Use [Table 16](#) to properly identify the appropriate drive bay.

Note: HP utilities report drive state through the SCSI ID. [Table 16](#) assists you in associating the appropriate drive bay to the corresponding SCSI ID.

Table 16: 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 dynamic disks or basic disk and corresponding volumes or partitions, 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.

Note: 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.



Caution: 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 16](#), "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 may be stored in volumes created from dynamic disks 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 volume 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 16](#).

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, and RAID levels.

In addition, the "Advanced 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 Array Configuration Utility (ACU). Logical Disk Manager then uses the LUNs to create dynamic disks, basic disks, volumes, partitions, and snapshots.

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

- ACU Overview
- Accessing the ACU
- Creating a New Array
- Creating Logical Drives (LUNs)
- Expanding the Capacity of an Existing Array
- Migrating an Existing LUN to a New RAID Level or Stripe Size

ACU Overview

The ACU is a graphical tool, incorporating wizard style interfaces used 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 volumes in Logical Disk Manager.

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.

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.

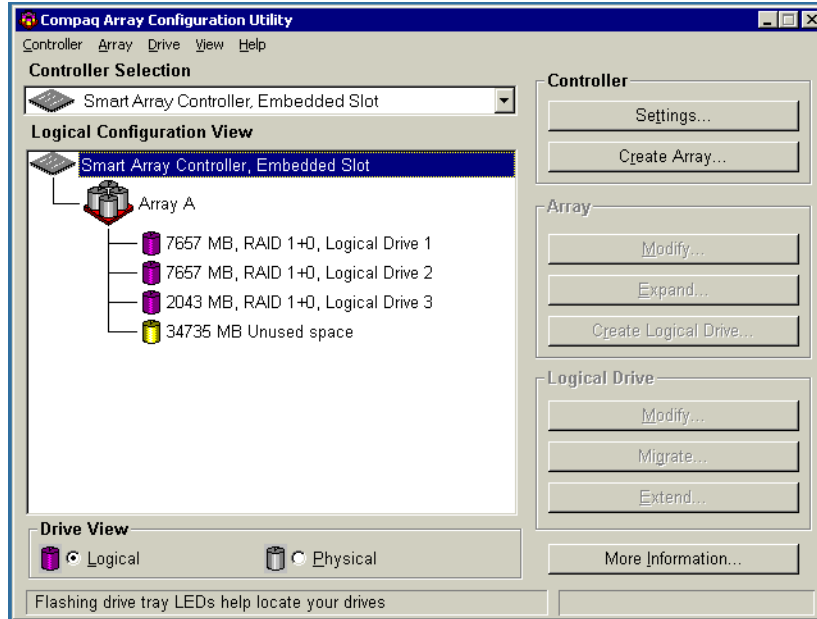


Figure 34: ACU Logical Drive view

Select either **Logical Drive View** or **Physical Drive View**. [Figure 34](#) and [Figure 35](#) are examples of the two presentation methods.

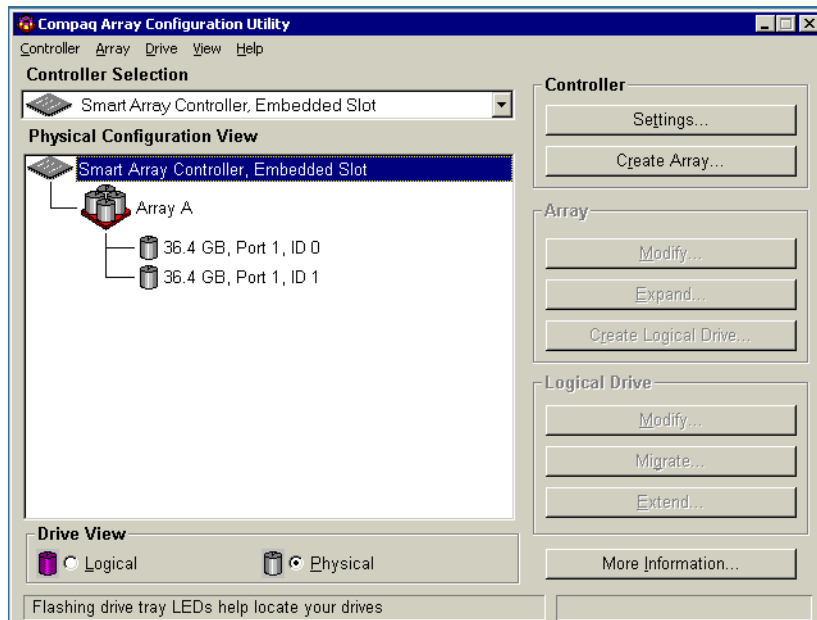


Figure 35: ACU Physical view

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.

- 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.
- 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:

- 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.

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

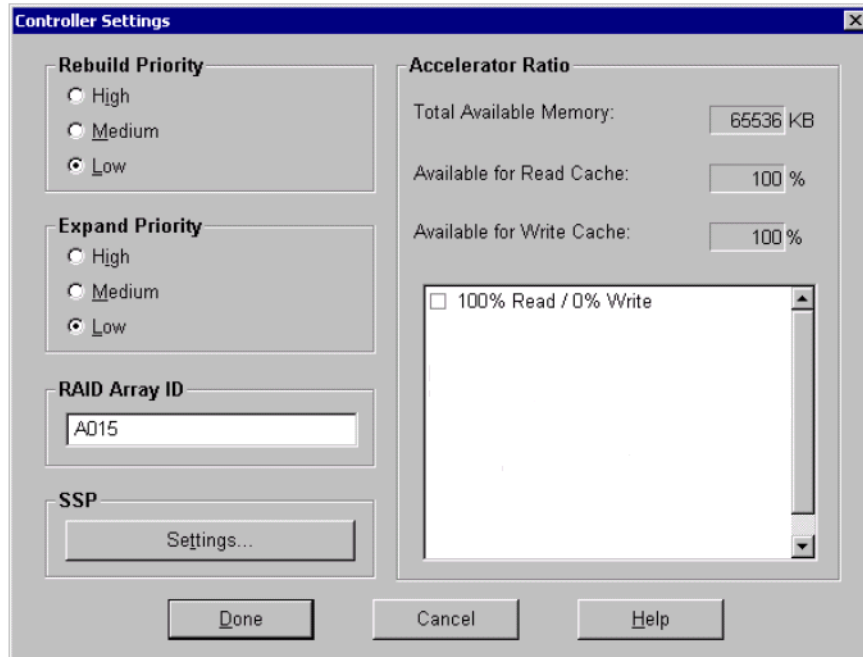


Figure 36: Controller Settings dialog box

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**. (Only applicable if you have the battery backed cache module.)

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 "Advanced 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.

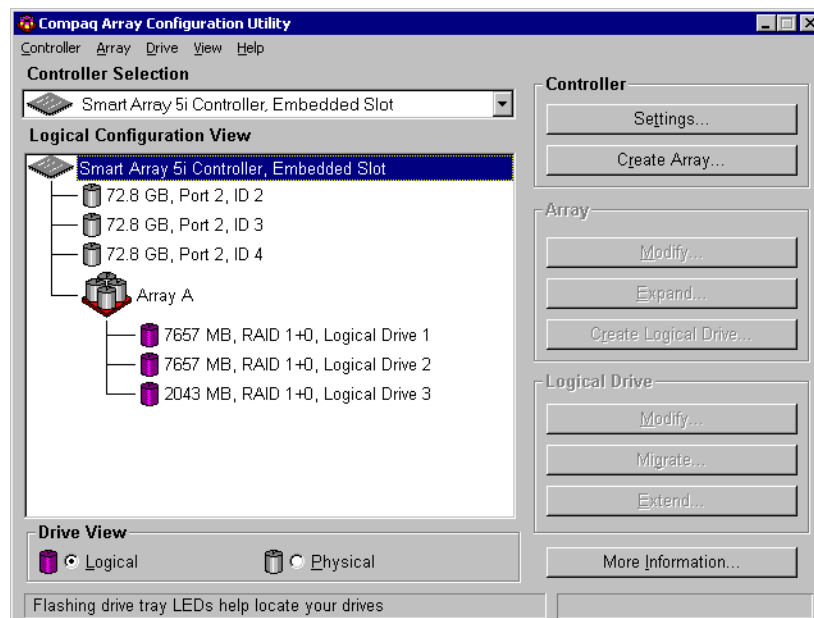


Figure 37: ACU main configuration screen

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

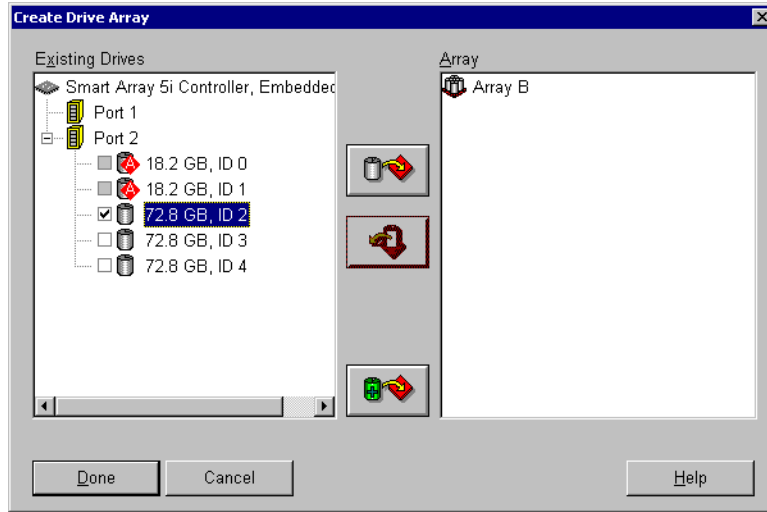


Figure 38: Create Drive Array screen

Note: 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 "Advanced 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.

The **Create Drive Array** screen will look similar to the following figure.

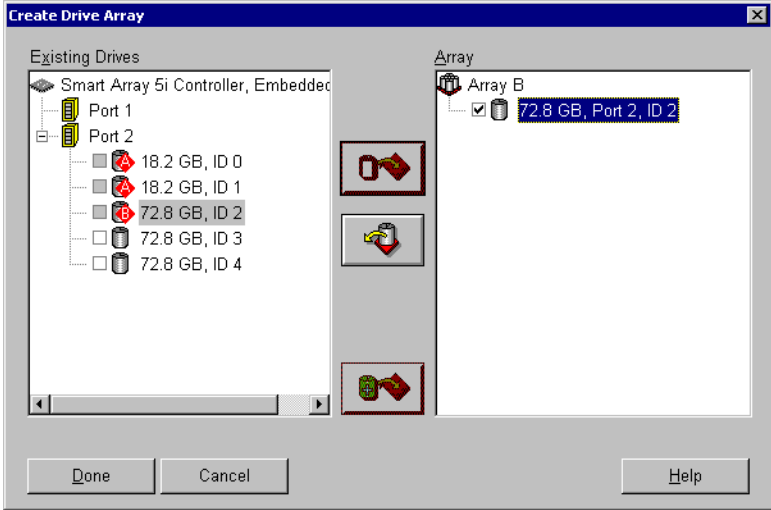


Figure 39: Example Array B

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.

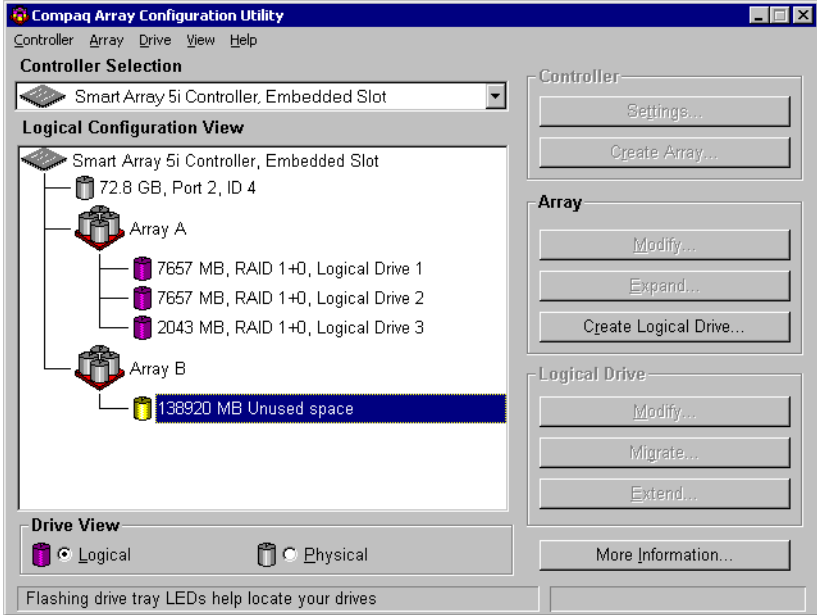


Figure 40: Example Array Logical Configuration view with two arrays



Caution: 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, HP 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.

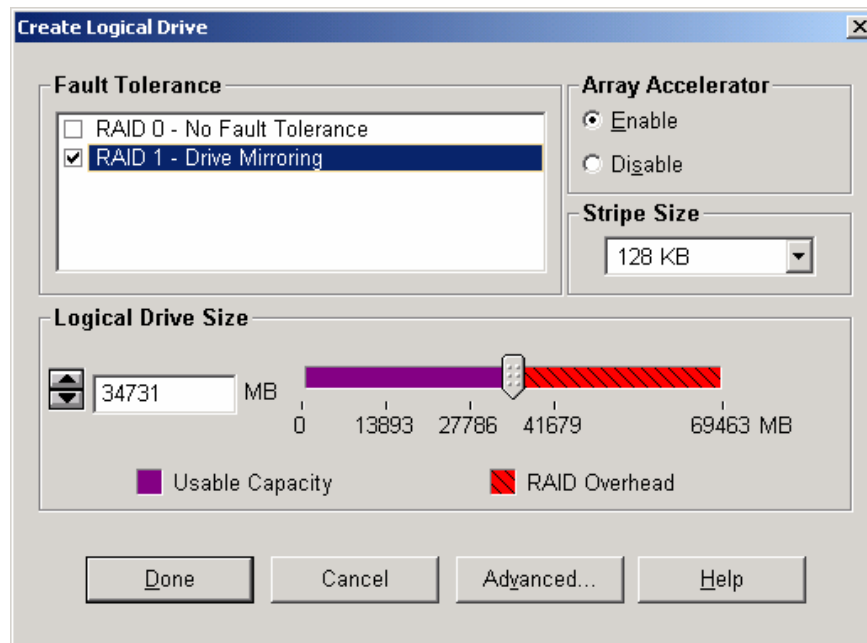


Figure 41: Create Logical Drive dialog box

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 17: 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+0.

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 should look similar to [Figure 42](#).

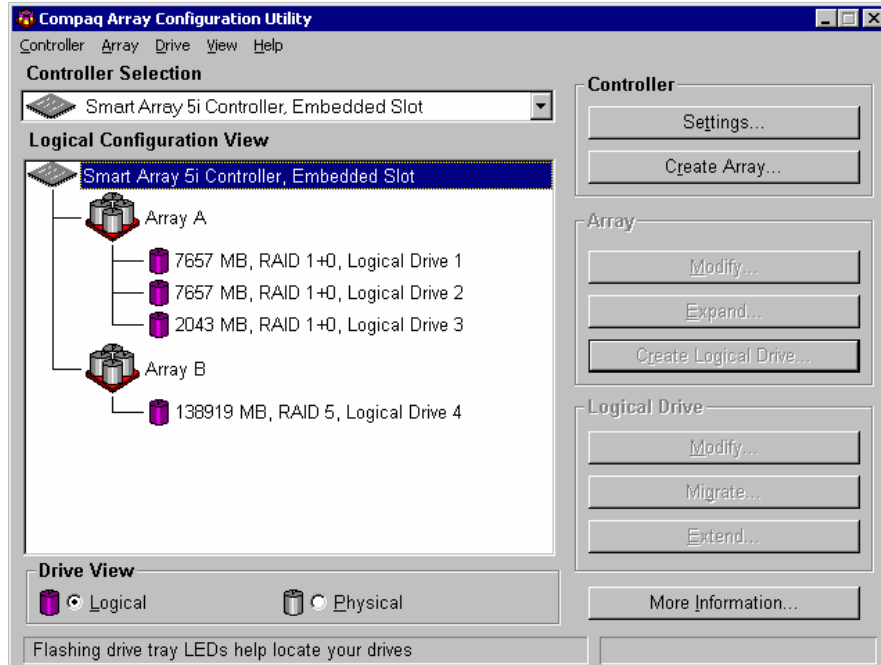


Figure 42: Example array - Configuration View screen with two arrays

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.



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

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.



Caution: 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**.

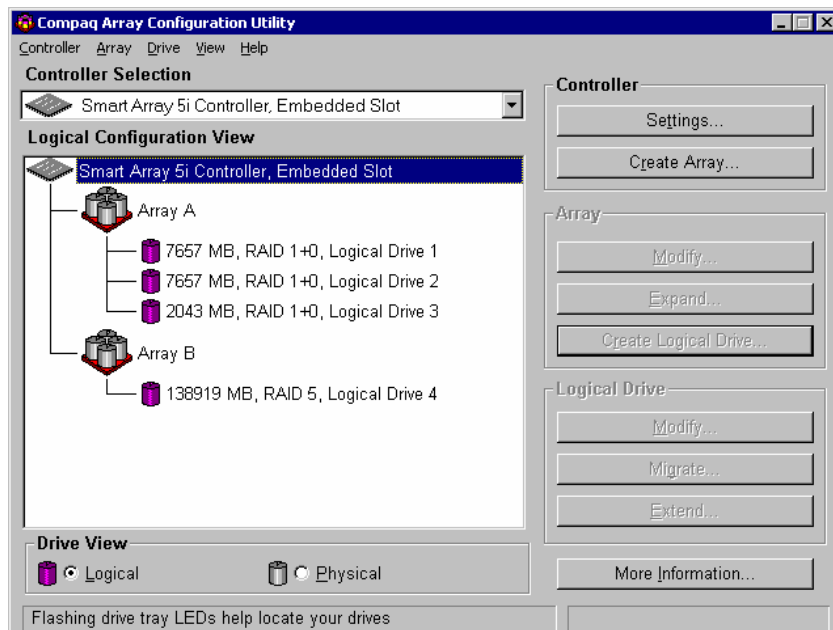


Figure 43: Array expansion example - Logical Configuration View screen

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




Figure 44: Expansion wizards - Logical Drive screen

8. To create a new LUN with the available capacity, click **Create Logical Drive**.
9. 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.
10. Click **Done**.
11. 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.



Caution: 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.

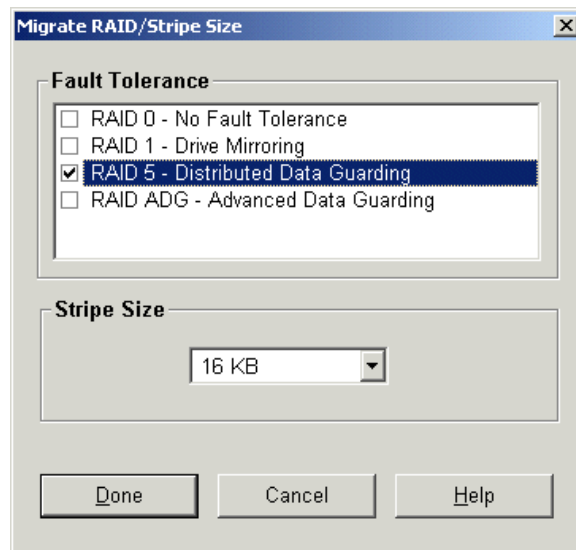


Figure 45: Migrate RAID/Stripe Size screen

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**.

Persistent Storage Manager

6

Persistent Storage Manager lets the administrator make replicas, called snapshots, of disks in a matter of seconds. Snapshots enable the 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 record data changes on volumes and are thus not a replacement for off-line backups.

This chapter covers the following items:

- Operational Overview
- Data Recovery
- Snapshot (Persistent Image) Considerations
- Accessing Persistent Storage Manager

Operational Overview

Each snapshot is a complete point-in-time representation of the data on the volumes. Each snapshot requires only a fraction of the hard-drive capacity of the original data. PSM does not keep all the data that was ever written. PSM maintains only the data required to maintain a snapshot.

PSM works below the operating system as a Filter Driver at the Volume block level. PSM maintains a library of snapshots, each representing a specific point-in-time. Snapshots can be accessed by users, administrators, or any Windows application, and look just like the familiar file/folder view.

With the first snapshot taken on a target volume, PSM establishes a cache file for that volume within which PSM retains overwritten data required to build a snapshot. The cache file size is based on a percentage of the volume it resides on and is configured through the WebUI; the default is 10 percent. As soon as the first snapshot is taken, PSM starts monitoring all writes on the target volume. When a write request occurs, PSM intercepts and pauses the write, reads the data that is to be overwritten, and saves the data in a Diff Directory within the PSM-specific cache file. After the original data is written to the Diff Directory, the new data is written on the active volume. This process is referred to as "copy-on-write." Only the first write forces a copy-out, subsequent writes to the same data block does not force a new copy-out, unless of course a new snapshot is taken between the initial and subsequent write.

PSM can create and manage up to 250 snapshots system wide. A snapshot can cover several volumes at once with an upper limit of 63 volumes within a single snapshot. However, when reverting from a "grouped" snapshot, the revert is non-selective and it reverts all volumes associated with the "grouped" snapshot.

Reading Snapshots

Users who have been granted access by the NAS Administrator see snapshots as network shares.

A snapshot is a representation of the NAS volume at the time it was created. During the copy-on-write operation, the data to be overwritten is preserved in the PSM Diff Directory. When reading a snapshot, PSM determines if the data has changed, meaning it is located in the Diff Directory, or if it is on the live volume. For data that has changed, PSM inserts the original data, held in the Diff Directory and, where no changes have occurred, PSM reads directly from the live volume.

Creating Snapshots

Creation of snapshots is scheduled through the SAK interface or may be generated by the NAS Administrator as a one-time request. When the command to create a snapshot is issued, PSM begins monitoring the file system looking for a quiescent period. A quiescent period is the amount of time a volume must be dormant before a snapshot is created. The default quiescent duration is five seconds but the NAS administrator may configure this, as can the amount of time PSM should search for this inactivity window. The quiescent period provides sufficient time for completion of writes and for the various software buffers to flush, the premise being that, by the end of the quiescent period, a volume will be produced which is in a stable state meaning that the volume is at rest and in a functional condition ready for users to access. If the volume is captured in a stable state, then that volume, or files and folder contained in the volume, will be returnable in a stable state or "useable condition" to users.

Following the quiescent period, PSM creates the snapshot.

PSM Snapshot Attributes

When creating PSM snapshots there are three basic attributes which affect the life and consistency of the snapshot. They are Read-only, Read/write and Always Keep. Read-only should be used to enforce the integrity of a snapshot so that changes can not be made to. Read/Write can be used in instances where test data is useful, such as developers altering a test website. Always Keep is useful when a snapshot needs to live indefinitely. These attributes are described in detail below.

Read Only

The default setting is for PSM to create "READ ONLY" snapshots which prohibits any modification to the snapshot - this is the most common parameter for snapshots. A READ ONLY snapshot allows users, who have been granted access, to view, open and save a copy of any file represent in the snapshot. The properties of a READ ONLY snapshot may be modified by the NAS Administrator to READ/WRITE or ALWAYS KEEP.

Read/Write

The READ/WRITE attribute may be assigned at the time of creation or the NAS Administrator may at any time change the attribute of any snapshot. READ/WRITE snapshots provide some unique capabilities to PSM.

READ ONLY snapshots changed to READ/WRITE snapshots and then modified return the data represented in the snapshot to the way it was originally, effectively acting as an UNDO.

Other applications for READ/WRITE snapshots: CFOs and auditors can run trial balances to accounting systems without affecting the actual systems. Prototyping, a new version of a program, can be installed in a READ/WRITE snapshot and its compatibility within the system tested with no adverse effects to the primary system.

Always Keep

ALWAYS KEEP snapshots are treated as untouchable by PSM. In a cache file fill situation PSM will cease writing to the cache file to avoid deleting or corrupting an ALWAYS KEEP snapshot. A "disk full" error will be returned to the user. ALWAYS KEEP allows the administrator to set some milestones that are not subject to the automatic deletion routines.

Automated Snapshot Deletion

PSM has a snapshot weighting system (low to highest) that helps set the priority of the snapshot. This weighting combined with the age of the snapshot determines the order by which it is deleted by PSM when the cache file fills up.

A key fact to consider is that PSM provides Primary Data Protection automatically. Once set up, PSM continues to provide Data Protection generating new scheduled snapshots or deleting older snapshots with little or no input required from system administrators.

Data Recovery

File/Folder/Volume Recovery

PSM facilitates instant data recovery from the stored on-line images. Individual files, groups of files, folders, groups of folders or complete volumes can be restored. Recovering the data can be accomplished by the NAS Administrator or the NAS Administrator can give individual users access to their data for that purpose through file share access over the network.

Security rights and privileges, as well as file and directory attributes, remain in effect as they were at the time the snapshot was created.

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 volumes are created they are initially contiguous as possible on the underlying storage units (RAID arrays and LUNs). If defrag utilities are used on volumes where snapshots exist, snapshots would grow as the defrag utility moves blocks from one part of the disk to another. PSM disables defrag on volumes that have current running snapshots to prevent the unnatural growth of the snapshot.

PSM (current versions) is fully compatible with the Windows 2000 system file defrag utility. On drives upon which snapshots are not installed or are not active, the defrag utility runs without interruption. If snapshots are active, by design, the drive is automatically marked as unavailable for defragmentation. In operation, the utility works as designed - providing defrag on volumes where it is allowed and omitting drives with active PSM Images. There is no user intervention required. This is consistent with the defragmentation handling of system and special files and is officially supported by the Microsoft defrag API. In the rare case when an existing volume requires defrag, disable scheduled snapshots, delete all snapshots on the volume and defrag the volume. When defrag completes, re-enable scheduled snapshots. Defrag is only effective when there are NO snapshots active on the volume being defragged.

Note: Defragmentation can not be performed if snapshots exist. To defragment a disk, first delete the snapshots. Drive defragmentation only operates on volumes formatted with a 4 KB or smaller allocation size. HP recommends larger allocation cluster sizes to improve performance.

PSM and Backup

Because snapshots are quick to create, it is possible to capture a coherent view of the volume 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 volume. 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 cache space becomes critical, snapshots are automatically deleted. See the "Automated Snapshot Deletion" section. In addition, snapshots are a short term convenience and may 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.

Although snapshots provide a mechanism for backup that does not require downtime, there are some considerations that should be given when performing backup and restore of a system using snapshots. HP recommends you review this section prior to establishing backup and restore policies. Backup and Restore programs are not trivial applications. As such they require effort to set up and use effectively. Given the nature of these products, it is critical that any backup and recovery plan be thoroughly tested before use on a live system.

Be sure to use a backup program that is PSM aware and has been certified for operation with PSM. This is especially true for open file options, system agents, and disaster recovery.

For backup:

- For base volumes that have snapshots in use or when backing up snapshots, archive bit resets and incremental backups should not be used. Archive bit resets are recorded as a change to the data and can fill the cache file with changes. Incremental backups make use of the archive bit set as well. Note if the snapshot is set to read only the backup will also fail.
- Be careful in the selection of folders, since snapshot folders provide a view into the data that can result in the backup of multiple views of the data. Forcing the backup to grow based on the number of snapshots in use.
- Junction points should be turned off to prevent the traversal of multiple snapshot directories of base volume backups.
- Junction points should be turned on when backing up a single snapshot. Be sure to pick the single snapshot and not the root folder. Selecting the root folder will cause multiple snapshot backups.

For restore:

- Delete all active snapshots as the restore will cause the cache file to grow.
- Select only the files representing the data of the volume and not the *.psm files.
- Be sure to restore to the root of the target volume.
- Restoration of operating system partitions does not restore the registry hive. System state backups should be utilized in these instances.

Snapshots Performance Impact

When using snapshots, performance of the disk may be affected, depending on the rate that data is changing and the number of snapshots kept for each disk. Read performance of the disk remains constant, regardless of the presence of snapshots. Read performance of the snapshot is identical to that of the disk. Write performance, however, may vary. PSM creates minimal additional I/O overhead which is limited to writes. The copy-on-write process adds one read (the write is paused to read the old data) and one write (the old data is written to the Diff Directory file) to each write system request. This only affects each initial write to a disk area that has a snapshot running on it. Copy out is not performed on subsequent writes to the same disk block, so write performance is unaffected after the initial write to each block.

Predicting the exact effect of snapshots on any particular 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 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. Disks with high write locality experience less performance degradation due to snapshots.

Recovering Snapshots after a System Restore or System Loss

The b2000 ships with a Quick Restore CD for circumstances that require a server rebuild. During the system restore or in the event of a complete system loss, registry information is lost with regard to the snapshots that were instantiated prior to system restore. Volume data will remain unaltered, only the snapshots will become invalidated. Even though all snapshot folders and cache files exist on the system volumes, the snapshots are not picked up by PSM and are orphaned. These files will need to be cleaned up. To delete the cache files and snapshot directories please see the section on "Clearing the Cache Files" from the system later in this chapter.

Granule Size Update Utility

PSM ships with a utility for adjusting the Granule size of the snapshots. Granules determine the largest cache size that can be managed by PSM. The default setting in the PSM product that ships on the b2000 is 64 K. This setting will allow for up to 1 TB of data to be written to the cache file. In order to gain greater cache file space, the granule size will need to be adjusted. The following table provides an overview of the addressable storage space and maximum cache size of each Granular size.

Table 18: Adjusting Granule Size

Granule Size	Largest Cache Size
64K granule	1 TB
128K granule	2 TB
256K granule	4 TB

When considering the granule size the following rules should be observed.

- Before altering the granule size, all snapshots should be removed from the target system.
- Cache File size is fixed as in the above table and the limit applies to the sum total of all cache files system wide.
- Granule size affects only the block size utilized for each change that is written to the cache files. Regardless of the setting, there is approximately 15.6 million blocks available for storing snapshot information system wide other system limitations may further limit this maximum such as memory consumption.
- If the changes occur in different underlying blocks, more blocks of larger space could get written for any set of changes, versus if the changes all occur in the same block. Therefore increased granule size does not necessarily lead to increased coverage for changes on the originating volumes. In theory, larger blocks should lead to fewer blocks consumed to record the original data due to write locality.
- Highly fragmented disk space could lead to increased separate cache writes and more consumption of the maximum available number of blocks system wide.
- Setting the value too low will limit the available space for cache file writes. For example, a 10 TB system undergoing change could only experience a 10 % change in original data if the granule size is set to 64 KB, assuming all of the changes fit neatly into the 64KB blocks.
- PSM now supports the PSM granule sizes of 64K, 128K, 256K with 64K as the default. This will allow for cache file to be 1TB, 2TB, and 4TB respectively. The program GRANSIZE.EXE, available in the directory c:\winnt\system32\serverappliance, is provided for setup - By increasing the granule size, PSM can be better suited to support very large terabyte systems. The command provides an error message if there are running snapshots on the system. Typing GRANSIZE ? will display the current granule size in use in the system. Typing just GRANSIZE will display the command usage. The command must be executed from a command prompt while residing in the directory stated above.

- When changing to a larger granule for systems, thus allowing for larger cache file sizes and accommodating larger amounts of storage, users should lower their percentage of volume space for the cache file. For example, if the percent is 30 and the supported amount of space in the system is 20 TB, then the cache file limit of 4 TB would get exceeded. Should the limit get exceed, PSM will issue an "Out of Memory" error in the event log and the WebUI status page. If the limit is exceeded, the cache file must be removed or reduced in size prior to system restart using either the `clearvol` command or by reducing the percent cache size under volume settings.

Clearing the Cache File from the System

The PSM interface allows the user to set the cache file to any percentage from 1 - 70 percent but it will not allow the deletion of the cache file in its entirety. It is possible to delete these files but the process must be done from the command prompt either through Terminal Services or from the NAS console. To delete the PSM cache files and cache directories the following command: `CleanVol.exe Vol:` must be performed for each existing volume where the cache file is no longer desired. The command may be found in `c:\winnt\system32\serverappliance`. Typing `cleanvol` will display the command usage. Prior to these steps the snapshots on the target volume need to be deleted as well or "access denied" error will be returned.

Re-extending Volumes from Old Snapshots

Volumes based on dynamic disks may be extended utilizing LDM. Corresponding snapshots can exist at points in time prior to the extension and after the extension. If a re-extended volume containing snapshots of the pre-extended volume is reverted, the re-extended area of the disk will be unusable. To reclaim this space, make sure the included utility `reextend.exe` is executed after reverting from a snapshot of the pre-extended volume.

This utility is available in the directory `c:\winnt\system32\serverappliance` and must be executed either through terminal services or at the NAS console.

Usage of this utility is available by typing `reextend -?`.

This program will extend a volume back to its original size after a restore operation of a smaller volume from a snapshot.

Volume Display in Persistent Storage Manager

PSM fully supports the use of all Logical Disk Manager storage elements this includes basic, dynamic, partitions, extended partitions, and volumes provided they are formatted as NTFS when created. PSM makes use of two items when displaying storage elements in the UI. These include the volume label and the GUID representing that volume or partition. In several web pages, the information displayed is limited with regard to the identification information and the volume label is essentially all that can be viewed. It is therefore important that volume labels be identifiable by the user to avoid confusing one volume over another. By default, Local Volume, followed by the drive letter is displayed, for mount points the GUID is displayed. This label should be updated to reflect a unique label either during volume/partition creation in LDM or post volume/partition creation via File Explorer and the properties tab of the target drive.

Persistent Storage Manager Storage Limitations

The version of PSM included in the b2000 is currently designed to work with 10 TB of storage with the ability to take 250 snapshots. However, the b2000 is capable of addressing 27 TB using the fully populated system and 146 GB drives. PSM will continue to function with larger systems but the snapshot coverage should only encompass, 10 TBs worth of storage. There are no safeguards to prevent the use of storage greater than 10 TB. HP is currently working with CDP to address this storage limitation. Please check the HP website for updates regarding this support.

Accessing Persistent Storage Manager

To access PSM, from the **WebUI Welcome** screen, select **Disks**, then **Persistent Storage Manager**.

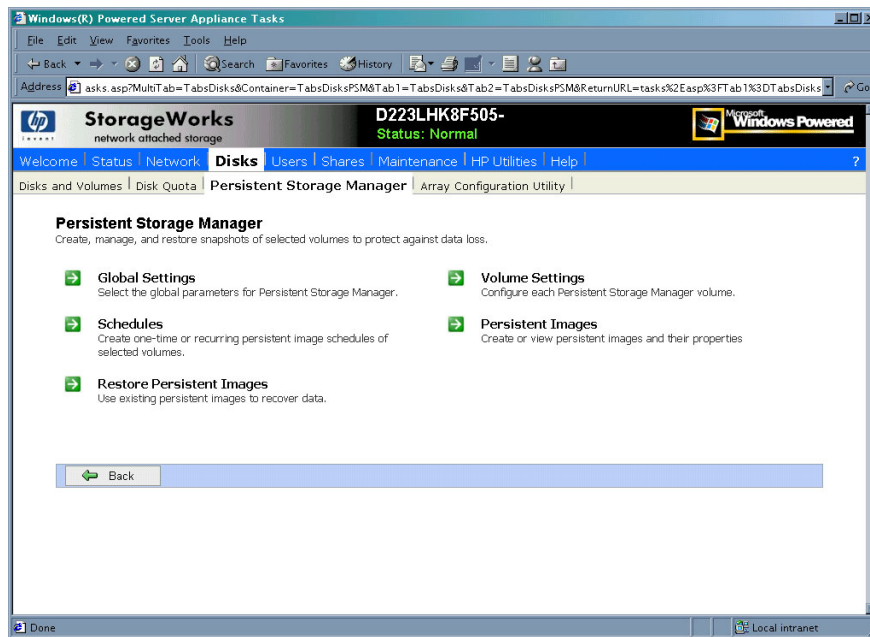


Figure 46: PSM Main screen

Global Settings

From the **Global Settings** screen you can control the overall environmental settings for Persistent Storage Manager. Some options will be disabled if there are already active snapshots.

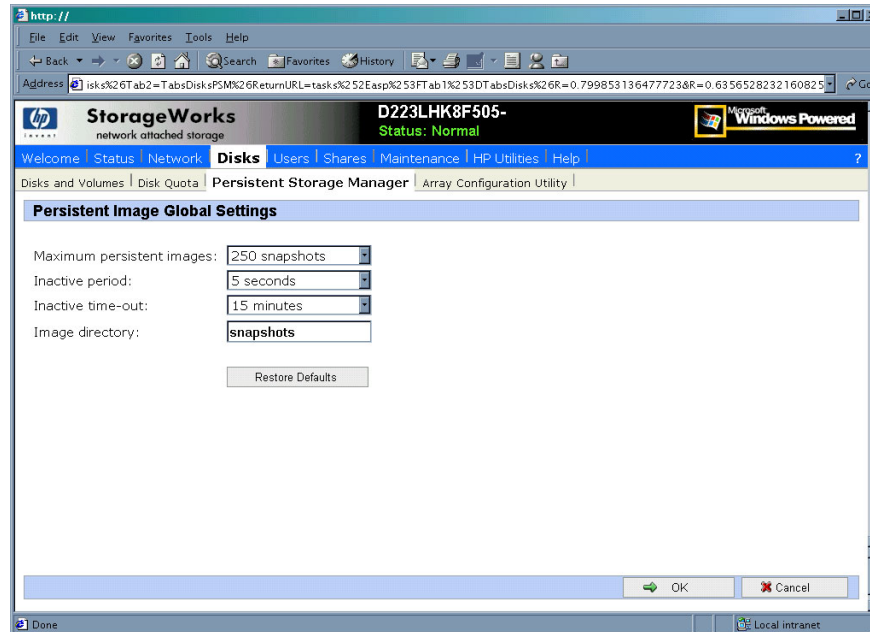


Figure 47: Global settings

Maximum Persistent Images

This option determines the maximum number of active Persistent Images (snapshots). PSM will support a maximum of 250 snapshots per server. The size of the cache file will determine the actual amount each server can hold.

If the creation of a new snapshot would cause the maximum number to be exceeded, the system will delete the oldest existing persistent image according to the deletion heuristics established by the user.

Inactive Period

This option specifies the amount of time a volume must be dormant before a snapshot is created. Before starting a snapshot, the system will wait for the volume being imaged to become inactive. The default value will allow systems to start an image with a consistent file set and a minimal time-out. Administrators can change this value for system optimization. Reducing the inactive period will allow you to create snapshots even on busy systems, but with possible synchronization problems within applications which are concurrently writing to multiple files.

Inactive time-out

This option specifies how long the server should try to create a snapshot. A snapshot will not begin until a period of relative inactivity set by the Inactive period has passed. If an interval passes that is longer than the Inactive time-out period, the snapshot will not be created and a notice generated to the system event log.

Image directory

This option specifies the root directory used for the snapshot. Each snapshot appears as a subdirectory of the volume that is being imaged. The entire content of the volume as it existed at the moment the snapshot was created will appear under this directory.

Restore Defaults

The **Restore Defaults** button will reset the system defaults.

Volume Settings

From the PSM screen select **Volume Settings**. From the **Volume Settings** screen you can view the Persistent Storage Manager attributes for each volume and change volume settings using the **Configure** button in the **Tasks** list.

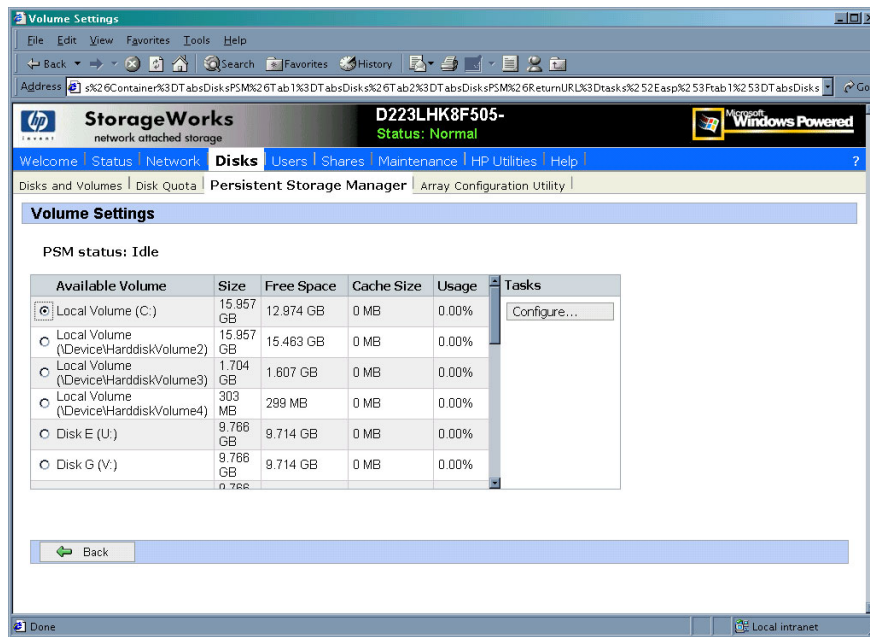


Figure 48: Volume settings

Available Volume

This field lists all of the volumes that can support snapshots. You can select the volume you want to configure.

Size

This column displays the size of the volume.

Free Space

This column displays the available storage size of the volume.

Cache Size

This column specifies the amount of space allocated to the cache file. Increasing this value will allow more and larger snapshots to be maintained.

Usage

This column displays the current cache file use as a percentage of the cache size.

Volume Configuration Settings

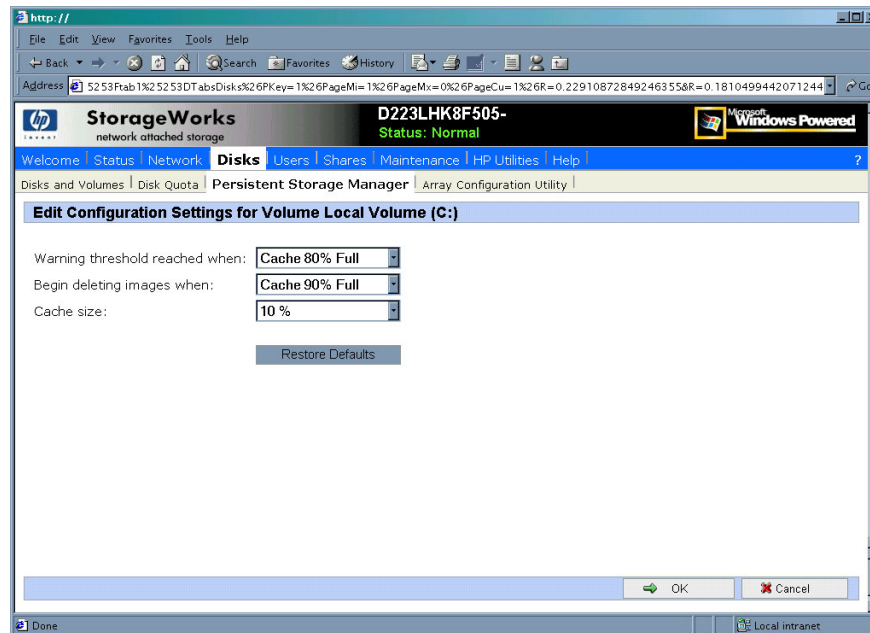


Figure 49: Volume configuration settings

Click **Configure** from the Volume Settings to modify the various aspects of the PSM volume attributes. Some of the fields will appear read-only if there are active snapshots. The **Restore Defaults** button will re-establish the system defaults. If it is desired to remove the cache files all together the *CLEANVOL.EXE* can be used to remove them; see the section on clearing the cache file. Also note the section on granular size in this chapter, prior to updating the percent reserved for cache size. The default value is 10 percent.

Note: Changing the values for the cache size can result in cache files that exceed the maximum cache file based on the current granule size. If the limit is exceeded "out of memory" notices appear in the event log and the WebUI status page when the first snapshot utilizing that cache file is taken. The snapshot will fail to create but the cache file is built regardless. It is important to reduce the cache file size via the above screen or clean the cache files prior to the restart of the NAS system if an oversized cache file is created.

Warning threshold reached when

This option defines the percentage of cache space which, when consumed, will trigger warning messages to the system event log.

Begin deleting images when

This option defines the percentage of cache space which, when consumed, will trigger the automatic deletion of the oldest snapshot on the system. Automatic snapshot deletions are recorded in the system log.

Cache size

This option specifies the amount of space allocated to the cache file. Increasing this value will allow more and larger snapshots to be maintained. Make sure that adequate space is available on the drive where snapshots are stored.

Schedules

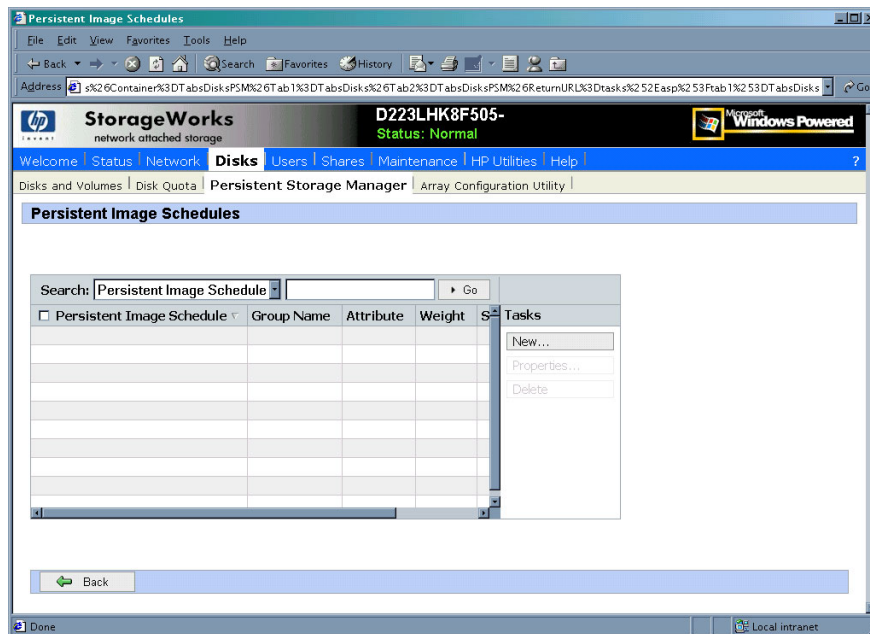


Figure 50: Persistent Image Schedules

The **Persistent Storage Manager Schedules** page displays a list of scheduled snapshots and associated tasks.

Each scheduled snapshot contains information such as its scheduled time, day, frequency, starting date, and group name.

Schedules screen allows you to create new schedules, delete existing schedules, and edit schedule properties.

Create a New Schedule

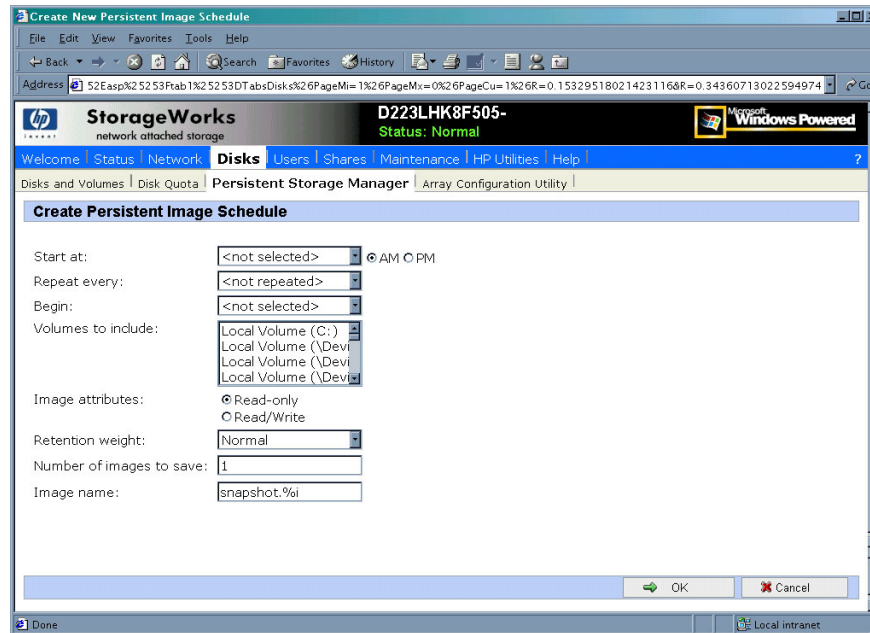


Figure 51: Create Persistent Image Schedule

To create a new schedule, you must supply a starting time, repeat period, starting day, volume, and the number of snapshots to make available to users.

To add a snapshot to the schedule:

1. Select **Schedules** from the **PSM Main** screen.
2. In the **Tasks** list, select **New**.
3. Select the parameters you want for the schedule.
4. Click **OK**.

Editing Persistent Image Schedule Properties

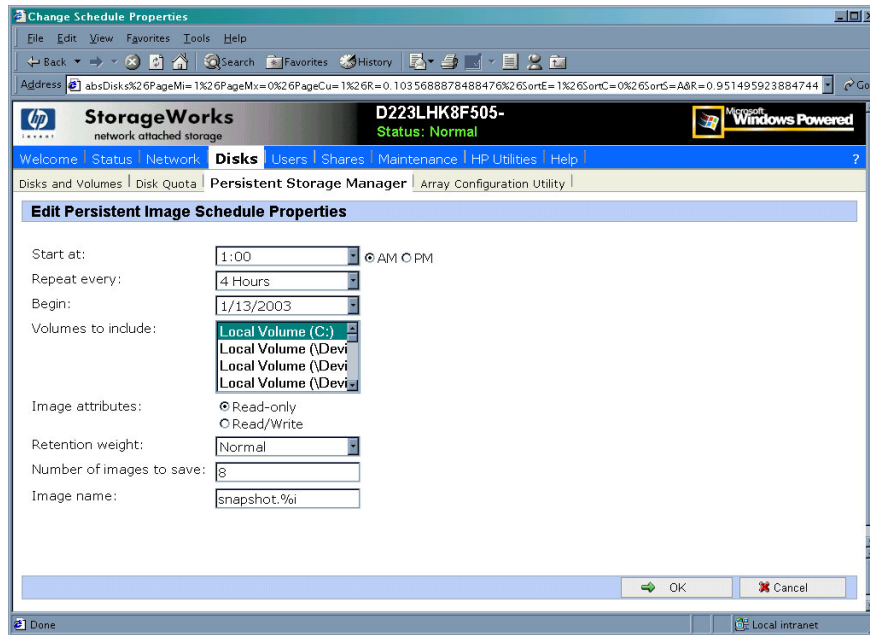


Figure 52: Edit schedule properties

To edit persistent image schedule properties:

1. Select **Schedules** from the **PSM Main** screen.
2. In the **Tasks** list, select **Properties**.
3. Select the changes you want to make to the schedule.
4. Click **OK**.

Deleting a Persistent Image Schedule

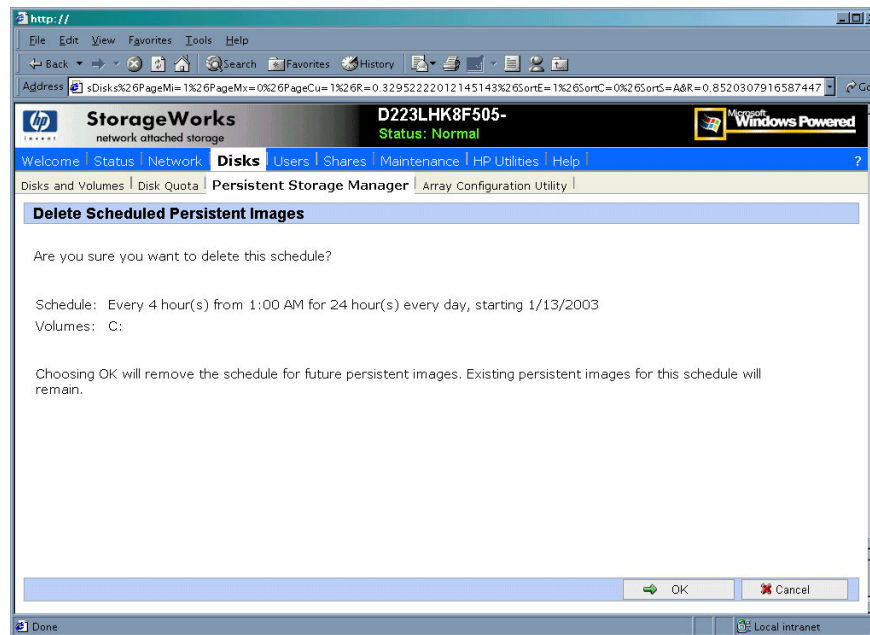


Figure 53: Delete scheduled images

To delete a persistent image schedule:

1. Select **Schedules** from the **PSM Main** screen.
2. Select the schedule you want to delete.
3. In the **Tasks** list, select **Delete**.
4. Click **OK**.

Persistent Image and Group Information

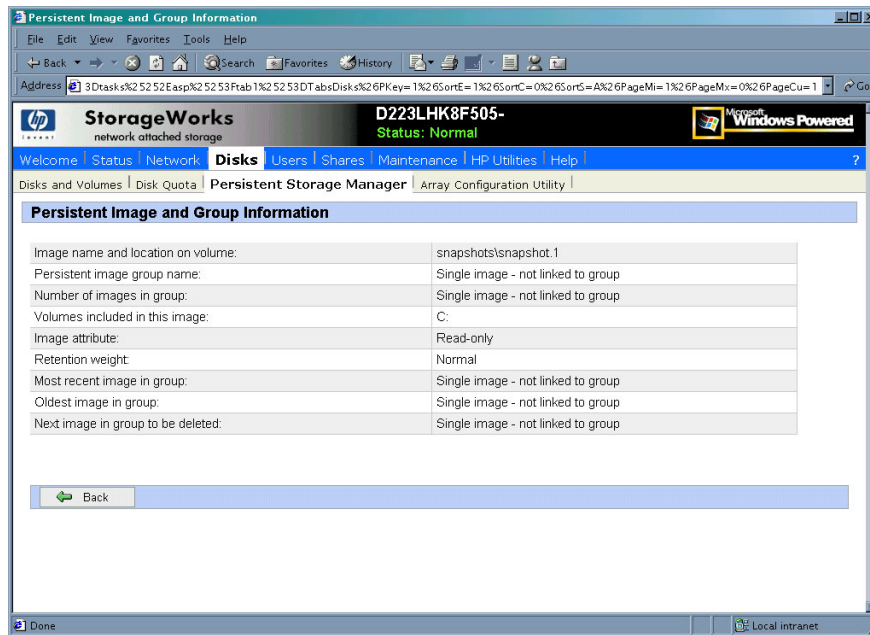


Figure 54: Persistent Image and Group Information

After a snapshot is created from the schedule you specify, it becomes a member of an image group. The **Persistent Image and Group Information** page can be accessed by selecting the desired snapshot and clicking **Details** on the **Persistent Images to Restore** screen. The screen displays the following information about the image group:

Image name and location on volume

This field displays the name of the image and its path.

Persistent image group name

This field displays the name assigned to this group.

Number of images in group

This field displays the maximum number of images that can be included in the group.

Volumes included in this image

This field displays each volume included in the image.

Image attributes

This field displays the read-only or read/write attribute of the image.

Retention weight

This field displays the relative retention weight of the image.

Most recent image in group

This field displays the date and time of the image most recently added to the group.

Oldest image in group

This field displays the chronologically oldest image in the group.

Next image in group to be deleted

This field displays the date and time of the image that will be deleted next so the system can stay within the saved images limit.

Managing Persistent Images

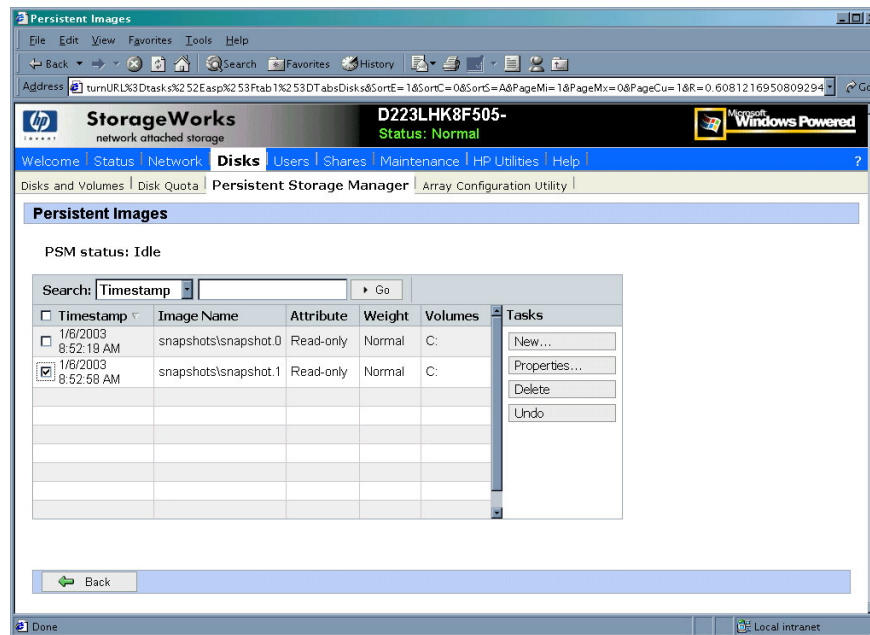


Figure 55: Managing persistent images

The **Persistent Images** page displays active persistent images. Each entry identifies the date and time the snapshot was created, the read-only or read/write attribute, the preservation weight, and the volume it preserves.

To manage snapshots:

1. From the **PSM Main** screen select **Persistent Images**.
2. Select the snapshot you want.

3. Choose one of the following tasks:
 - a. Choose **New** to create a new snapshot.
 - b. Choose **Properties** to view or change the image read/write attribute or retention weight.
 - c. Choose **Delete** to delete the image from the system.
 - d. Choose **Undo** to undo changes to a read/write image.

Creating a New Persistent Image

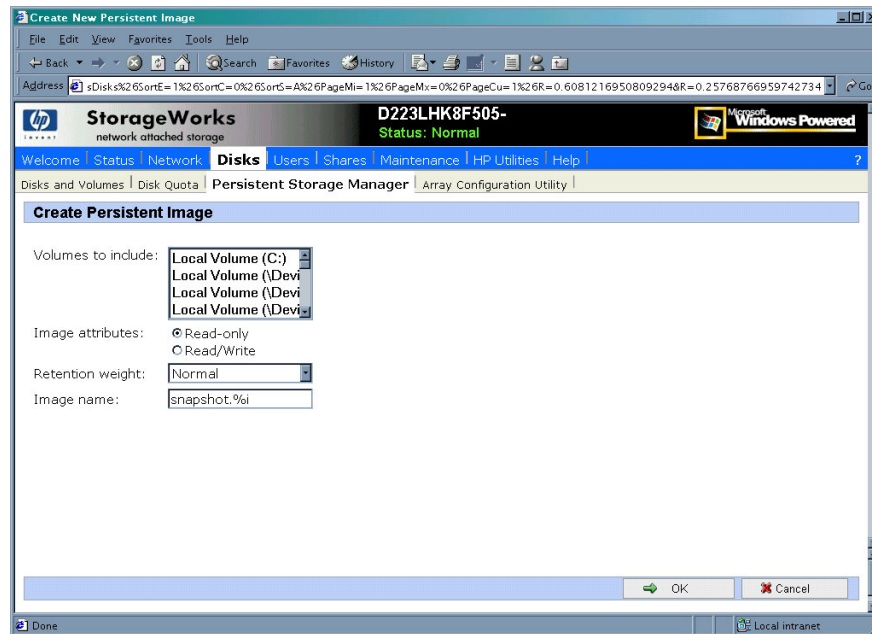


Figure 56: Create new persistent image

Snapshots may be created directly through the **Persistent Images** page. You can also use the **Schedules** page to schedule future or recurring snapshots. To create a new snapshot:

1. From the **PSM Main** screen select **Persistent Images**.
2. In the **Tasks** list, choose **New**.
3. In the **Volumes to include** list, choose volumes to be included in the image.
4. Select the **Read-only** or **Read/Write** button.
5. Select a retention weight from the **Retention weight** list.
6. Type the image name in the Image name box.
7. Choose **OK**.

Deleting a Persistent Image

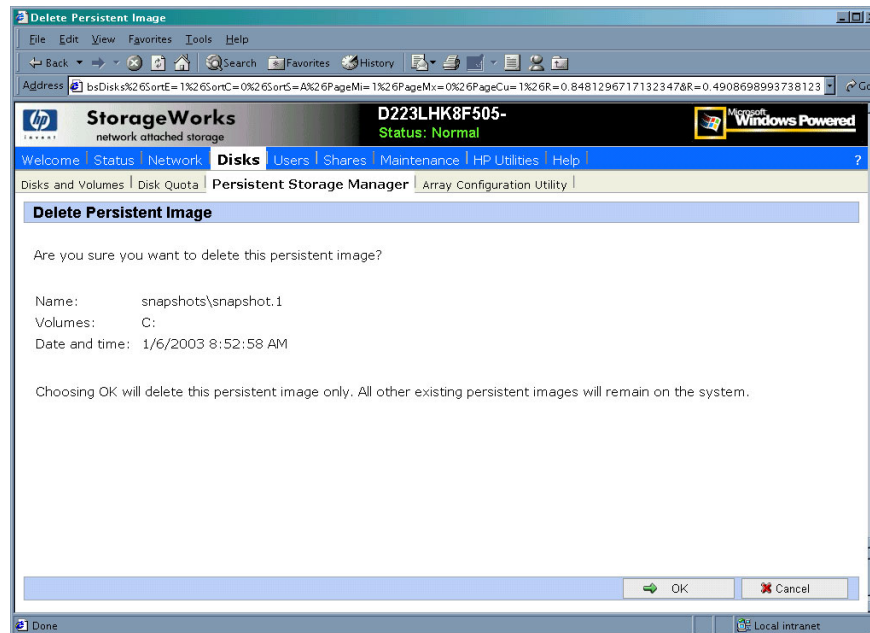


Figure 57: Delete Verification

To delete a persistent image:

1. From the **PSM Main** screen select **Persistent Images**.
2. Select the snapshot you want to delete.
3. In the **Tasks** list, choose **Delete**.
4. Choose **OK**.

Editing Persistent Image Properties

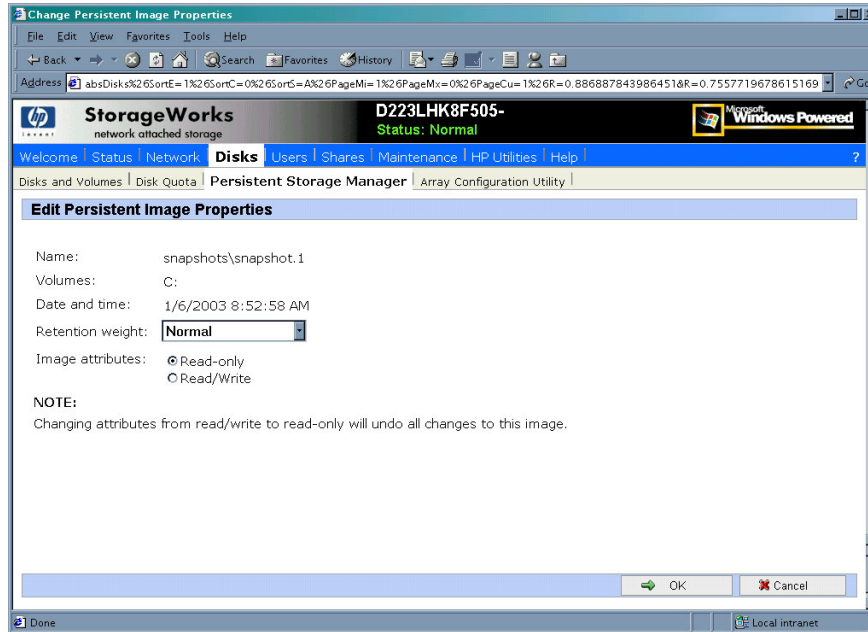


Figure 58: Edit Persistent Image Properties

You can change properties such as the read-only attribute or preservation weight of an image.

To edit persistent image properties:

1. From the **PSM Main** screen select **Persistent Images**.
2. In the **Tasks** list, choose **Properties**.
3. Select a retention weight from the **Retention weight** list.
4. Select the **Read-only** or **Read/Write** button.
5. Choose **OK**.

Undo Persistent Image Changes

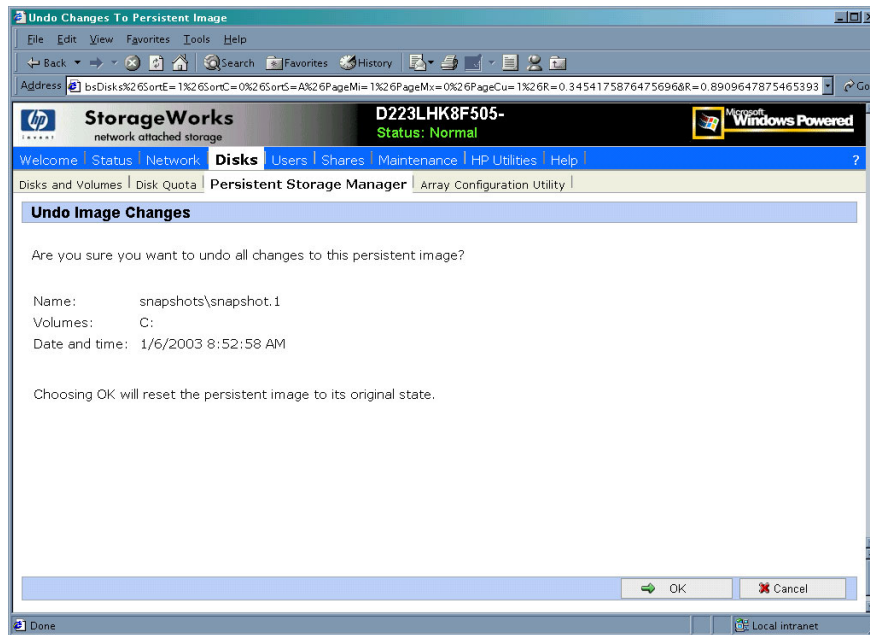


Figure 59: Undo Image Changes

After you create a read/write snapshot, you can make changes to the image, for example, you can modify files in the image, add new files, or delete existing files. If you make a change to an existing image and later want to revert to the original file contents, you can use the following procedure to restore the original snapshot.

To undo snapshot changes:

1. From the **PSM Main** screen select **Persistent Images**.
2. Select the snapshot you want to restore to its original state.
3. In the **Tasks** list, choose **Undo**.
4. Choose **OK**.

Restoring an Image

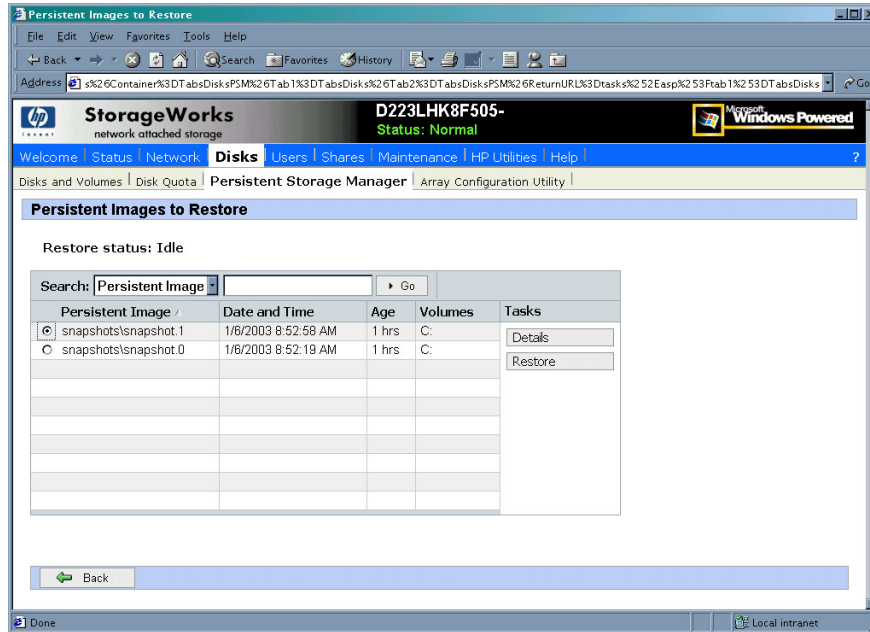


Figure 60: Images available to restore

The **Persistent Images to Restore** page displays a list of all snapshots. You can choose to view an image or restore your server appliance to an image you have previously created.

To restore a snapshot:

1. On the PSM Main screen select **Restore Persistent Images**.
2. Select the snapshot you want to restore.
3. Choose **Restore**.

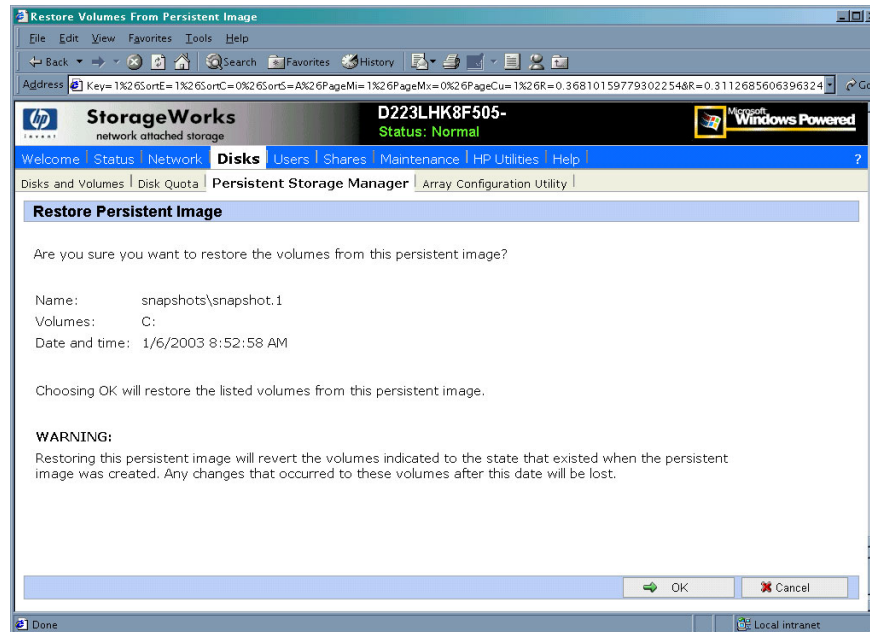


Figure 61: Restore confirmation screen

4. After selecting **Restore** the **Are you sure** screen will appear.
5. Choose **OK**.

Note: PSM will not allow the restoration of the system partition from a snapshot. No error is issued, it simply will not revert the volume. PSM protects the system partition against the revert operation, since it would potentially lead the operating system in an inconsistent state.

Known Issues

These were the known issues at time of publication. Please refer to the release notes for the b2000 for updated information regarding known issues.

Event log error at cache full

The eventlog error a driver below this one has failed in some way may occur when the cached file is full.

Display Error on SAK

Status events not rendered properly on SAK. The percent signs not displayed value substitutions missing in displayed message

Always Keep error at cache file full

If all your snapshots on C:\ are tagged as **Always Keep** and the cache file fills up, the system may experience a BSOD at reboot.

It is not recommended to flag all snapshots as **Always Keep** because this disallows the PSM deletion logic to delete the older snapshots to free up cache file space.

Improper display of default Cache File Size

You must delete all snapshots before changing the cache size.

Page file setting

The **Page file size** must not change and the initial size must be set equal to the maximum size. This setting is located in the **Virtual Memory** settings under **System Properties**.

No Boot - No Revert

If the system cannot boot, a revert operation cannot be performed.

Reverting of System Drive Prohibited

PSM does not allow the ability to revert the system boot drive.

No support for mount points in UNIX, AppleTalk, or NetWare

Microsoft confirmed that the Microsoft NFS Services for UNIX, Services for Macintosh, and Services for NetWare do not support volume mount points. These clients will not be able to access data on volumes mounted using a volume mount point. Since snapshots for a volume are mounted as directory junctions (AKA mount points), and even though they are shared these clients will not be able to access the snapshots.

Please refer to the Microsoft Release Notes for Microsoft Server Appliance Kit dated June 2001.

User and Group Management

7

The HP 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
 - Managing Local Groups
- Drive Quotas
 - Managing quotas
 - Enabling and disabling quota management
 - Creating new quota entries for a user or group
 - Deleting new quota entries for a user or group
 - Modifying new quota entries for a user or group

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:

<http://www.microsoft.com/windows2000/technologies/directory/ad/default.asp>

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

Username 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 19](#) provides examples of group names.

Table 19: Group Name Examples

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 two choices:

- 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 Local Users

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.

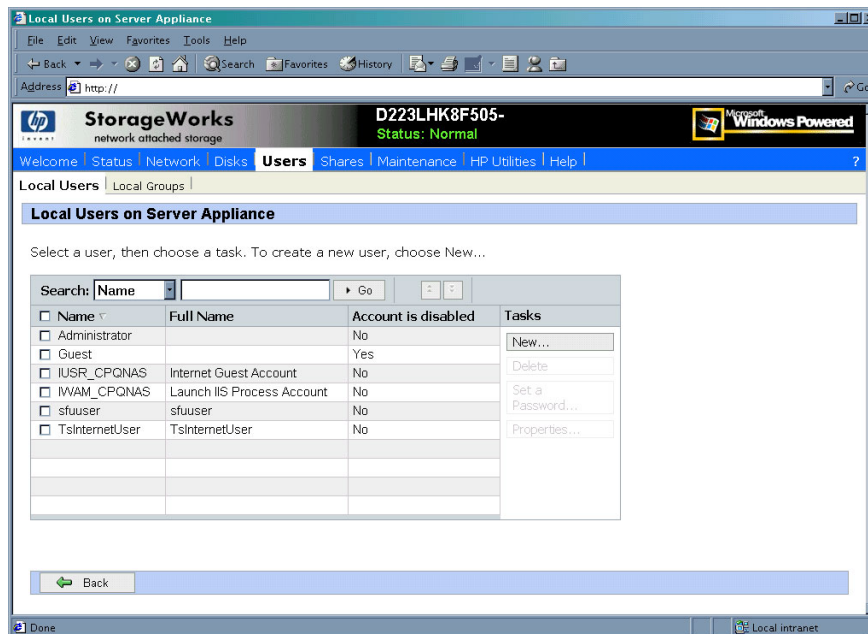


Figure 62: Local Users dialog box

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.

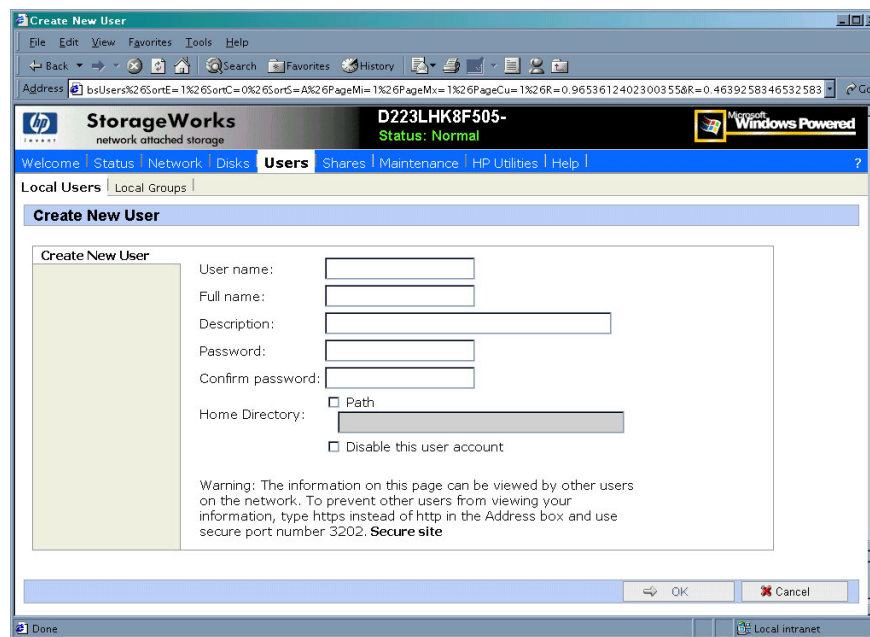


Figure 63: Create New User dialog box

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 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 64](#) is an illustration of the **User Properties** dialog box.

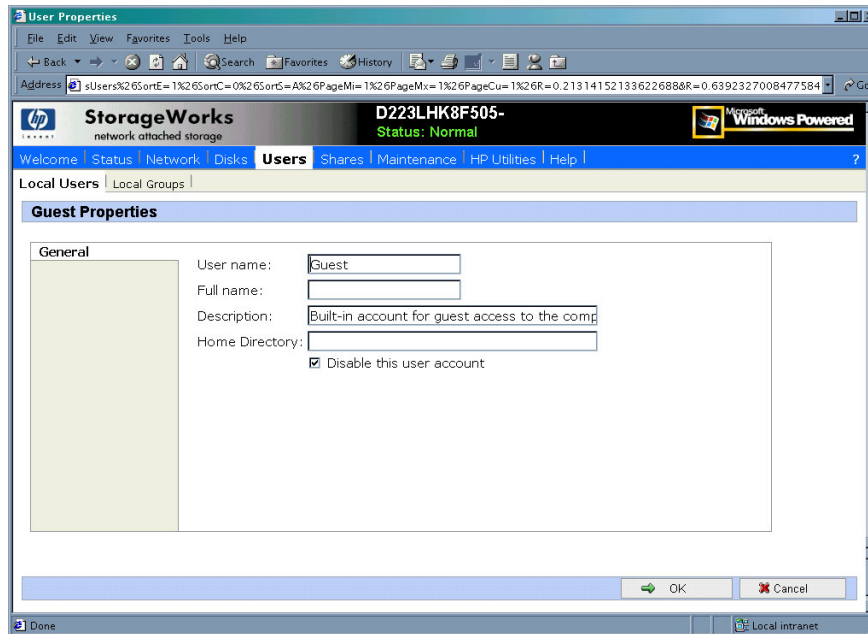


Figure 64: User Properties dialog box

2. The following information can be changed or set:
 - User name
 - Full name
 - Description
 - Home Directory
 - Disable this user account
3. After completing the changes, click **OK**. The **Local Users** dialog box is displayed again.

Managing Local Groups

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.

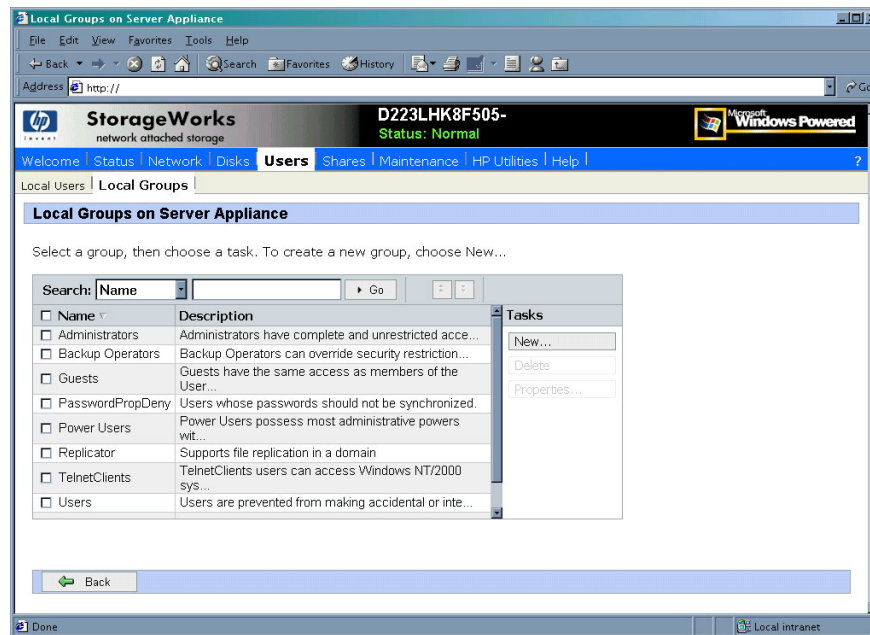


Figure 65: Local Groups dialog box

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.

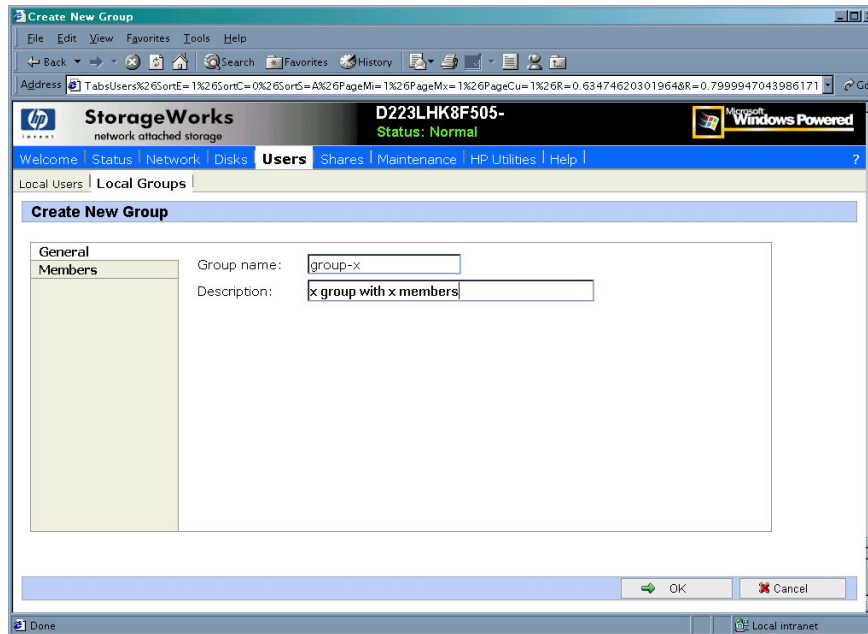


Figure 66: Create New Group dialog box, General tab

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.

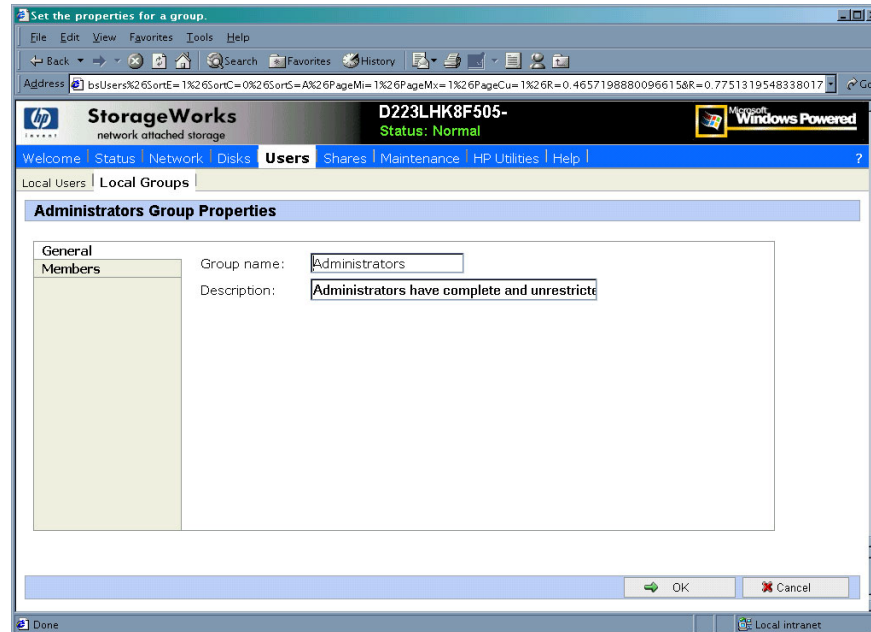


Figure 67: Group Properties dialog box, General tab

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 user or group** 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 68 is an example of the **Members** tab.

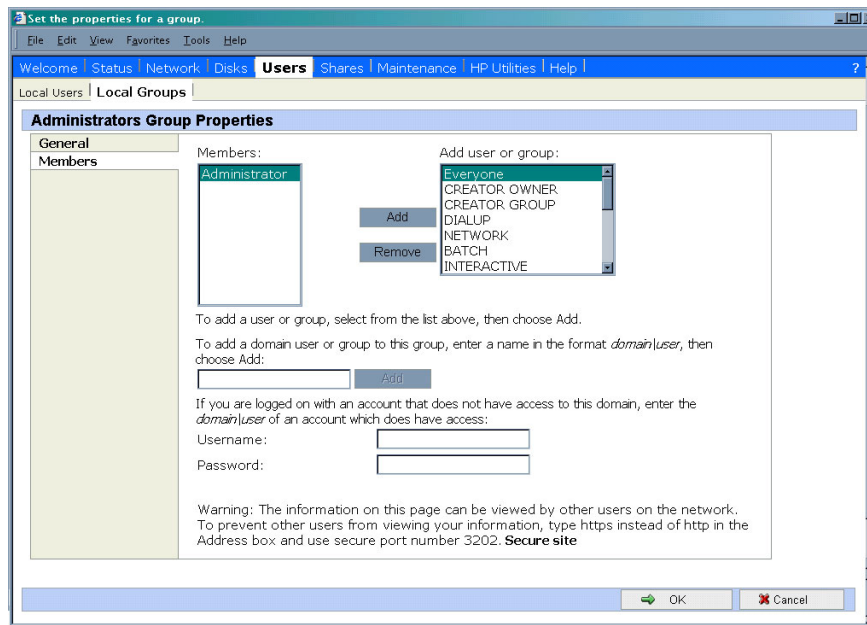


Figure 68: Group Properties dialog box, Members tab

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.

Managing Quotas

Managing quotas includes:

- Enabling and disabling quota management
- 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 69 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.

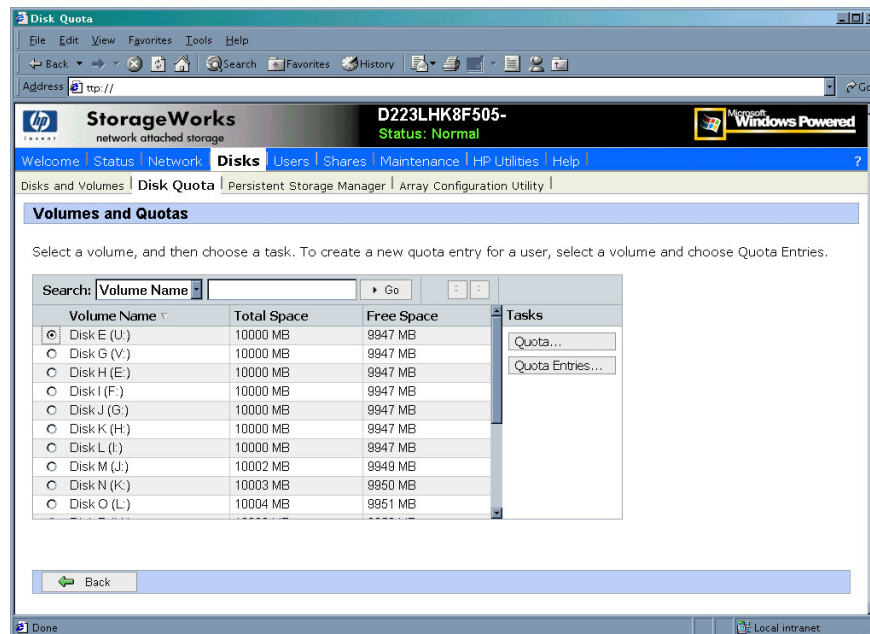


Figure 69: Disk Quota dialog box

Enabling and Disabling Quota Management

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.

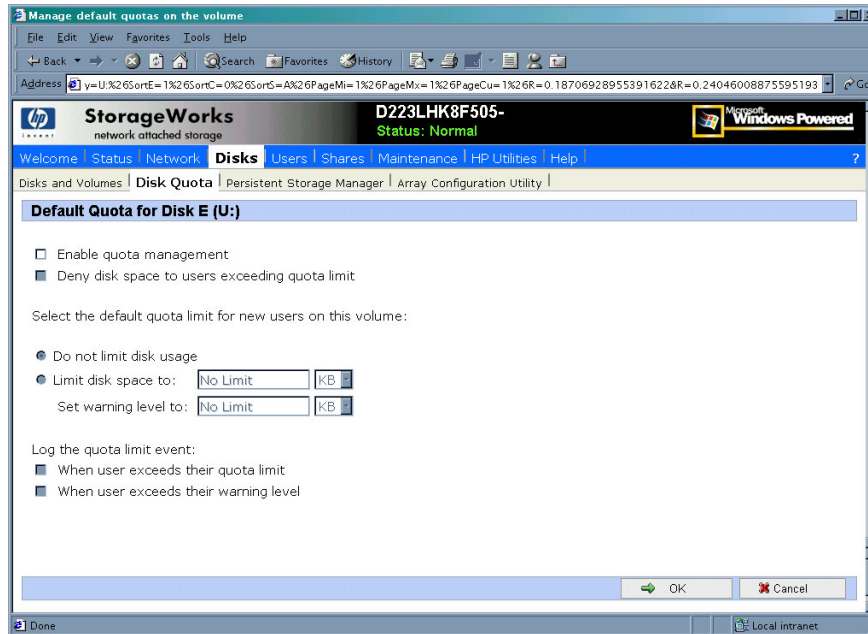


Figure 70: Default Quota dialog box

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.

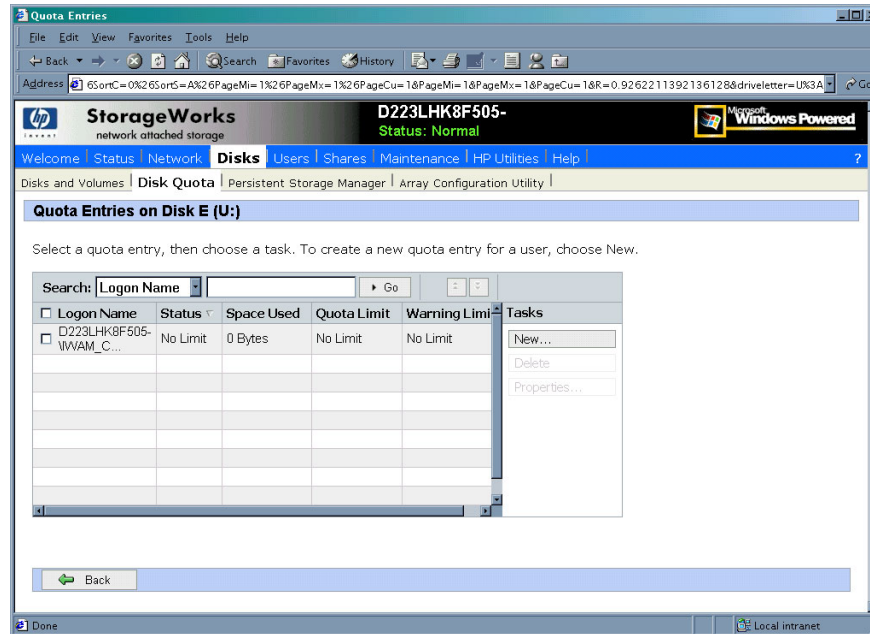


Figure 71: Quota Entries dialog box

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.

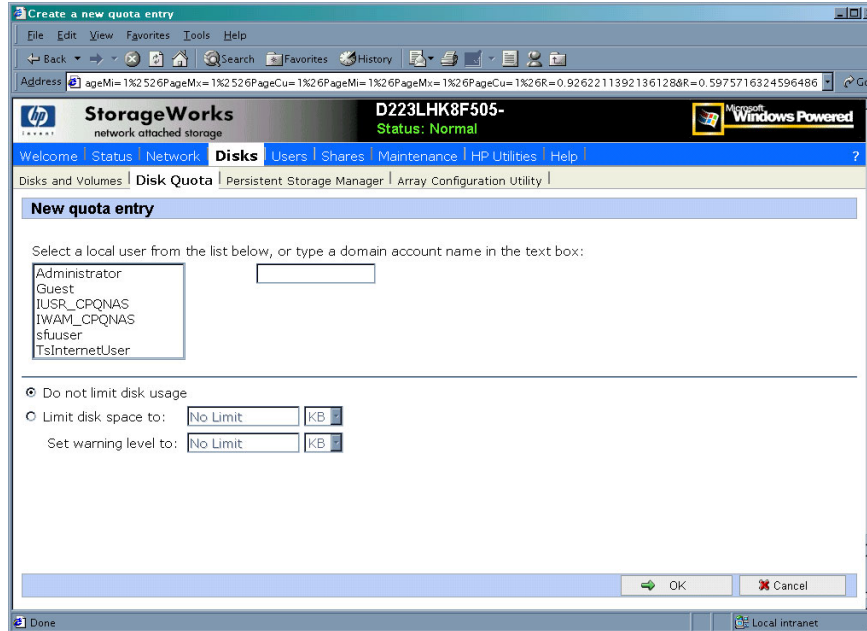


Figure 72: New Quota Entry dialog box

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 user quota settings:

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.

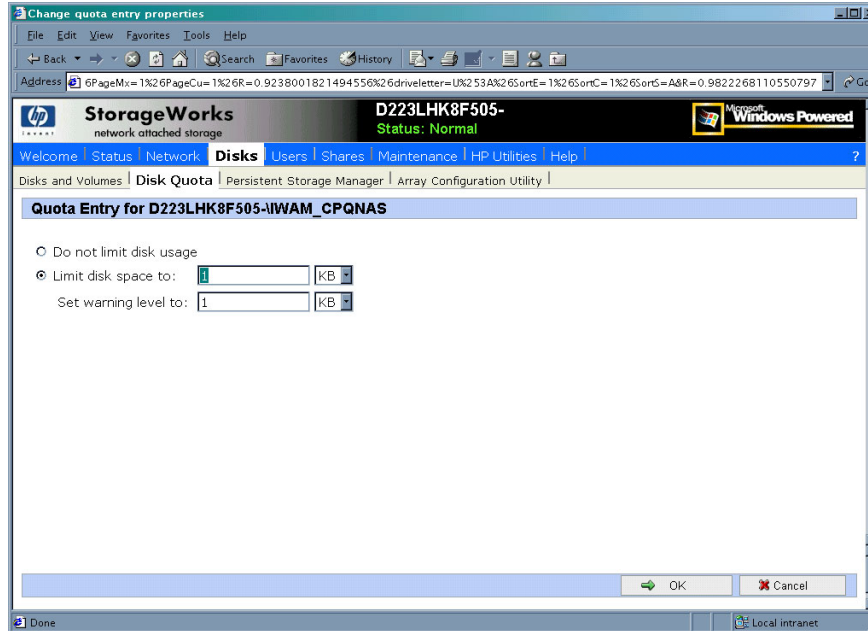


Figure 73: Quota Entry dialog box for a user

3. Enter the new disk limit information, and then click **OK**. The **Quota Entries** dialog box is displayed again.

Folder and Share Management

8

The HP 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 NAS Management Console user interface. For information on managing NCP file shares, see the "NetWare 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
- Share Management
 - Share Considerations
 - Defining Access Control Lists
 - Integrating Local File System Security into Windows Domain Environments
 - Comparing Administrative (Hidden) and Standard Shares
 - Planning for Compatibility between File Sharing Protocols

- Managing Shares
 - > Creating a new share
 - > Deleting a share
 - > Modifying share properties
 - > CIFS sharing
 - > NFS sharing
 - > FTP sharing
 - > Web sharing (HTTP)
 - > Netware sharing (NCP)
 - > AFP (AppleTalk) sharing
 - > Installing services for AppleTalk
 - > Installing Windows NT Services for Macintosh

- Protocol Parameter 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.

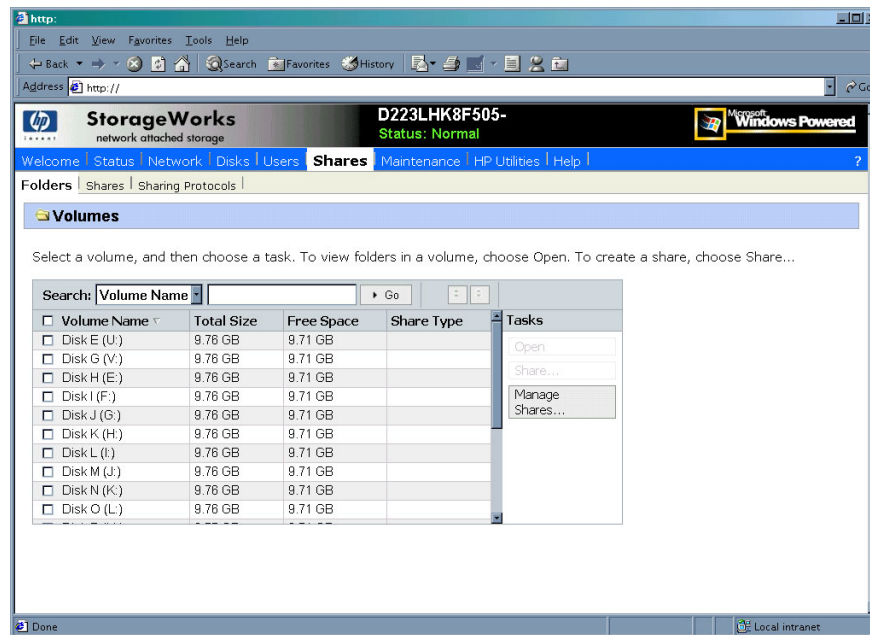


Figure 74: Volumes dialog box

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.
3. 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 75](#) for an example of **Folders** dialog box.

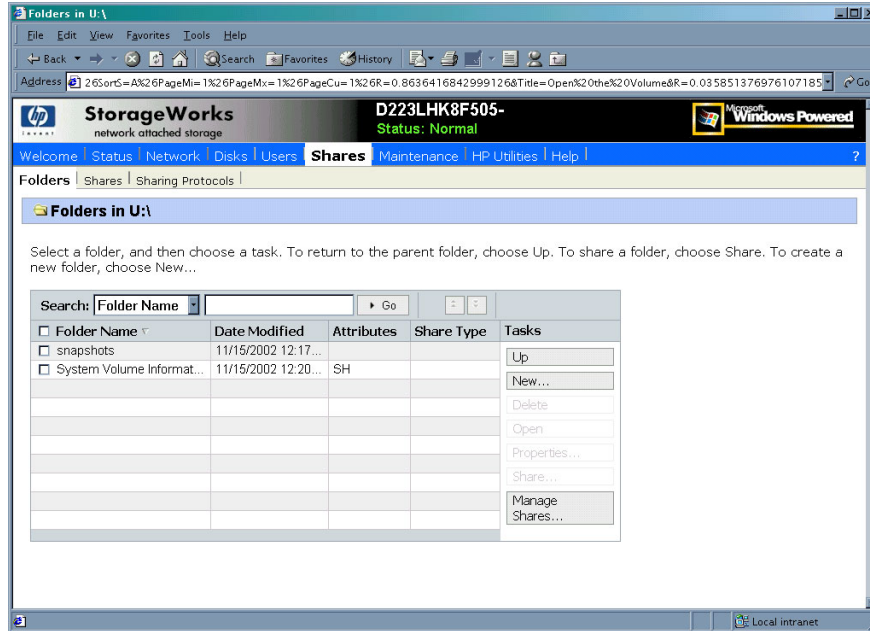


Figure 75: Folders dialog box

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.

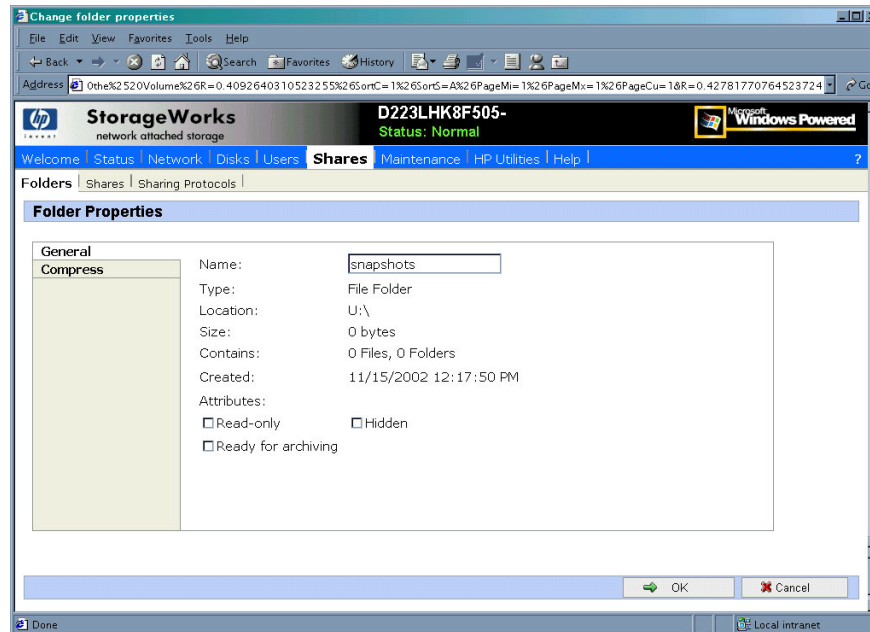


Figure 76: Create a New Folder dialog box, General tab

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.

Note: 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.

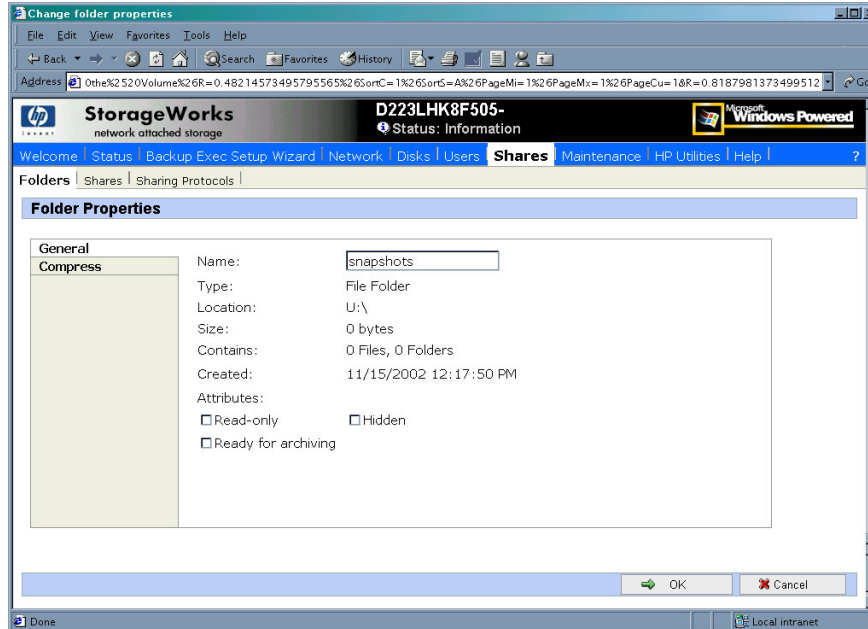


Figure 77: Folder Properties dialog box, General tab

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.

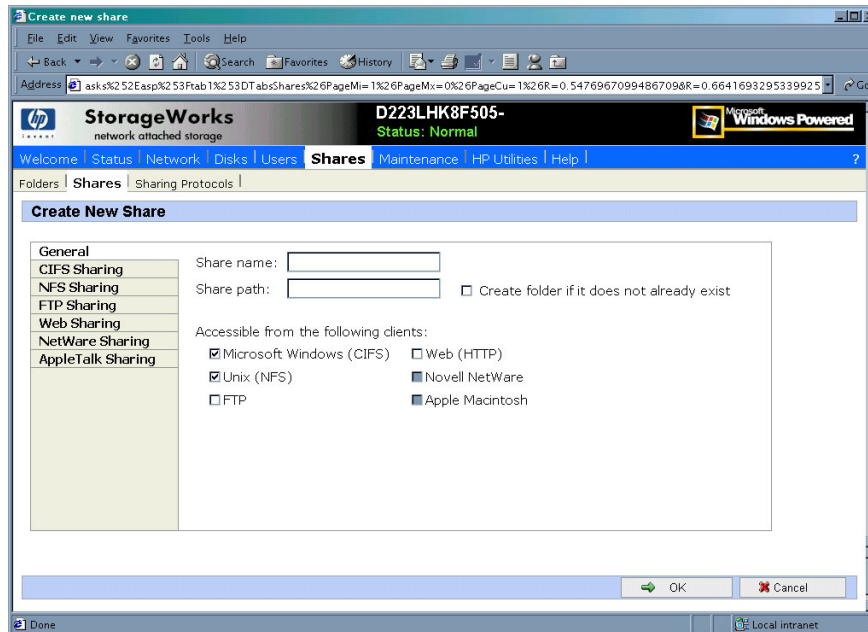


Figure 78: Create New Share dialog box, General tab

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 the "Managing Shares" 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 "Deleting 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 "Moifying Share Properties" 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 79](#) illustrates the properties available on the **Security** tab.

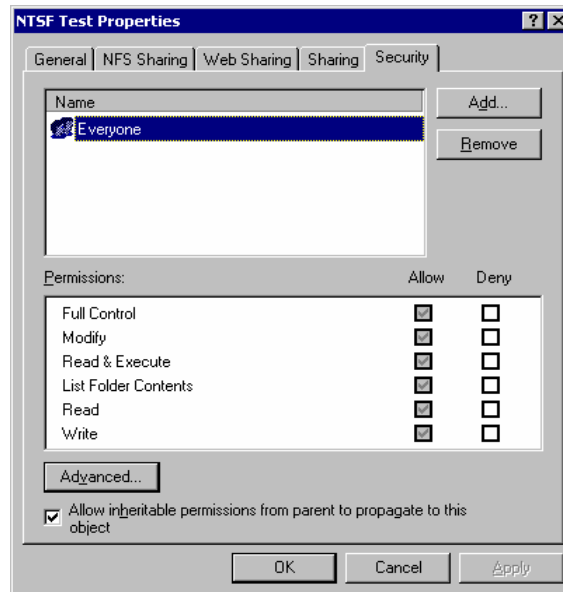


Figure 79: Security Properties dialog box for folder name NTFS Test

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 **Advanced**.




Figure 80: Access Control Settings dialog box for folder name NTFS 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.
2. Click **View/Edit**.
3. 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 81](#) illustrates the **View/Edit** screen and some of the permissions.

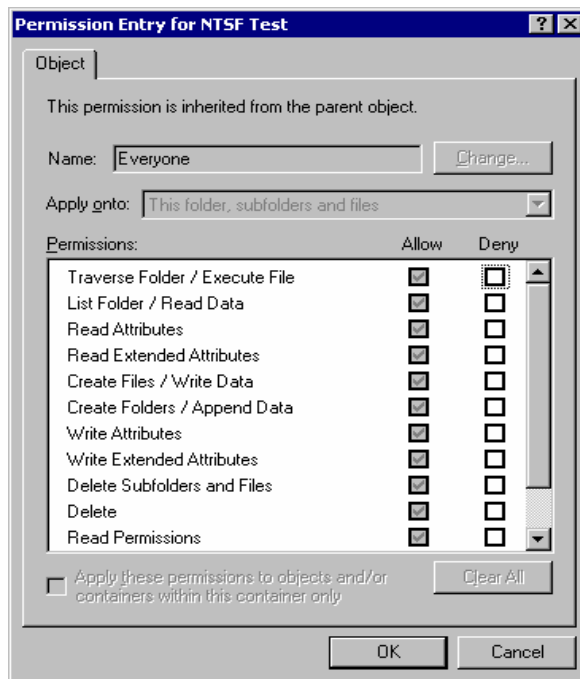


Figure 81: User or Group Permission Entry dialog box for folder name NTFS Test

Other functionality available in the **Advanced Access Control Permissions** tab is illustrated in [Figure 81](#) 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 82](#).

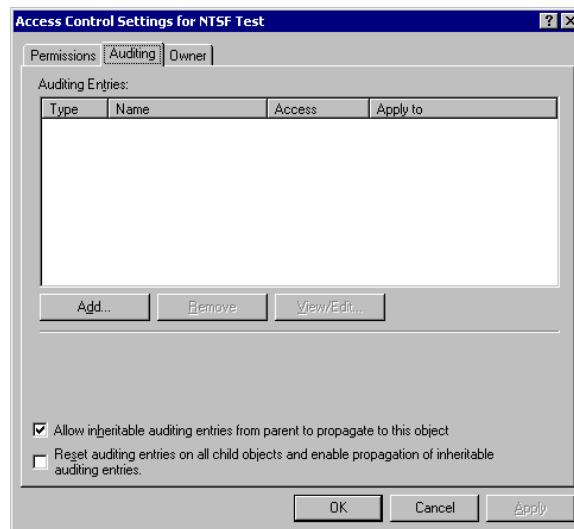


Figure 82: Access Control Settings, Auditing tab dialog box for folder name NTSF Test

[Figure 83](#) illustrates the screen that is displayed when a user or group to be audited is added.

4. 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.

5. Select the user or group.

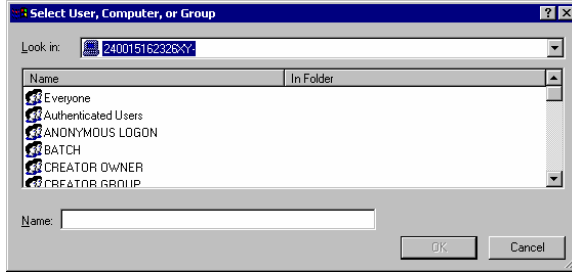


Figure 83: Select User, Computer, or Group dialog box

6. Click **OK**. [Figure 84](#) illustrates the **Auditing Entry** screen that is displayed.

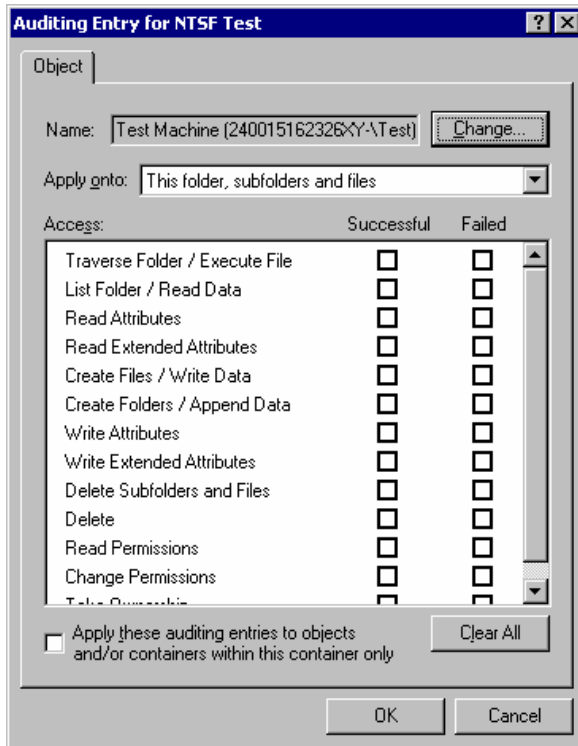


Figure 84: Auditing Entry dialog box for folder name NTSF Test

7. Select the desired **Successful** and **Failed** audits for the user or group as shown in [Figure 84](#).
8. 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 85](#) illustrates the **Owner** tab.

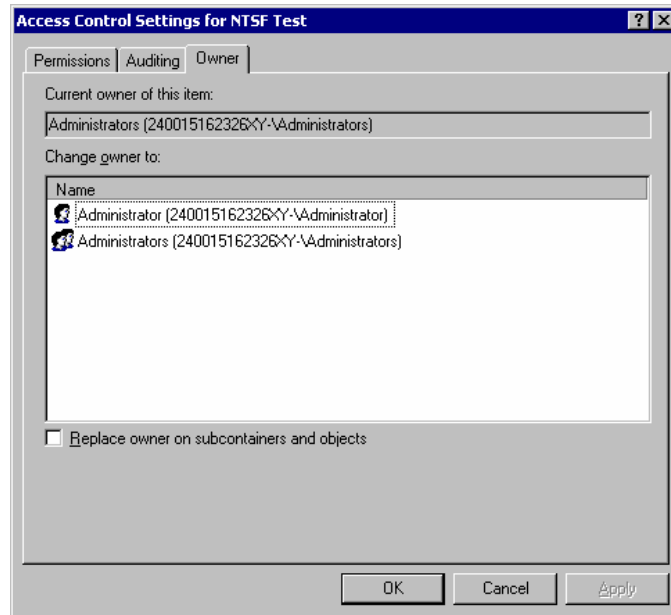


Figure 85: Access Control Settings, Owner tab dialog box for folder name NTSF Test

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 NAS Management Console. 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 Permissions" earlier in this chapter for information on file security.

Shares management topics include:

- Share Considerations
- 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

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.

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 "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

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.

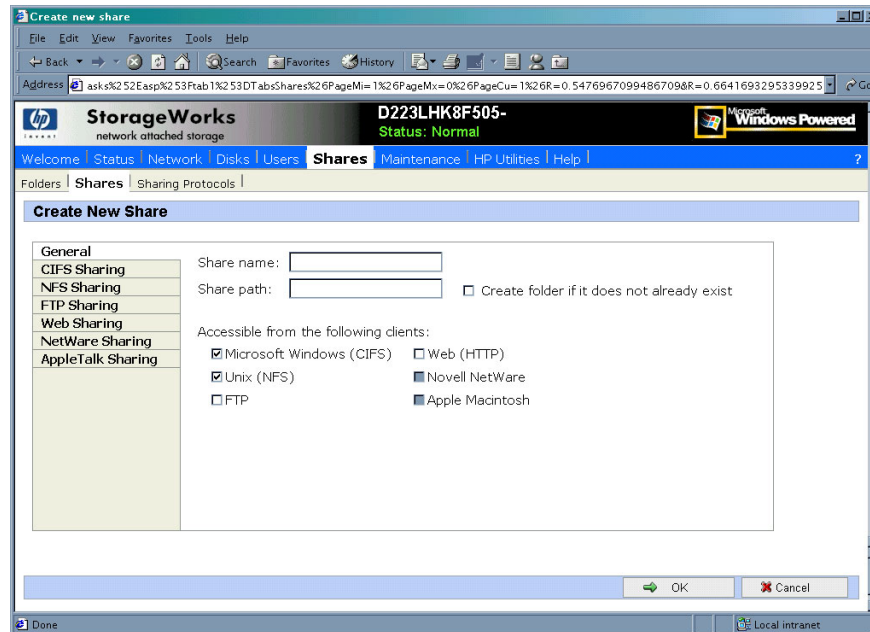


Figure 86: Create a New Share dialog box, General tab

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



Caution: 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.

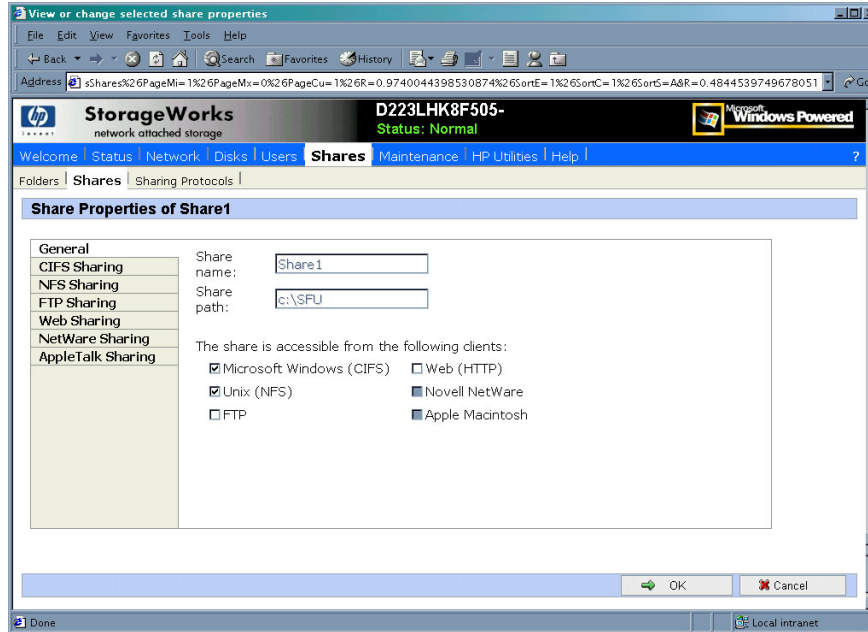


Figure 87: Share Properties dialog box, General tab

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 88](#) 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.

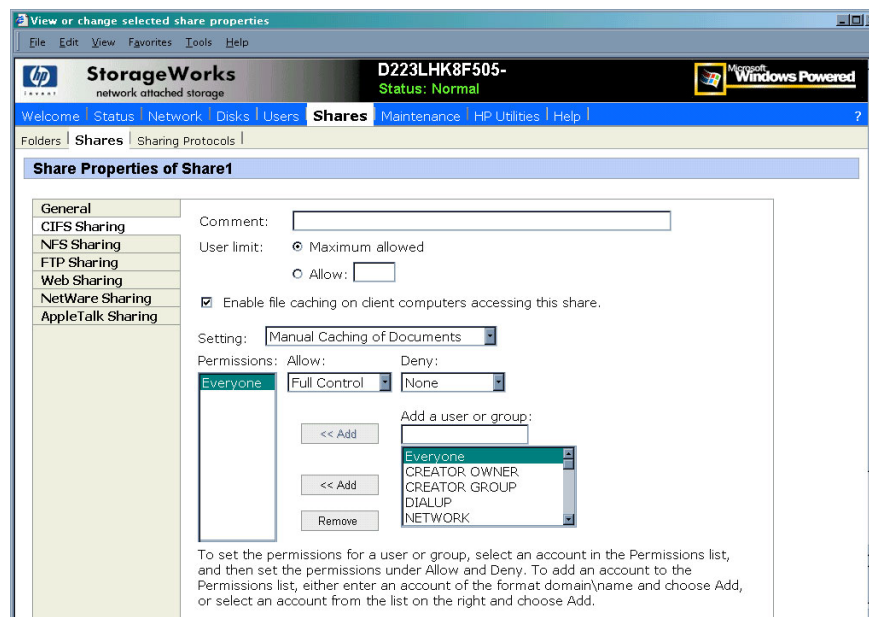


Figure 88: Share Properties dialog box, CIFS Sharing tab

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**.

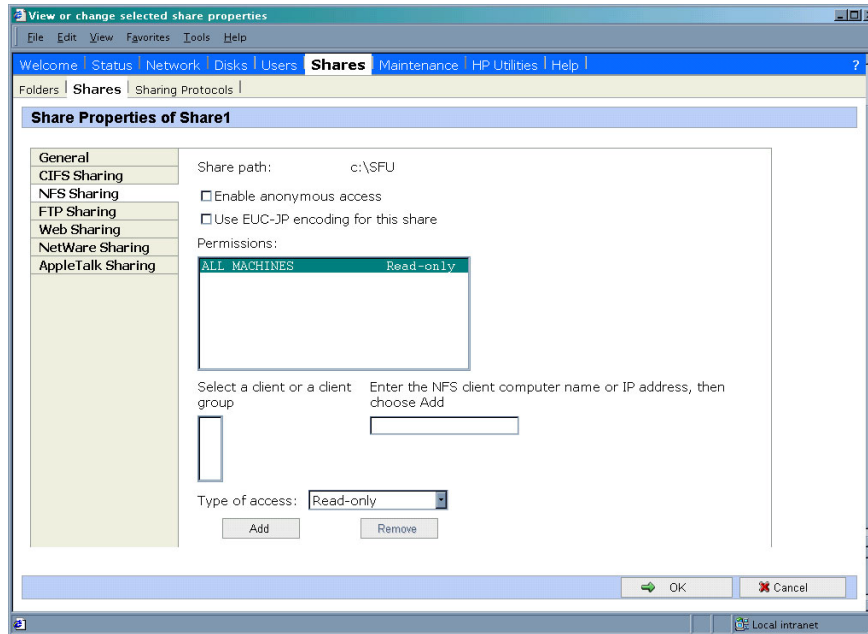


Figure 89: Share Properties dialog box, NFS Sharing tab

2. Indicate whether to allow anonymous access to the NFS share.
3. 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.

4. 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 90](#), 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**.

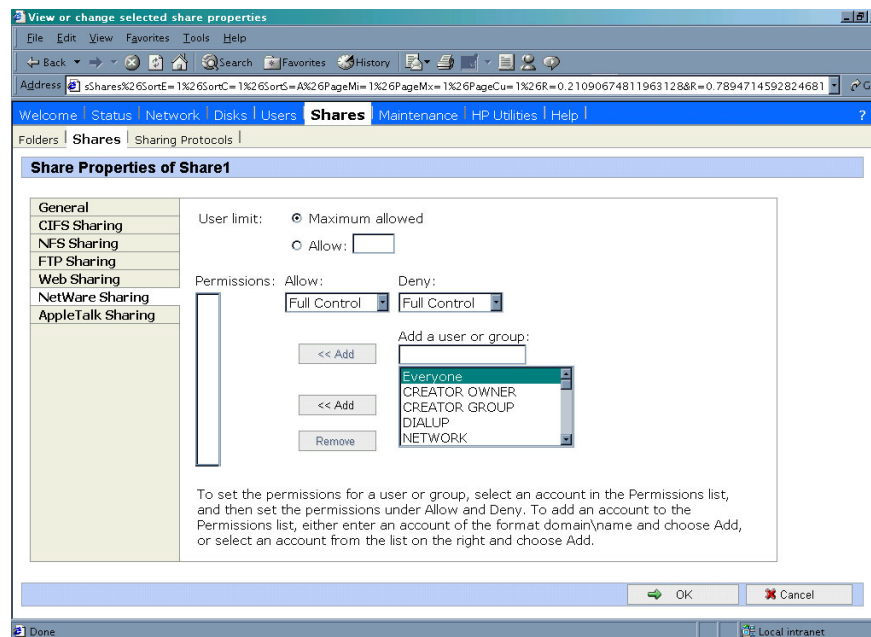


Figure 90: Share Properties dialog box, NetWare Sharing tab

AFP (AppleTalk) Sharing

AppleTalk shares can be set up only after Service for AppleTalk and Microsoft Windows NT Services for Macintosh have been installed on the NAS b2000.

Installing Services for AppleTalk

To install Services for AppleTalk:

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 91](#) is an example of the **Select Network Component Type** dialog box.

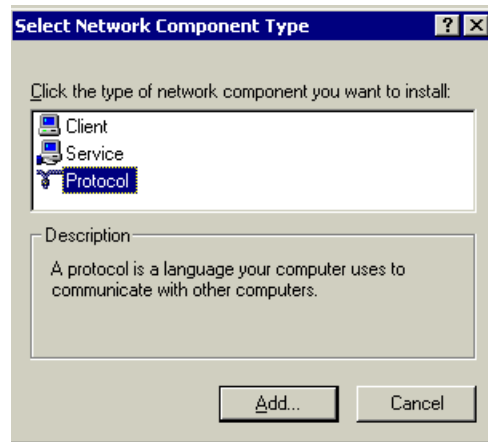


Figure 91: Local Area Connection Properties page, Install option

3. Select **Protocol** and click **Add**.
4. Select **AppleTalk Protocol** and click **OK**.

Installing Windows NT Services for Macintosh

To install Windows NT Services for Macintosh:

1. Select **Maintenance** from the WebUI interface.
2. Select **Terminal Services**.
3. Open **Add/Remove Programs** from the Control Panel.
4. Click **Add/Remove Windows Components**.
5. Double-click **Other Network File and Print Services**.
6. Select **File Services for Macintosh** then click **OK**.
7. Click **Next**.
8. Click **Finish**.

To set up AppleTalk shares, 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.

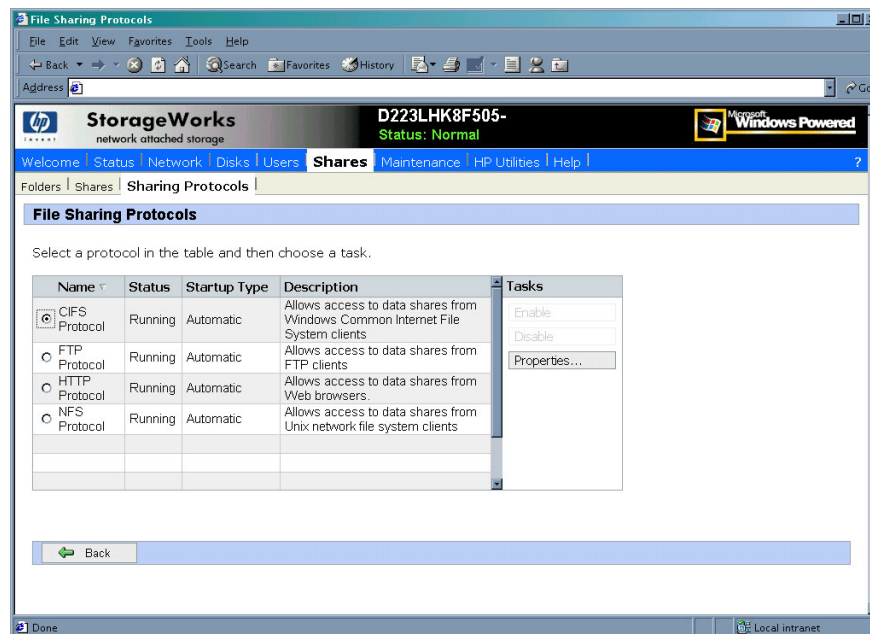


Figure 92: Sharing Protocols dialog box

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 Mappings

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

9

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
- Server for NFS
 - Authenticating User Access
 - Indicating the Computer to Use for the 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 Service, and Remote Shell Service
 - Using Terminal Services
 - Using Telnet Service
 - 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™ 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 "NFS 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 that will store user name mappings, the name of that computer must be indicated, as detailed below:

1. Use Terminal Services to access the **NAS Management Console**, click **File Sharing, Services for UNIX**. Click **Server for NFS**. [Figure 93](#) 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.

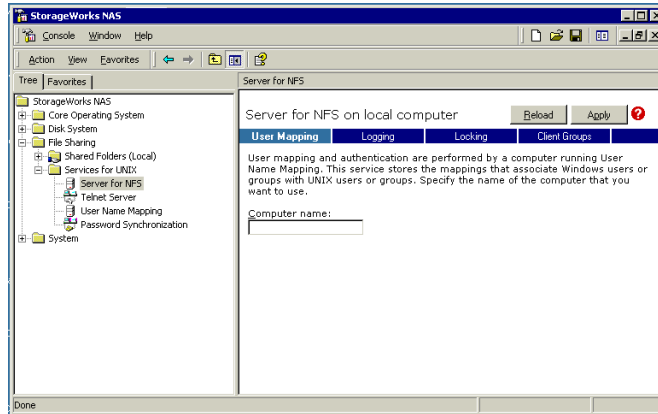


Figure 93: NAS Management Console Server for NFS screen, User Mapping tab

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 NAS Management Console, click **File Sharing, Services for UNIX, Server for NFS**. Click the **Logging** tab.
2. To log events to the event viewer application log, click the check box for **Log events to event log**.
3. To log selected event types, click the check box for **Log events in this file** on the screen.
4. Enter a filename or use the default filename provided (*rootdrive\SFUlog\nfssvr.log*) and log file size (7-MB default). The default log file is created when the changes are applied.

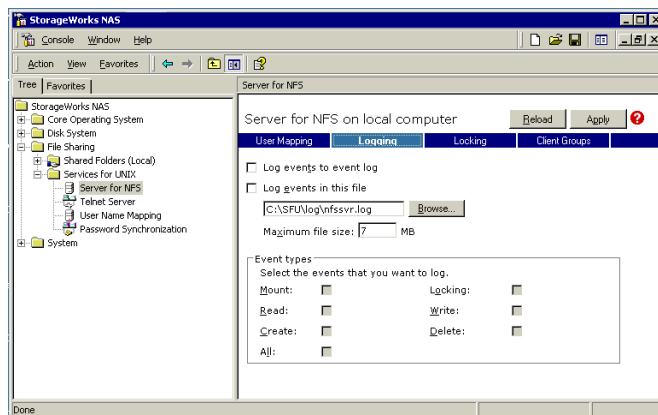


Figure 94: NAS Management Console Server for NFS screen, Logging tab

Installing NFS Authentication Software on the Domain Controllers and Active Directory 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. This includes active directory domains. For instructions on setting up user mappings, see "NFS 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 "Folder and 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.

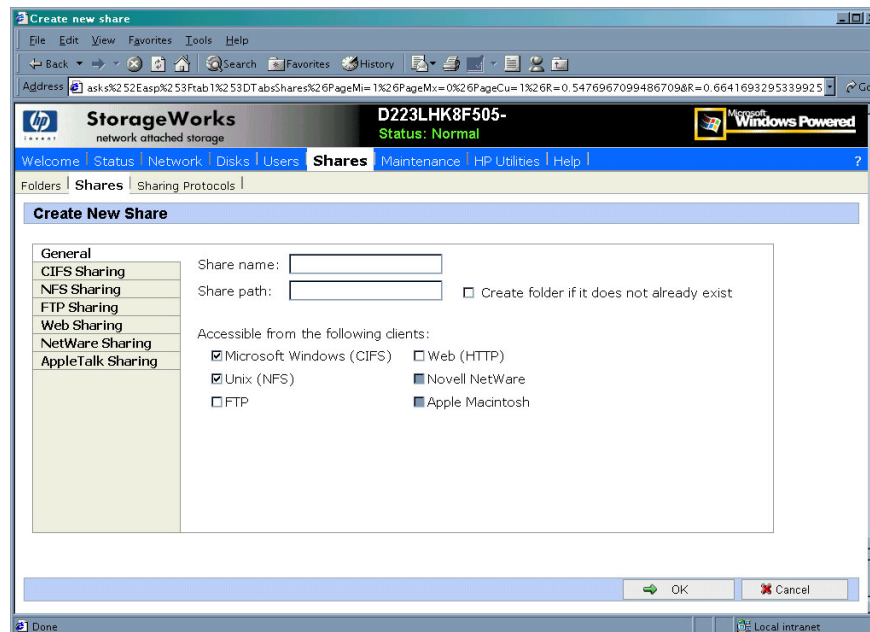


Figure 95: Create a New Share dialog box, General tab

2. In the **General** tab, enter the share name and path. Check the **Unix (NFS)** client protocol check box.

Note: Uncheck the Microsoft Windows (CIFS) option if you do not want to allow CIFS access to the share.

Note: 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



Caution: 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.

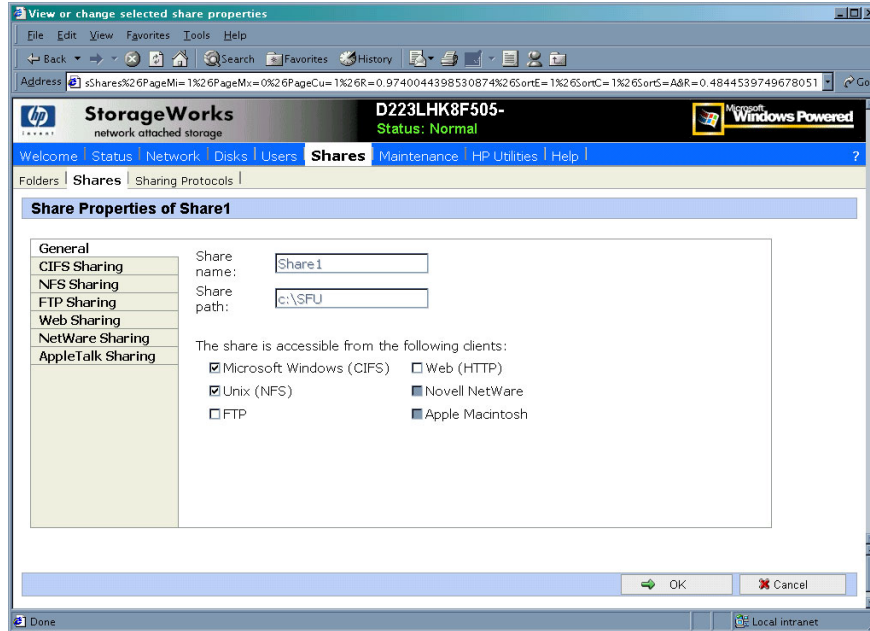


Figure 96: Share Properties dialog box, General tab

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.

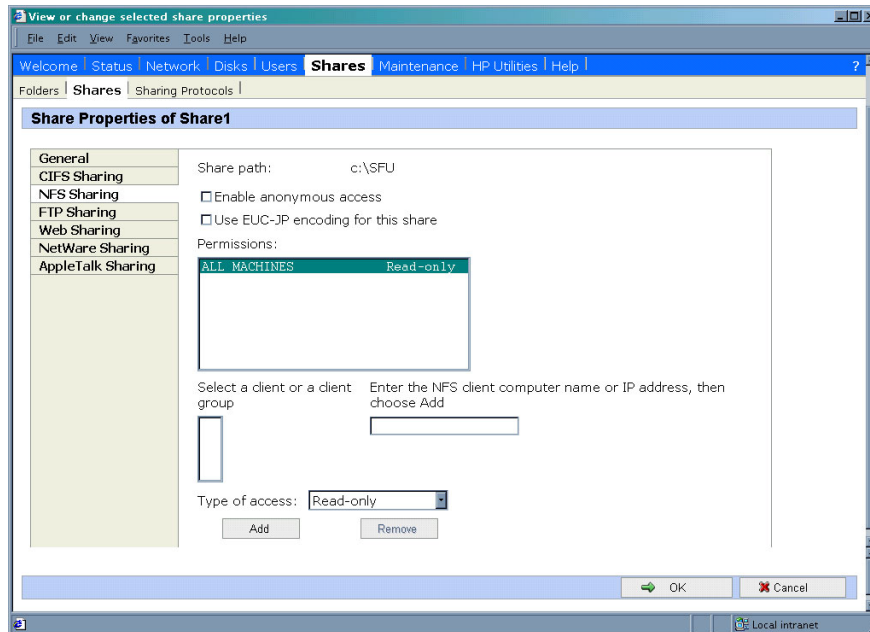


Figure 97: NFS Sharing tab

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 whether to allow anonymous access to the NFS share.

Note: The default values for Anonymous UID and Anonymous GID are -2. Non-default IDs can be specified for the NFS share using Terminal Services.

- c. 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:
 - **Read-only**—Use this permission to restrict write access to the share.
 - **Read-write**—Use this permission to allow clients to read or write to the share.
 - **Read-only + Root**—Use this permission to restrict write access to the share. 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 + Root**—Use this permission to allow clients to read or write to the share. 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.
 - **No access**—Use this permission to restrict all access to the share.
4. After all NFS sharing information is entered, click **OK**.

Encoding Types

Two encoding types can be selected using the WebUI. These include the default ANSI as well as EUC-JP. Other encoding types can be assigned to the NFS share using Terminal Services. The encoding choices are:

- ANSI (default) - able to assign with the WebUI
- BIG5 (Chinese)
- EUC-JP (Japanese) - able to assign with the WebUI
- EUC-KR (Korean)
- EUC-TW (Chinese)
- GB2312-80 (Simplified Chinese)
- KSC5601 (Korean)
- SHIFT-JIS (Japanese)

If the option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:

- Japanese: SHIFT-JIS
- Korean: KS C 5601-1987
- Simplified Chinese: GB
- Traditional Chinese: BIG5

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.

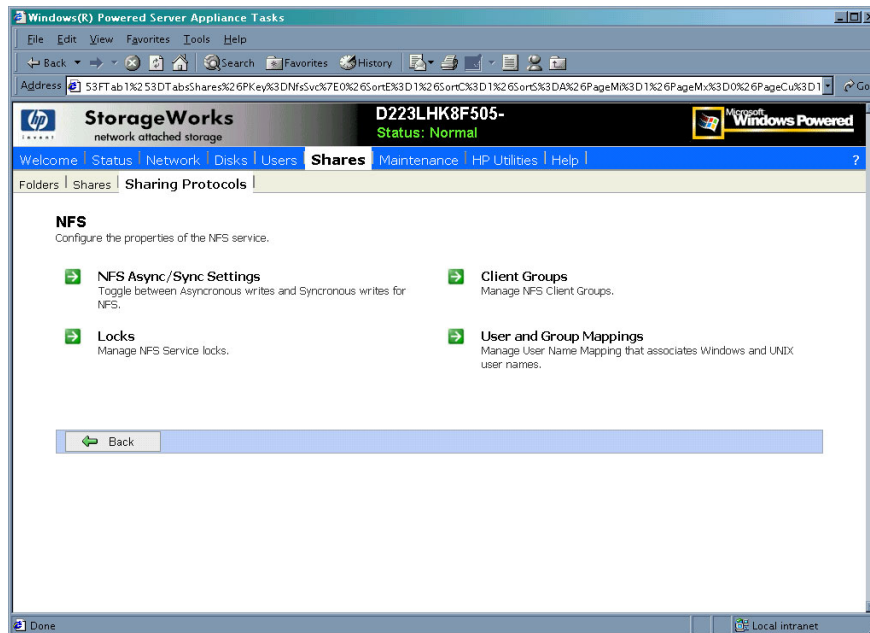


Figure 98: NFS Sharing Protocols menu

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.

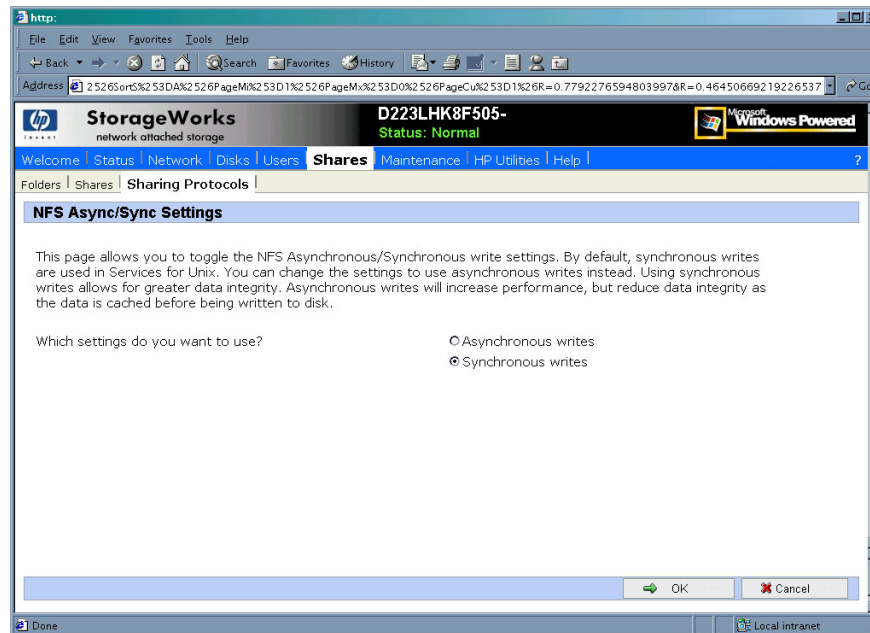


Figure 99: NFS Async/Sync Settings dialog box

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 100** 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.

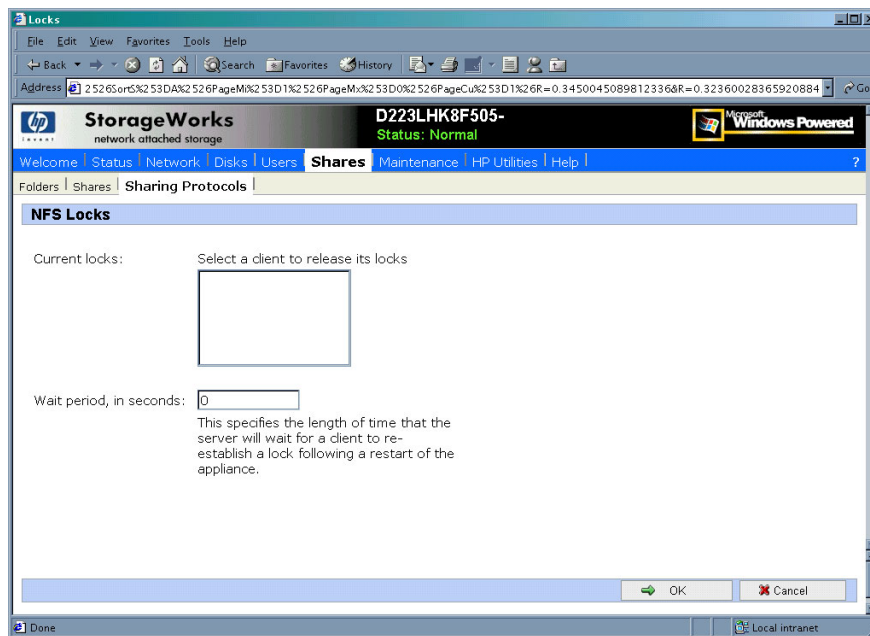


Figure 100: NFS Locks dialog box

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.

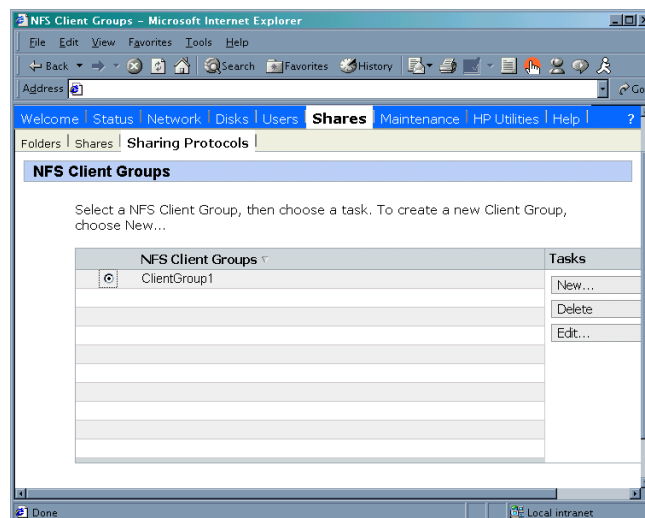


Figure 101: NFS Client Groups dialog box

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.

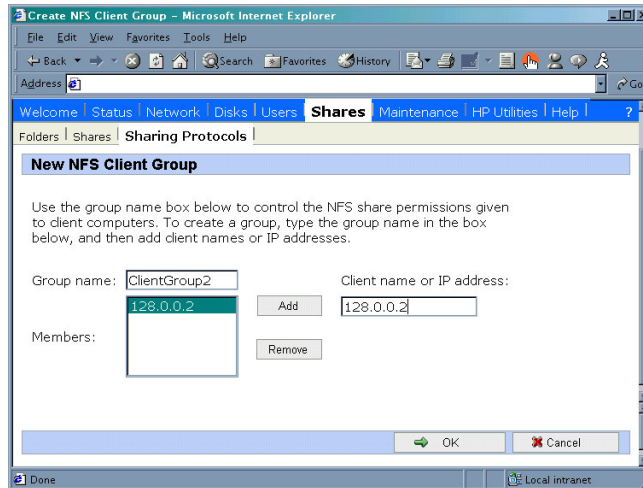


Figure 102: New NFS Client Group dialog box

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.

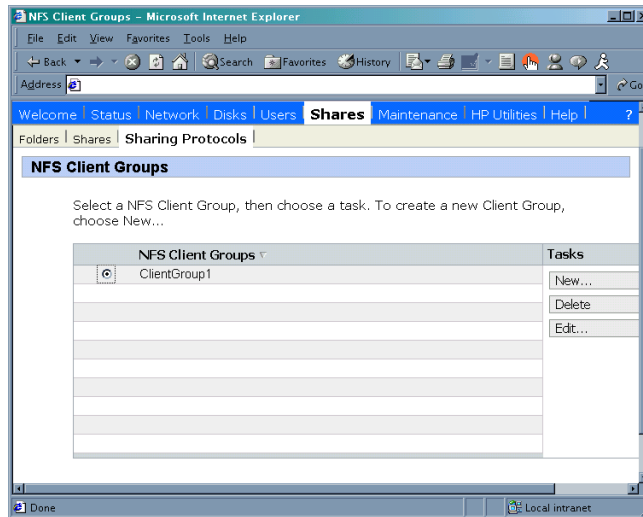


Figure 103: Client Groups dialog box

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.

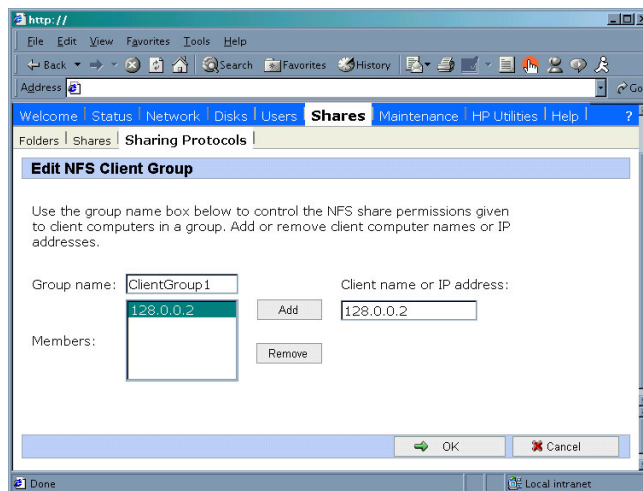


Figure 104: Edit NFS Client Groups dialog box

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 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 "OEM Supplemental Help" chapter of the SFU help, found on the NAS b2000.

Figure 105 is a diagram showing an example of how the mapping server works for an `ls -al` command.

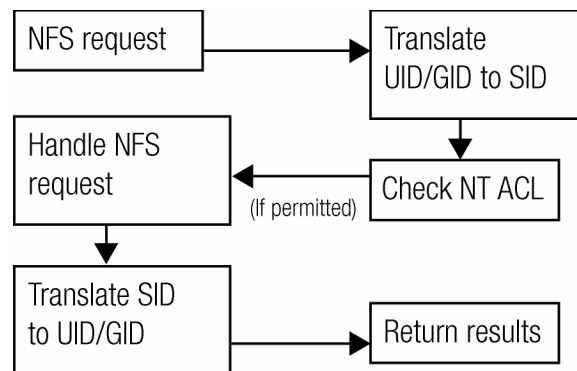


Figure 105: Mapping Server "ls -al" Command example

A double translation, as illustrated in Figure 105, 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**.

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.

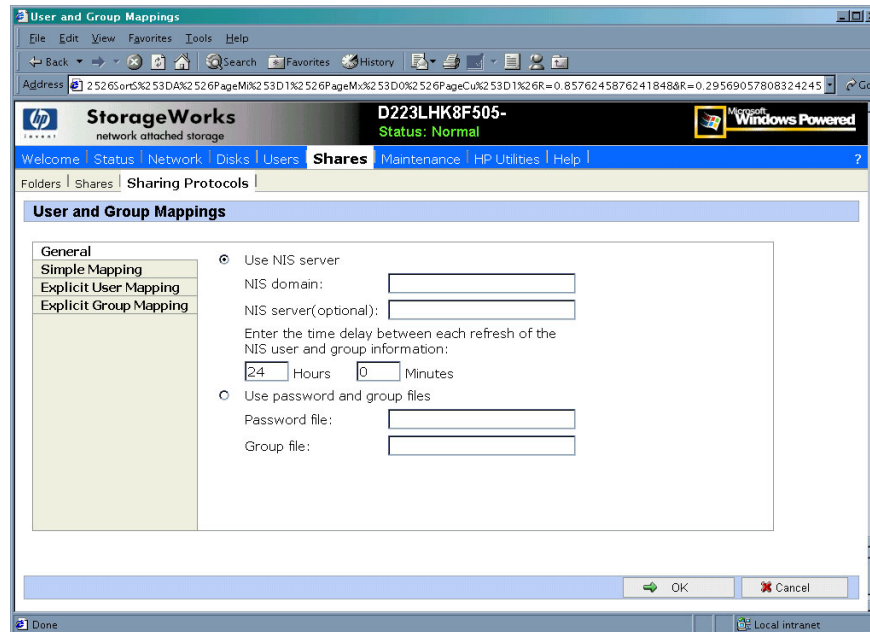


Figure 106: User and Group Mappings dialog box, General tab

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 select the Windows domain name.

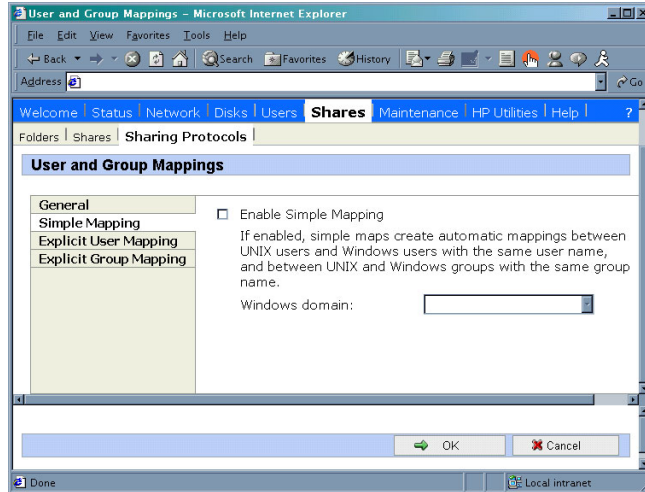


Figure 107: User and Group Mappings dialog box, Simple Mapping tab

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 108](#) is an example of the **Explicit User Mapping** tab.

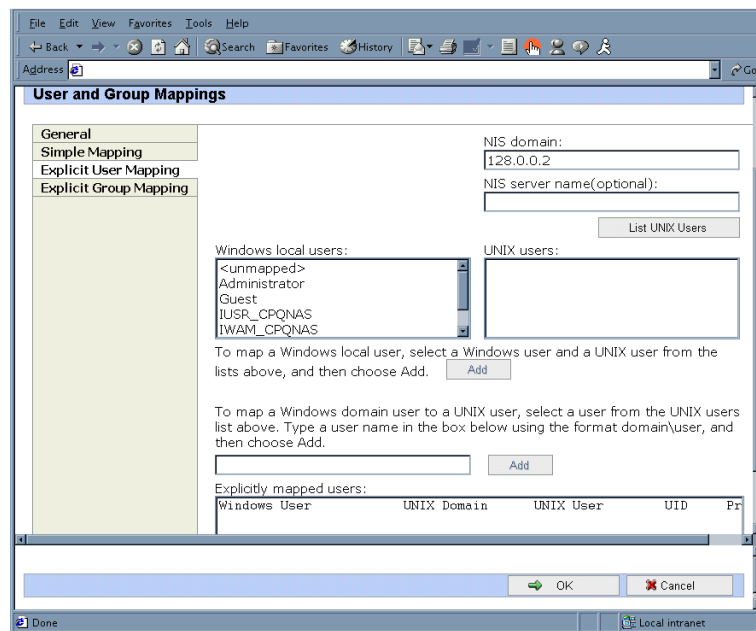


Figure 108: User and Group Mappings dialog box, Explicit User Mapping tab

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 109](#) 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.

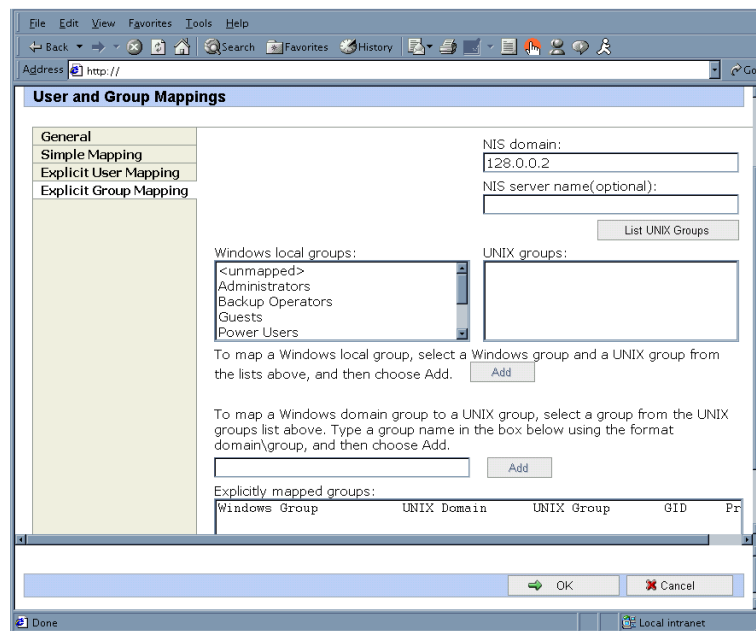


Figure 109: User and Group Mappings dialog box, Explicit Group Mapping tab

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 NAS Management Console under the **Map Maintenance** tab of the **User Name Mapping** screen, as shown in [Figure 110](#).

To access the NAS Management Console, use Terminal Services. To open a Terminal Services session, from the WebUI, select **Maintenance**, **Terminal Services**.

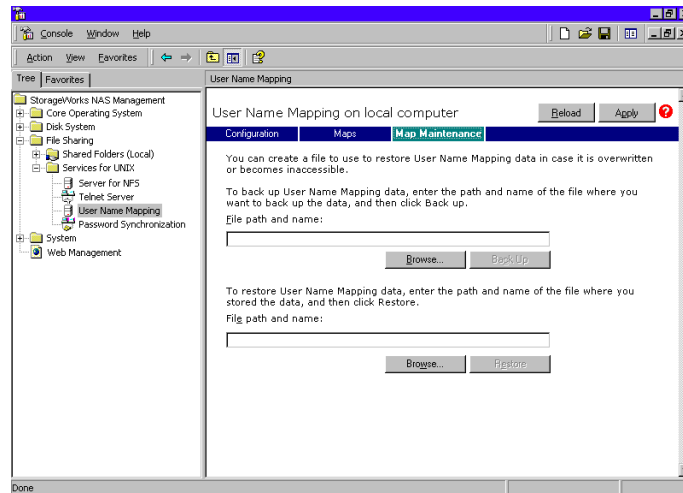


Figure 110: NAS Management Console User Name Mapping screen, Map Maintenance tab

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.
4. 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

HP 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 Linux/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`)
 - 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 20](#) 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 NAS Management Console 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 Server

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 20](#) describes some common SFU commands.

Table 20: Command Line Interface Command Prompts

Command	Function
<code>nfsstat /?</code>	Learn about viewing statistics by NFS operation type
<code>showmount /?</code>	View the format of the command to display NFS export settings on NFS servers
<code>showmount a</code>	View users who are connected and what they currently have mounted
<code>showmount e</code>	View exports from the server and their export permissions
<code>rpcinfo /?</code>	Learn how to display Remote Procedure Call (RPC) settings and statistics
<code>mapadmin /?</code>	View how to add, delete, or change user name mappings
<code>tnadmin /?</code>	View how to change Telnet Server settings
<code>nfsshare /?</code>	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.

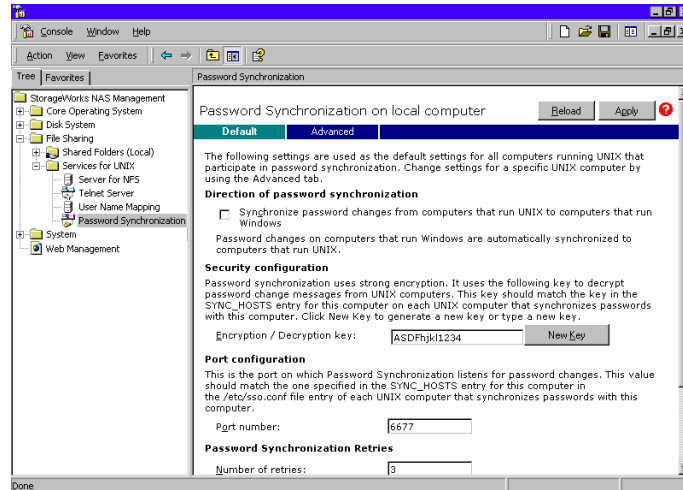


Figure 111: Password Synchronization screen

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.

- 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 **NAS Management Console**. From the **NAS Management Console**, 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.

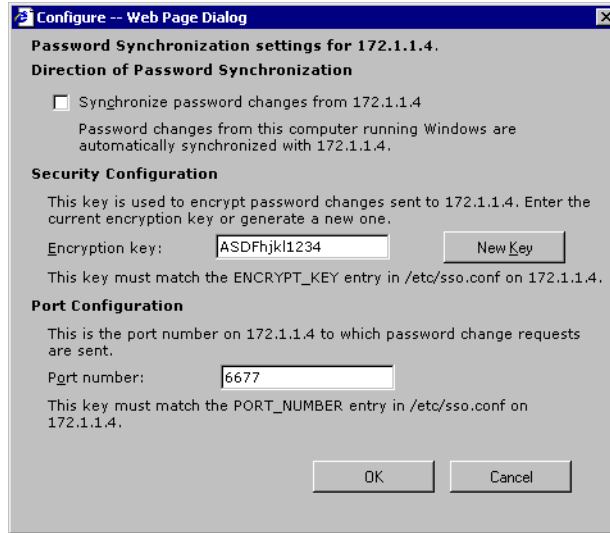


Figure 112: Password Synchronization screen, Advanced Settings dialog box

Installing Password Synchronization on Domain Controllers and Active Directory Domain Controllers

The password synchronization service must be installed on all primary domain controllers (PDCs) and backup domain controllers (BDCs) in a domain that will implement the password synchronization service. This includes Active Directory domains. 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.



Caution: 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

10

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.

Note: 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
- Managing NCP 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:

www.microsoft.com/WINDOWS2000/guide/server/solutions/NetWare.asp

Note: 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 113 is an example of the **Select Network Component Type** dialog box.




Figure 113: 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\SFNFNWA*.
5. Select the *NETSFNTRV* 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**.

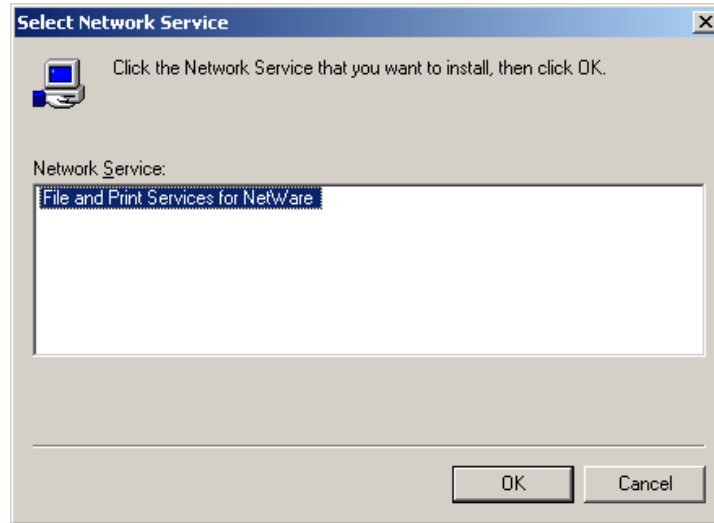


Figure 114: Installing File and Print Services for NetWare

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**.

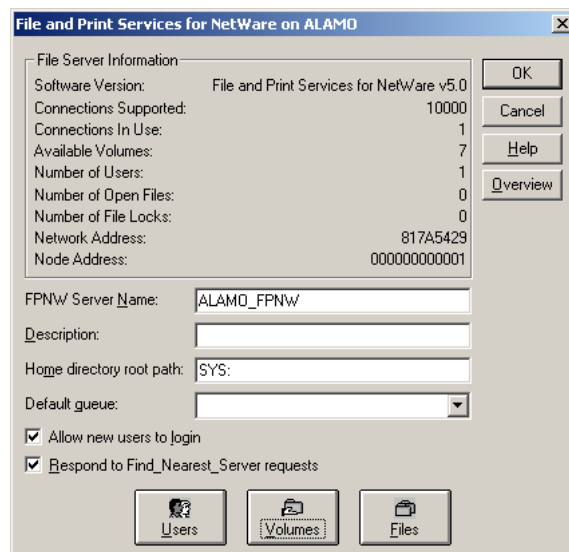


Figure 115: File and Print Services for NetWare screen

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 115](#) 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 **Users** 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 **Volumes** to:
See users connected to specific volume and to disconnect users from a specific volume.
6. Click **Files** to:
View 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**.

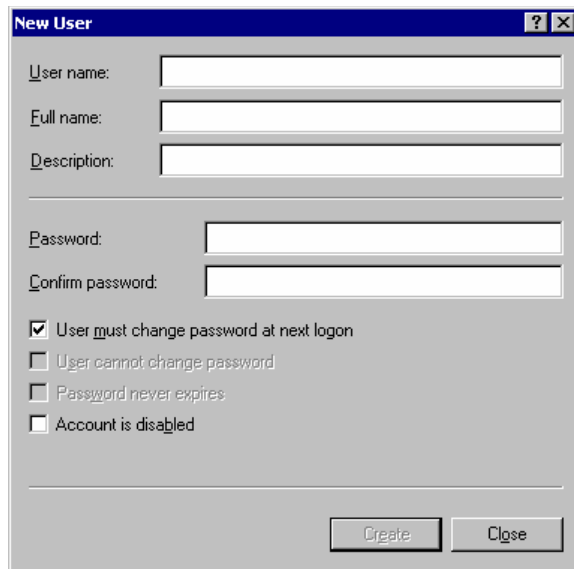


Figure 116: New User dialog box

3. Enter the user information, including the user's User name, Full name, Description, and Password. Click **Create**.
4. Repeat these steps until all NetWare users have been entered.

Enabling Local NetWare User Accounts

1. In the **Users** folder (NMC, 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.

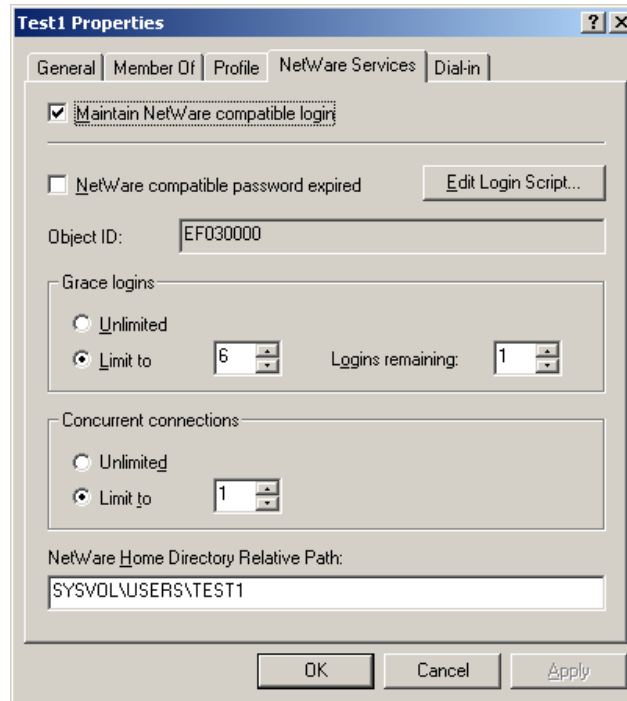


Figure 117: NetWare Services tab

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)

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
- NAS Management Console

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.

Shares can be managed through the Shares menu option of the WebUI. Tasks include:

- Creating a new NCP share
- Deleting an NCP share
- Modifying NCP 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.

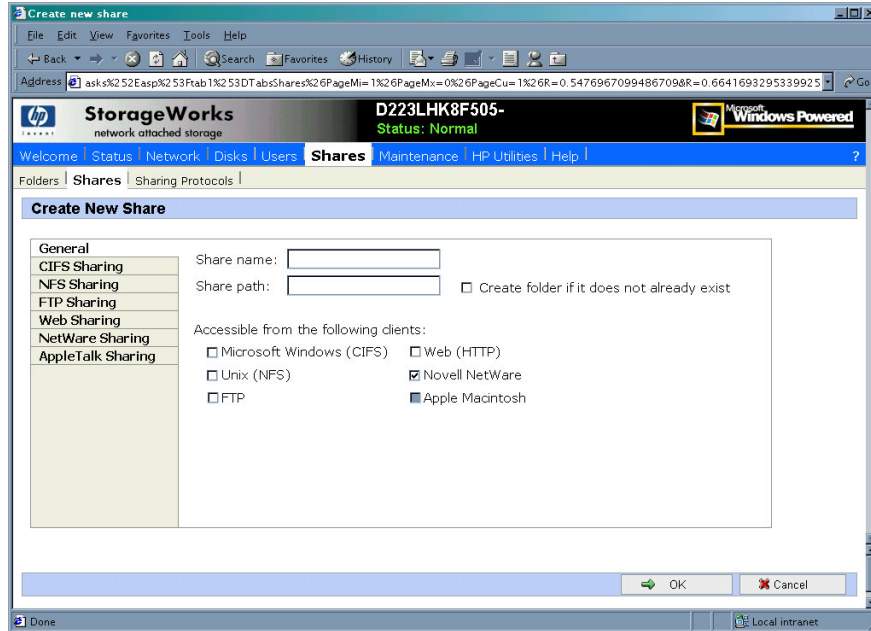


Figure 118: Create a New Share dialog box, General tab

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



Caution: 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 are displayed.

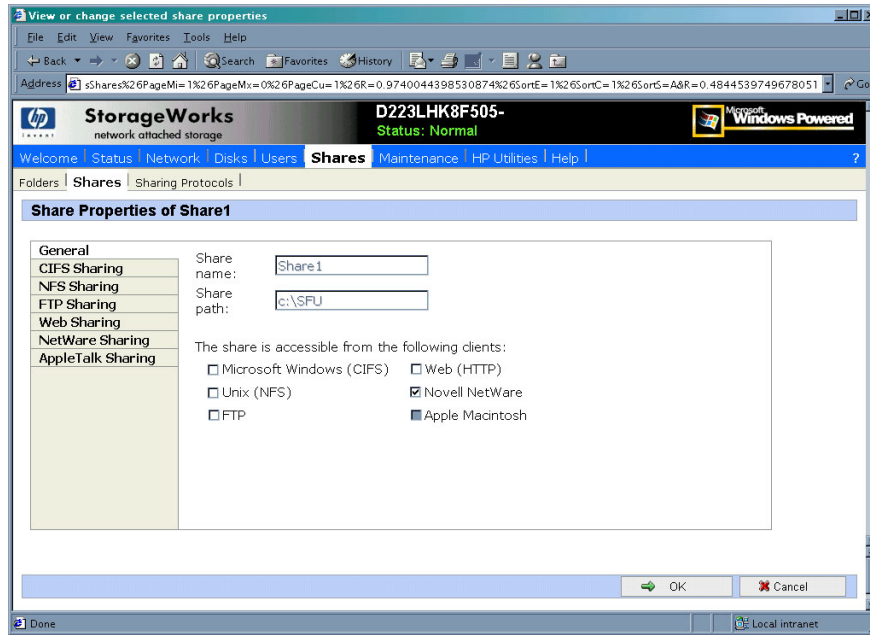


Figure 119: Share Properties dialog box, General tab

2. To enter or change client protocol information, check the **Novell NetWare** client type box and then click the **NetWare Sharing (NCP)** tab.

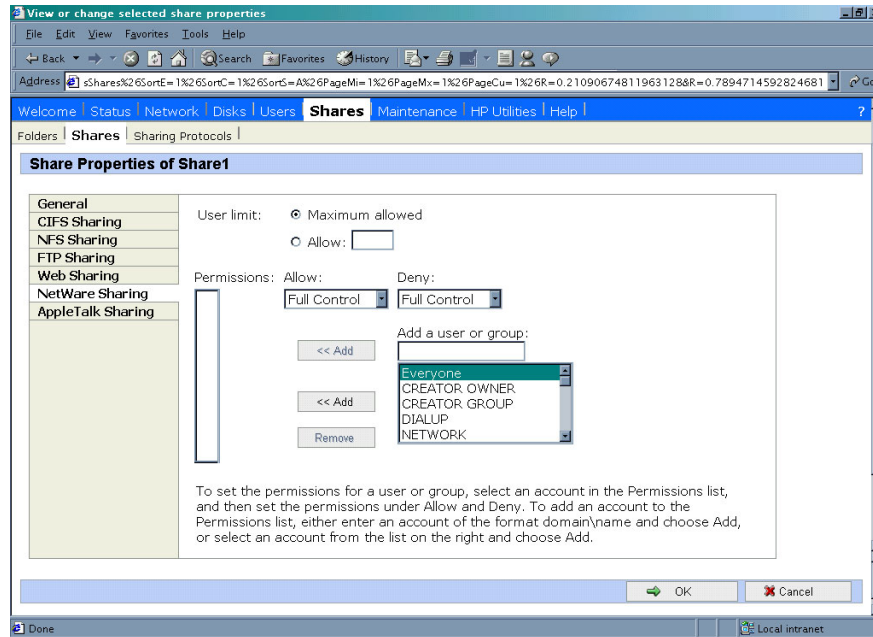


Figure 120: Share Properties dialog box, NetWare Sharing tab

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 NAS Management Console

In addition to the WebUI available on the NAS b2000, shares can be managed through the NAS Management Console. 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 NAS Management Console

To create a new file share:

1. From the NAS b2000 desktop, click the **NAS 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.

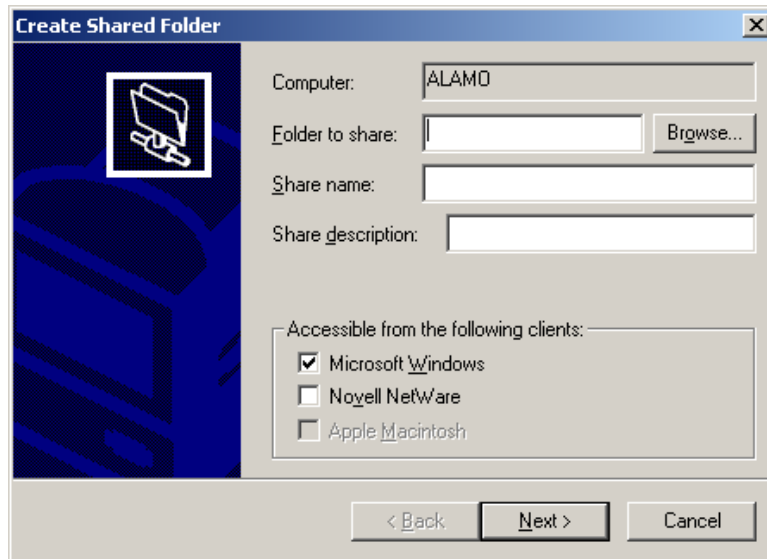


Figure 121: Create Shared Folder dialog box

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 122](#) is displayed.

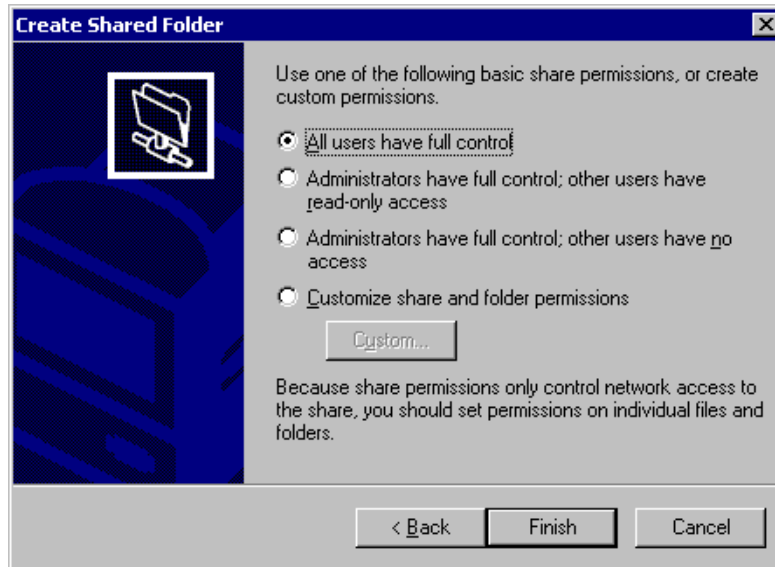


Figure 122: NetWare Basic Share Permissions dialog box

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 123](#) is an illustration of the **Customize Permissions** dialog box.

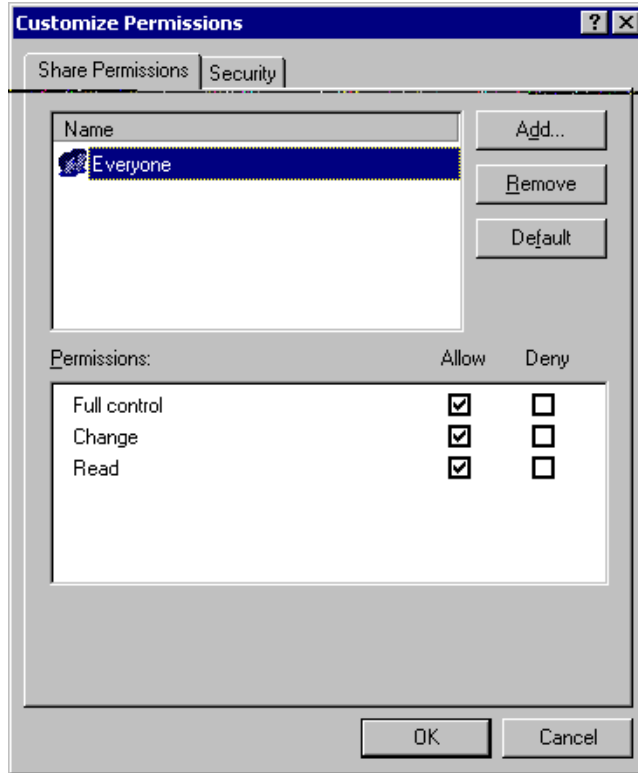


Figure 123: Customize Permissions dialog box, Share Permissions tab

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.

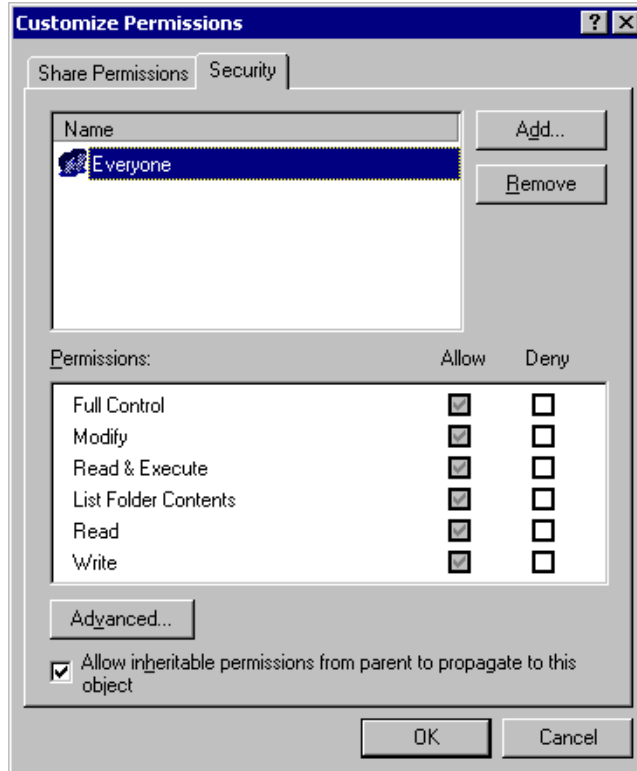


Figure 124: Customize Permissions dialog box, Security tab

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 NAS Management Console

To change share settings through the NAS Management Console:

1. From the NAS b2000 desktop, select the **NAS 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

11

The HP StorageWorks NAS b2000 comes from the factory with full remote manageability. Several methods of remote access are provided:

- Web based user interface
- Terminal services
- Integrated Lights-Out Port
 - Features
 - Integrated Lights-Out Port Configuration
 - Using the Integrated Lights-Out Port to Access the NAS b2000
- Telnet Server
 - Enabling Telnet Server
 - Configuring Telnet Server
- Remote Shell Daemon
- Insight Manager
 - Insight Manager Console
 - Insight Manager Agent Web Interface
- Enterprise management applications
 - HP OpenView (Windows-Based Operating System)
 - 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.

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 Management Console
- 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.

Integrated Lights-Out Port

The following information provides an overview of the integrated Lights-Out port capabilities. For further information, refer to the *Integrated Lights-Out Port Installation and Users Guide* on the Documentation CD.

The integrated Lights-Out port is an ASIC-based Web interface that provides 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 integrated Lights-Out port independent of the host server and its operating system. The integrated Lights-Out port 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 integrated Lights-Out port provides the following features:

- Hardware based graphical remote console access

Note: The remote client console must have a direct browser connection to the integrated Lights-Out port without passing through a proxy server or firewall.

- Remote restart
- Server failure alerting
- Integration with 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 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 *Integrated Lights-Out Port User Guide* for more information about the integrated Lights-Out port features and functionality.

Integrated Lights-Out Port Configuration

The integrated Lights-Out port on the NAS b2000 is initially configured through the Rapid Startup Utility. SNMP is enabled and the Insight Management Agents are preinstalled.

The integrated Lights-Out port 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 *Integrated Lights-Out Port User Guide* for information about changing these settings.

There are several methods for performing integrated Lights-Out port configuration changes:

- Web interface
- Integrated Lights-Out port configuration utility accessed by pressing **F8** during a system restart.

Note: You must connect locally with a monitor, keyboard, and mouse.

- Integrated Lights-Out port access using the default DNS name
- The integrated Lights-Out port 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 integrated Lights-Out port)
 - IP Address: The IP address entered during system setup

Using the Integrated Lights-Out Port 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 integrated Lights-Out port.
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 before use.



Caution: 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 NAS Management Console. 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.

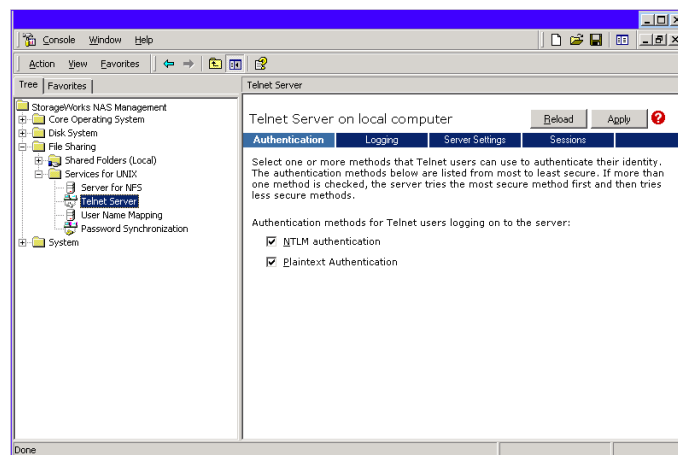


Figure 125: Telnet Server interface screen

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: A `.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.

Insight Manager

The NAS b2000 is equipped with the latest Insight Management Agents for Servers, allowing easy manageability of the server through Insight Manager, HP OpenView, and Tivoli NetView.

Insight Manager is a comprehensive management tool that monitors and controls the operation of HP servers and clients. Insight Manager Version 6.0 or later is needed to successfully manage the NAS b2000. 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 Insight Manager XE.

Insight Manager Console

System monitoring applications such as 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 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 agents.

For more information about using Insight Manager, refer to the HP Management CD.

Insight Manager Agent Web Interface

There are two options for accessing the Insight Manager Agent Web interface.

- From the 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 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.

For more information about Insight Manager, see the HP 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 Insight Manager for HP OpenView Version 2.0 or later onto the client machine.
2. Modify *CPQCONFIG.DAT* in HP OpenView.

Insight Manager for HP OpenView, Version 2.0

Insight Manager for HP OpenView integrates HP hardware management and event notification into the HP OpenView Network Node Manager (NNM) network management console.

Depending on the version of 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.

Insight Manager for HP OpenView introduces an integrated browser launch from the NNM console to the home page of the Web enabled Management Agents, as shown in [Figure 126](#). The interface is used to collect in depth information about installed HP hardware.




Figure 126: Web Enabled interface

The information collected includes system status, system health, prefailure monitors, performance and environmental data, event alarms, and Windows based OS statistics. Insight Manager for HP OpenView can be obtained from the HP website:

<http://www.hp.com/country/us/eng/prodserv/servers.html>

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 Insight Manager for Tivoli NetView Version 2.0 or later onto the client machine.
2. Modify *CPQCONFIG.DAT* in Tivoli NetView.

Insight Manager for Tivoli NetView (AIX), Version 2.0

Insight Manager for Tivoli NetView integrates HP 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 Management Agents, which is used to collect in depth information about HP hardware. The information collected includes system status, system health, prefailure monitors, performance and environmental data, event alarms, and Windows based OS statistics. Insight Manager for Tivoli NetView can be obtained from the HP website:

<http://www.managementsoftware.hp.com/>

Depending on the version of 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 Insight manager console onto the client machines.

Insight Manager:

1. Insert the HP Management CD Version 6.0 or later into the management console of the client machine.
2. From the HP Management CD window, click **Insight Manager**.
3. Select **Insight Manager** and follow the installation directions.

Insight Manager for HP OpenView (Windows 2000 Operating System)

For detailed instructions on downloading and installing Insight Manager for HP OpenView (Windows 2000 operating system), visit the following website:

<http://www.managementsoftware.hp.com/>

For detailed instructions on downloading and installing Insight Manager for HP OpenView (HPUX) visit the management area website:

<http://www.hp.com/country/us/eng/prodserv/servers.html>

Insight Manager for Tivoli NetView (AIX)

For detailed instructions on downloading and installing Insight Manager for Tivoli NetView (AIX) visit the management area of the HP website:

<http://www.hp.com/country/us/eng/prodserv/servers.html>

Backup Management

A

This appendix is a backup guide for HP StorageWorks NAS devices. This appendix 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. HP 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 of device and connection ensures a reliable backup of data that is well suited to the particular computing environment. HP recommends several tape solutions for use with the NAS b2000.

For a full list of qualified tape solutions, refer to the HP website:

www.hp.com

Additional backup recommendations and information is available in the Backup whitepapers, also available at the HP 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 HP 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 HP 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

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.

HP 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 volume, it is not possible to directly take advantage of multiple tape drives. If possible, make multiple, smaller volumes. 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 volume, the administrator configures several backup groups to contain the various directory trees at the root, so that more than one tape device can work in parallel.

Also, note how the volumes are constructed when setting up backup jobs. To increase the performance of the backups, schedule the back up of volumes 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 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. 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
 - Member storage units (LUNs)
 - 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

Persistent Storage Manager (PSM) provides instant data recovery from hundreds of online snapshots. Once the first snapshot is taken PSM monitors all drive activity retaining the deleted data required to recreate the last snapshot. Individual files, groups of files, folders, groups of folders or complete volumes can be restored. Security rights and privileges, as well as file and directory attributes, remain in effect as they were at the time the snapshot was created.

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.



Caution: 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 data. Snapshots will be automatically deleted without warning by PSM 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 volume.

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 and logical drives 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.

PSM Error Codes



If you experience a problem using Persistent Storage Manager, the following list of event log messages can be used to troubleshoot. Error codes are logged to the system event log by the file system driver for Persistent Storage Manager, PSMAN5 driver; each entry appears with "psman5" as the source name.

Table 21: PSM Error Codes

Error Code	Description
0x00000001	An invalid IOCTL was sent to the driver. Action: Save the system eventlog and contact technical support
0x00000002	Device name is not recognized by PSM. Action: Save the system eventlog and contact technical support.
0x00000003	An invalid path was given for the cache file. Explanation: This error will appear if the cache file cannot be created because the cache file drive is not present. Action: Save the system eventlog, contact technical support.
0x00000005	An exception occurred. Action: Save the system eventlog, contact technical support.
0x00000005	You do not have sufficient rights to the cache file directory. Action: Make sure you have full access to the cache file directory
0x00000005	The cache file specified is a directory instead of a file. Action: Give a full path and filename for the cache file.
0x00000005	PSM was told to shut down. Action: Save the system eventlog and contact technical support.
0x00000006	User performing PSM function without opening PSM. Action: Programmatically, PSM must be opened before a command can be submitted.
0x00000015	Access to a virtual volume has been attempted after it has been destroyed. Action: Do not access virtual volumes after they have been destroyed.
0x00000016	Something has gone wrong with PSM. Action: Save the system eventlog and contact technical support.
0x00000017	Bad sector was detected in the cache file. Action: Save the system eventlog and contact technical support.
0x0000001F	General failure. Action: Save the system eventlog and contact technical support.

Table 21: PSM Error Codes

Error Code	Description
0x00000057	An invalid parameter was passed to a function. Action: Programmatically, verify the parameters being passed to PSM are correct.
0x00000079	I/O timed out while reading from the cache file. Action: Verify the hard drive is operational.
0x0000007A	Buffer size supplied is insufficient to hold requested information. Action: Save the system eventlog and contact technical support.
0x000000A1	An invalid path was given for the cache file. Action: Save the system eventlog and contact technical support.
0x000000EA	Buffer size supplied is insufficient to hold requested information. Action: Save the system eventlog and contact technical support.
0x000003E6	An exception occurred. Action: Save the system eventlog and contact technical support.
0x00000456	PSM was stopped because the media of a device being PSM'ed was changed. Action: You can take a new snapshot now
0x0000045D	An error occurred on the device. Action: Save the system eventlog and contact technical support.
0x000005AA	There is insufficient memory available. Action: Close unnecessary applications or add more memory.
0x000006F8	Buffer size supplied is insufficient to hold requested information. Action: Save the system eventlog and contact technical support.
0x000006F8	Invalid buffer address passed for I/O. Action: Save the system eventlog and contact technical support.
0x80000005	Specified buffer size is too low. Action: Save the system eventlog and contact technical support.
0x8000001C	PSM was stopped because the media of a device being PSM'ed was changed. Action: Take a new snapshot.
0xA0000004	The cache file is <x>% full. The oldest snapshot(s) will automatically be deleted at <y>%. Explanation: This is a warning that the cache file size is approaching the threshold at which some snapshots will be deleted automatically to free up some cache file capacity. <x> is the percentage for which the warning message will be generated, and <y> is the percentage which represents the threshold. (By default, these values are 80% and 90%, respectively, and can be modified in Windows 2000 for NAS (Disks/Persistent Storage Manager).) Action: In Windows 2000 for NAS (Disks/Persistent Storage Manager), delete some (non-critical) snapshots before the system does to guarantee that critical snapshots do not get deleted accidentally.
0xC0000001	General failure. Action: Save the system eventlog and contact technical support.
0xC0000002	Function is not yet implemented. Action: Save the system eventlog and contact technical support.

Table 21: PSM Error Codes

Error Code	Description
0xC0000005	An Access Exception occurred. Action: Save the system eventlog and contact your vendor's technical support.
0xC0000008	User performing PSM function without opening PSM. Action: Save the system eventlog and contact technical support.
0xC000000D	An invalid parameter was passed to a function. Action: Save the system eventlog and contact technical support.
0xC000000E	Device name is not recognized by PSM. Action: Save the system eventlog and contact technical support.
0xC0000010	An invalid IOCTL was sent to the driver. Action: Save the system eventlog and contact technical support.
0xC0000013	Access to a virtual volume has been attempted after it has been destroyed. Action: Do not access virtual volumes after they have been destroyed.
0xC000001C	An invalid IOCTL was sent to the driver. Action: Save the system eventlog and contact technical support.
0xC0000022	An access exception occurred. Action: Save the system eventlog and contact technical support.
0xC0000022	You do not have sufficient rights to the cache file directory. Action: Save the system eventlog and contact technical support.
0xC0000023	Specified buffer size is too small. Action: Save the system eventlog and contact technical support.
0xC0000034	Cache file name is invalid. Action: Save the system eventlog and contact technical support.
0xC000003A	An invalid path was given for the cache file. Action: Save the system eventlog and contact technical support.
0xC000003B	An invalid path was given for the cache file. Action: Save the system eventlog and contact technical support.
0xC000003E	Bad sector was detected in the cache file. Action: Save the system eventlog and contact technical support.
0xC0000043	A file cannot be opened because the share access flags are incompatible. Action: This occurs when the very last snapshot is deleted. PSM initializes its files when the last snapshot is deleted. While it is initializing, a new snapshots can not be created. Try again in a few minutes.
0xC000009A	There is insufficient memory available. Action: Save the system eventlog and contact technical support.
0xC00000B5	I/O timed out while reading from the cache file. Action: Save the system eventlog and contact technical support.
0xC00000BA	The cache location must be a file rather than a directory. Action: Save the system eventlog and contact technical support.
0xC00000E8	Invalid buffer address passed for I/O. Action: Save the system eventlog and contact technical support.

Table 21: PSM Error Codes

Error Code	Description
0xC000010A	PSM was told to shut down. Action: Save the system eventlog and contact technical support.
0xC0000184	Something has gone wrong with PSM. Action: Save the system eventlog and contact technical support.
0xC0000185	An error occurred on the device. Action: Save the system eventlog and contact technical support.
0xC0000206	Buffer size supplied is insufficient to hold requested information. Action: Save the system eventlog and contact technical support.
0xE0001001	PSM could not start due to the server being constantly busy for minutes. Action: Take a snapshot when the system demands are lower.
0xE0001002	PSM detected a deadlock. Action: Check what other filter drivers you are running (i.e., virus scanners, etc.) Save the system eventlog and contact technical support.
0xE0001003	Specified volume not active or deleted. Action: Do not delete volumes with active snapshots.
0xE0001004	PSM was specified for a volume that is currently not being PSM'ed. Action: Save the system eventlog and contact technical support.
0xE0001005	Cache file overflow caused all existing snapshots to be deleted. Action: Increase the cache file size in Windows 2000 for NAS (Disks/Persistent Storage Manager), or take/schedule snapshots when fewer users are online.
0xE0001006	The application tried to enable PSM without first calling Psm_Register. Action: Programmatically, a program must register with PSM prior to sending it commands.
0xE0001007	Invalid license code. Action: Contact vendor for a valid license.
0xE0001008	Another application already has PSMed locked exclusively. Action: Save the system eventlog and contact technical support.
0xE0001009	PSM needs to be locked exclusive for this function to work. Action: Save the system eventlog and contact technical support.
0xE000100A	Wrong version of the driver has been loaded on this system. Action: Verify the PSM version, save the system eventlog and contact technical support.
0xE000100B	A reboot is required before PSM can operate. Action: Reboot the machine, and try taking a snapshot again. If this still fails, save the system eventlog and contact technical support.
0xE000100C	PSM is not installed. Action: Save the system eventlog and contact technical support.
0xE000100D	An incompatible DLL from another version of PSM is already loaded. Action: Verify the PSM version, save the system eventlog and contact technical support.
0xE000100E	Out of memory. Action: Close unnecessary applications or add more memory.

Table 21: PSM Error Codes

Error Code	Description
0xE000100F	Invalid parameter. Action: Save the system eventlog and contact technical support.
0xE0001010	Invalid handle. Action: Save the system eventlog and contact technical support.
0xE0001011	Not implemented yet. Action: Save the system eventlog and contact technical support.
0xE0001012	Object type is not expected object. Action: Save the system eventlog and contact technical support.
0xE0001013	User buffer is not large enough. Action: Save the system eventlog and contact technical support.
0xE0001014	Out of available structures. Action: In Windows 2000 for NAS (Disks/Persistent Storage Manager), delete some snapshots.
0xE0001015	PSM is shutting down. Action: This is not an error but is a status message.
0xE0001016	The device, volume or object does not exist. Action: Verify that the device, volume, or object exists.
0xE0001017	Unsuccessful. Action: Save the system eventlog and contact technical support.
0xE0001018	The device does not have any media loaded. Action: If the snapshot has been deleted, it cannot be accessed
0xE0001019	Object already exists. Action: Save the system eventlog and contact technical support.
0xE000101A	Specified path is a directory and not a file. Action: Provide a full path and filename
0xE000101B	Invalid path was specified. Action: Ensure the CacheFile name is correct
0xE000101C	The static volume was not mounted. Action: Look at the system event log for a warning message (from the PSMAN5 service) whose code should appear this list. The action depends on the message.
0xE000101D	The static volume had errors during mount. Action: Look at the system event log for a warning message (from the PSMAN5 service) whose code should appear in this list. The action depends on the message.
0xE000101E	The static volume could not be found. Action: Save the system eventlog and contact technical support.
0xE000101F	The volume the cache file resides on is out of space. Action: The cache file for each volume resides on the volume itself. Free some space on the volume.
0xE0001020	The volume the cache file resides on was dismounted. Action: The cache file for each volume resides on the volume itself. Do not dismount the volume.

Table 21: PSM Error Codes

Error Code	Description
0xE0001021	The server was shutdown. Action: Do not shut down the machine while snapshots are in progress.
0xE0001022	Unable to create cache file. Action: Save the sysktem eventlog and contact support.
0xE0001023	PSM recovery could not find a snapshot entry. Explanation: A snapshot was lost during the recovery process. It is unknown which snapshot it was. Action: Save the system eventlog and contact technical support.
0xE0001024	PSM recovery could not open the index file. Explanation: All snapshots are corrupt. Action: Save the system eventlog and contact technical support.
0xE0001025	PSM recovery encountered error <x> inserting key (<y>:<z>) into dictionary. Explanation: <x> is the error that occurred and can be found in this list of errors. Action: Look up the error in this list and take the specified action.
0xE0001026	PSM recovery encountered corrupt index sector %2. Explanation: An index entry was found to be corrupt during the last boot. Action: Save the system eventlog and contact technical support.
0xE0001027	A snapshot could not be created due to error 0x<x>. Explanation: <x> is the error that occurred. Action: Look up the error in this list and take the specified action.
0xE0001028	The cache file is <x>% full. Snapshots have been deleted. Explanation: The oldest snapshots have been deleted. Action: In Windows 2000 for NAS (Disks/Persistent Storage Manager), delete snapshots to make sure specific (critical) snapshots are not destroyed by mistake.
0xE0001029	The maximum (<x>) allowed snapshots has been reached. A snapshot was not created. Explanation: PSM cannot create any more snapshots because the configured maximum number of snapshots that PSM can keep concurrently has been reached. Action: In Windows 2000 for NAS (Disks/Persistent Storage Manager) increase the number of snapshots allowed, or edit the schedules to not make so many snapshots.
0xE000102A	The evaluation period has expired. Action: Contact your vendor's technical support for a non-evaluation version.
0xE000102B	There is not enough free cache space to perform the operation. Action: Delete some snapshots to free up some cache space of enlarge the cache file.
0xE000102D	The maximum number of snapshots has been reached. The oldest snapshot was deleted to allow creation of a new snapshot. Action: Increase the maximum snapshot number. This is a status message.

Table 21: PSM Error Codes

Error Code	Description
0xE0001030	Could not dismount volume before starting snapshot restore. The restore operation was canceled. Action: Save the system eventlog and contact technical support.
0xE0001033	An attempt was made to differentiate volumes of unequal length. Action: Save the system eventlog and contact technical support.
0xE0001034	The volume image backup contains one or more corrupt or missing files. Action: Save the system eventlog and contact technical support.
0xE0001036	An exception has occurred. The data contains the exception record. Action: Save the system eventlog and contact technical support.
0xE0001037	Cannot log on to remote server. Action: Save the system eventlog and contact technical support.
0xE0001038	A backup could not be started because a backup was already in progress. Action: None. This is a status message only.
0xE0001039	Canceled by user. Action: None. This is a status message only.
0xE000103A	The restore of the multiple-volume snapshot was disabled. Action: None. This is a status message only.
0xE000103B	The volume does not have enough free cache to perform the restore. Action: Save the system eventlog and contact technical support.
0xE000103C	The restore operation failed. Action: Save the system eventlog and contact technical support.
0xE000103D	Cannot find space to extend cache file because free space detection is disabled. Action: Save the system eventlog and contact technical support.
0xE000103E	Cannot find space to extend cache file because volume contains no snapshots. Action: Save the system eventlog and contact technical support.

index

A

- accelerator read/write ratio [111](#)
- access control list. See ACL
- ACL
 - defined [174](#)
 - translating [203](#)
- ACU
 - accessing [107](#)
 - controller settings [109](#)
 - features [106](#)
 - overview [106](#)
- adaptive load balancing [38](#)
- ADG (Advanced Data Guarding) [61](#)
- ADU (Array Diagnostics Utility) [102](#)
- AFP
 - installing services for AppleTalk [182](#)
 - installing services for Macintosh [182](#)
 - protocol settings [185](#)
 - shares, setting up [183](#)
 - sharing
 - volume mount points [144](#)
- alerts, e-mail, setting up [51](#)
- AppleTalk. See AFP.
- Array Configuration Utility. See ACU
- array controller
 - purpose [56](#)
- arrays
 - configuration methods [69](#)
 - creating new [111](#)
 - defined [56](#)
 - expanding capacity [116](#)
 - horizontal striping [72](#)
 - moving [105](#)
 - NSPOF horizontal configuration [74](#)
 - vertical striping [69](#)
- audience [16](#)
- Authentication software, installing [191](#)

- authorized reseller, HP [20](#)

B

- backup
 - best practices [243](#)
 - hardware options [242](#)
 - mappings [208](#)
 - multiple devices [243](#)
 - readiness testing [246](#)
 - schedules [244](#)
 - snapshots [124](#), [245](#)
 - software options [242](#)
 - solutions [241](#)
- basic disk [63](#)

C

- cache [22](#)
 - changing size [132](#)
 - clearing from system [127](#)
 - deleting images [132](#)
 - full [144](#)
 - size [131](#)
 - usage [131](#)
- caching
 - automatic for documents [179](#)
 - automatic for programs [179](#)
 - manual for documents [179](#)
- capacity utilization [77](#)
- CIFS
 - add a new user or group [179](#)
 - modify access [179](#)
 - protocol settings [184](#)
 - remove access [179](#)
 - share support [175](#)
 - sharing [178](#)
- CIFS administration [146](#)
- Cisco Fast EtherChannel [38](#)

- client groups
 - adding NFS 200
 - deleting NFS 200
 - editing NFS 201
 - managing NFS 199
- configuration
 - recommended RAID methods 76
- controller settings, ACU 109
- conventions
 - document 17
 - equipment symbols 17
 - text symbols 17
- CPQTeam Utility
 - installing 34
 - opening 36

D

- data blocks 56
- data guarding explained 60
- data recovery, PSM 123
- data replication software
 - installing 45
- data striping 56, 58, 69
- date, system, changing 48
- deployment 26
- disaster recovery 246
- disk quotas
 - creating new entries 157
 - deleting entries 158
 - disabling 156
 - enabling 156
 - modifying entries 158
- document
 - conventions 17
 - prerequisites 16
 - related documentation 17
- domain controller
 - configuring 146
- domain environment 27
- drive defragmentation 123
- drive mirroring explained 59
- drive quotas
 - defined 154
 - managing 154
- dynamic disk 63

E

- e-mail alerts, setting up 51
- encoding types 195
- environments
 - domain compared to workgroup 146
 - overview 27
- equipment symbols 17

- error codes 247
- Ethernet copy 96
- Ethernet NIC teams
 - adding 36
 - checking status 41
 - configuring 37
 - configuring properties 39
 - configuring TCP/IP 40
 - renaming the connection 39
 - setting up 34
 - showing connection icon 40
 - troubleshooting 43
- events, SFU, logging 190
- examples, storage planning 83
- expand priority 110
- explicit mapping 206
- explicit mappings 202

F

- fail on fault setting 37
- fault tolerance
 - compromised 104
 - for NIC teams 37
 - methods supported 58
 - most important for configuration 76
- features
 - cache
 - size 22
 - hard drives 22
 - hardware 22
 - memory 22
 - optional 23
 - processor 22
 - redundancy 24
 - software 23
- FEC 38
- File and Print Services for NetWare. See FPNW.
- file level permissions 168
- files, ownership 173
- folders
 - auditing access 171
 - compress tab 165
 - creating new 164
 - creating new share 166
 - deleting 165
 - general tab 164
 - managing 162
 - managing shares for 167
 - modifying properties 165
 - navigating to 163
- FPNW
 - accessing 219
 - described 217
 - installing 218

- FTP
 - protocol settings 184
 - sharing 180
- G**
- getting help 19
- granule size
 - rules 126
 - update utility 126
- group names
 - examples 147
 - managing 147
- groups
 - adding from a domain 154
 - adding local users 153
 - adding to permissions list 169
 - local, adding 152
 - local, deleting 152
 - local, managing 151
 - local, modifying properties 153
 - properties, general tab 153
 - properties, members tab 153
 - removing local users 154
- H**
- hard drives
 - best practices 63
 - in server 22
 - LEDs 100
 - managing 100
 - migration 96
 - moving 104
 - online spares 58
 - physical 55
 - planning, size and type 78
 - RAID 24
 - replacing failed 102
- hardware, backup 242
- help, obtaining 19
- horizontal striping 72
- HP
 - authorized reseller 20
 - OpenView 238
 - storage website 20
 - technical support 19
- HTTP protocol settings 184
- I**
- I/O performance 77
- iLO. See Integrated Lights-Out Port
- image directory 130
- image groups 136
- inactive period 129
- inactive time-out 130
- Insight Manager
 - agent web interface 237
 - console 237
 - defined 24
 - described 237
 - for HP OpenView 238
 - for Tivoli NetView (AIX) 239
- Integrated Lights-Out port
 - accessing NAS B2000 234
 - activating 46
 - configuration 234
 - described 232
 - features 233
 - license key 46
- L**
- LDM (Logical Disk Manager) 63, 64
- LEDs, defining 100
- license key, iLO port 46
- load balancing 38
- localhost 189
- locks, NFS 197
- logging, SFU events 190
- logical drives. See LUNs
- logical storage elements 63
- logs
 - accessing 50
 - audit 50
 - options 50
- LUNs
 - and storage controller subsystems 57
 - cannot be extended 79
 - creating 114
 - expanding 57
 - largest size 57
 - management under Windows Powered 80
 - managing
 - maximum number 57
 - migrating to new RAID level 119
 - sizing 79
- M**
- Macintosh, installing services for 182
- management software, installing software
 - installing management 239
- management, storage 53
- managing system storage 43

- mapping
 - best practices 203
 - data stored 204
 - explicit 206
 - simple 205
- mappings
 - backup and restore 208
 - creating 204
 - explicit 202
 - NFS 202
 - simple 202
 - squashed 203
- media
 - offsite storage 244
 - rotation 244
- memory 22
- Microsoft Services for UNIX. See SFU
- migration
 - backup and restore outline 95
 - departmental 94
 - developing a plan 94
 - hard drive 96
 - performing 95
 - system wide 94
- mount points 144
- mount points, creating 63

N

- NAS B2000
 - defined 22
 - desktop 30
 - hardware features 22
 - restarting 49
 - setup information archive 245
 - shutting down 49
 - software features 23
 - storage capacity 55
 - supported fault tolerance methods 58
 - using iLO to access 234
 - utilities 24
- NAS Data Copy
 - described 45
 - installing 46
- NAS Management Console 31
- NCP
 - adding new user 181
 - creating new share 222, 226
 - properties, modifying 224
 - protocol settings 184
 - removing access 181
 - shares, deleting 223

- shares, modifying properties 229
- sharing 181
- NetWare
 - adding local users 220
 - enabling user accounts 221
 - installing services for 218
 - supervisor account 221
 - volume mount points 144
- Network File System. See NFS.
- network interface controllers 22
- network settings, changing 52
- NFS
 - async/sync settings 197
 - client groups 199
 - adding 200
 - deleting 200
 - editing 201
 - compatibility issues 175
 - deleting shares 193
 - file share, creating 192
 - file sharing tests 210
 - group mappings 202
 - locks 197
 - modifying share properties 193
 - network protocols 188
 - protocol properties settings 196
 - protocol settings 184
 - sharing 180
 - user mapping server 189
 - user mappings 202
- NIC teams. See Ethernet NIC teams
- NSPOF
 - horizontal array configuration 74
- NTFS partition size limit 57

O

- offsite storage, media 244
- online spares 58

P

- partitions 63
- passwords 212
 - modifying local user's 150
 - synchronization
 - advanced settings 213
 - best practices 212
 - customizing 215
 - implementing 213
 - installing 214
 - requirements 213
- performance, snapshots 125

- permissions
 - file level [168](#)
 - list
 - adding users and groups [169](#)
 - removing users and groups [169](#)
 - modifying [169](#)
 - resetting [171](#)
- persistent image. See snapshots.
- Persistent Storage Manager. See PSM
- physical storage best practices [63](#)
- prerequisites [16](#)
- processor [22](#)
- properties
 - editing PSM schedule [134](#)
 - editing snapshot [140](#)
- protocols
 - AFP settings [185](#)
 - CIFS settings [184](#)
 - FTP settings [184](#)
 - HTTP settings [184](#)
 - NCP settings [184](#)
 - NFS properties settings [196](#)
 - NFS settings [184](#)
 - parameter settings [183](#)
 - planning for compatibility [175](#)
 - supported [27](#), [145](#), [183](#)
- PSM
 - accessing [128](#)
 - and creating mount points
 - creating new schedule [133](#)
 - creating snapshots [138](#)
 - data recovery [123](#)
 - deleting schedule [135](#)
 - deleting snapshots [139](#)
 - editing schedule properties [134](#)
 - editing snapshot properties [140](#)
 - elements overview [65](#)
 - error codes [247](#)
 - global settings [129](#)
 - image directory [130](#)
 - image groups [136](#)
 - inactive period [129](#)
 - inactive time-out [130](#)
 - managing snapshots [137](#)
 - overview [121](#)
 - restore defaults [130](#)
 - restore snapshot [142](#)
 - schedules [132](#)
 - storage limitations [128](#)
 - undo snapshot changes [141](#)

- volume configuration settings [131](#)
- volume display [127](#)
- volume settings [130](#)

R

- rack stability, warning [19](#)
- RAID
 - ADG advantages [62](#)
 - ADG disadvantages [62](#)
 - ADG explained [61](#)
 - horizontal striping [72](#)
 - level on server [24](#)
 - migrating LUN to new level [119](#)
 - NSPOF horizontal configuration [74](#)
 - RAID 0 [56](#)
 - RAID 0 advantages [58](#)
 - RAID 0 disadvantages [59](#)
 - RAID 0 explained [58](#)
 - RAID 1 advantages [59](#)
 - RAID 1 disadvantages [60](#)
 - RAID 1 explained [59](#)
 - RAID 1+0 explained [59](#)
 - RAID 5 advantages [61](#)
 - RAID 5 disadvantages [61](#)
 - RAID 5 explained [60](#)
 - storage enclosure options [75](#)
 - summary of methods [62](#)
 - vertical striping [69](#)
- rapid startup utility
 - defined [24](#)
- rebuild priority [110](#)
- redundancy [24](#)
- related documentation [17](#)
- remote access
 - HP OpenView [238](#)
 - iLO port [232](#)
 - Insight Manager [237](#)
 - methods listed [231](#)
 - remote shell daemon [236](#)
 - Telnet Server [235](#)
 - Terminal Services [232](#)
 - Tivoli NetView [239](#)
 - WebUI [232](#)
- remote shell [236](#)
- Remote Shell Service [211](#)
- restarting the server [49](#)
- restore
 - PSM defaults [130](#)
 - snapshot [142](#)

S

- scalability 25
- scheduled shutdown 49
- schedules
 - creating new PSM 133
 - deleting PSM 135
 - editing PSM properties 134
 - PSM 132
- security
 - auditing 171
 - file level permissions 168
 - ownership of files 173
- services for AppleTalk, installing 182
- services for Macintosh, installing 182
- setup
 - completing 34
 - e-mail alerts 51
 - Ethernet NIC teams 34
- SFU
 - commands 211
 - event logging 190
- SFU, described 188
- shares
 - administrative 175
 - AFP
 - CIFS tab 178
 - creating new 166, 176
 - creating new NCP 222, 226
 - deleting 177
 - deleting NCP 223
 - FTP 180
 - managing 174
 - managing for a volume or folder 167
 - modifying NCP properties 224, 229
 - modifying NFS properties 193
 - modifying properties 178
 - NCP 181, 222
 - NFS 180
 - NFS tests 210
 - NFS, creating 192
 - NFS, deleting 193
 - path 167
 - setting up AppleTalk 183
 - standard 175
 - web (HTTP) 180
- shutting down the server 49
- simple mapping 205
- simple mappings 202
- smart switch 37
- snapshots
 - always keep 123
 - attributes 122
 - automated deletion 123
 - backup 124
 - creating 122, 138
 - deleting 139
 - drive defragmentation 123
 - editing properties 140
 - facts 65
 - managing 137
 - maximum number 129
 - performance impact 125
 - read/write 122
 - reading 122
 - read-only 122
 - recovering 125
 - restoring 142
 - undo changes 141
- software
 - backup 242
 - data replication 45
 - installing Authentication 191
 - updating 52
- spare drives
 - size consideration 78
 - use and number 78
- squashed mappings 203
- squashing 189
- storage capacity 55
- storage capacity expansion 97
- storage configuration planning 68
- storage controller subsystems and LUNs 57
- storage enclosures
 - bay configuration 103
 - configuration options 75
- storage management elements 54
- storage planning examples 83
- storage sizing considerations 80
- subfolder, navigating to 163
- symbols in text 17
- symbols on equipment 17
- synchronization 212
- system characteristics, prioritizing 68
- system date, changing 48
- system storage
 - managing 43
- system time, changing 48

T

- tape libraries 243
- TCP/IP, configuring on NIC team 40
- technical support, HP 19
- Telnet Server
 - auditing log 236
 - authentication tab 235
 - configuring 235
 - enabling 235
 - sessions information 236
 - settings 236
 - using 211
- Terminal Services
 - defined 51
 - described 232
 - exiting 51
 - opening 51
 - using 211
- text symbols 17
- time, system, changing 48
- Tivoli NetView (AIX) 239
- troubleshooting
 - PSM error codes 247
 - PSM known issues
 - PSM known issues 144

U

- UNIX
 - authenticating user access 189
 - converting ACL 203
 - group ID 189
 - remote shell 236
 - See also NFS
 - Telnet Server 211
 - user ID 189
 - volume mount points 144
- user access, authenticating 189
- user interfaces 28

users

- adding to permission list 169
- local
 - adding 149
 - deleting 149
 - managing 148
 - modifying properties 150
- names, managing 147
- NetWare
 - adding 220
 - enabling 221

V

- vertical striping 69
- virtual storage 54
- volume settings, PSM 130
- volumes
 - available for snapshots 130
 - creating new share 166
 - creating Novell 217
 - managing shares for 167
 - mount points 144
 - navigating to 163
 - NCP 222
 - planning 64
 - PSM configuration settings 131
 - re-extending 127

W

- warning
 - rack stability 19
 - symbols on equipment 17
- web sharing 180
- websites
 - HP storage 20
- WebUI
 - accessing 28
 - defined 24
 - launching 232
- workgroup environment 27
- worksheet
 - array configuration storage needs 92
 - drive and enclosure requirements 93
 - example array configuration requirements 87
 - example drives required 88
 - example enclosures required 89
 - example storage need 84
 - usable storage need 91

