

administration guide

# hp StorageWorks

## NAS b3000 v2 and e7000 v2

Second Edition (March 2003)

**Part Number:** 326196-001

This guide provides information on performing the administrative tasks necessary to manage the HP StorageWorks NAS b3000 v2 server and the NAS e7000 v2 server. Overview information as well as procedural instructions are included in this guide.



© Copyright 1983-2003 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Development 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.

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 b3000 v2 and e7000 v2 Administration Guide  
Second Edition (March 2003)  
Part Number: 326196-001



- About this Guide. . . . . 11**
- Overview. . . . . 12
  - Intended Audience . . . . . 12
  - Prerequisites . . . . . 12
- Conventions . . . . . 13
  - Document Conventions . . . . . 13
  - Text Symbols . . . . . 13
  - Equipment Symbols . . . . . 13
- Rack Stability . . . . . 15
- Getting Help . . . . . 15
  - HP Technical Support . . . . . 15
  - HP Storage Website . . . . . 16
  - HP Authorized Reseller . . . . . 16
- 1 System Overview . . . . . 17**
- Product Definition and Information . . . . . 18
  - Server Hardware Features . . . . . 18
  - Software Features . . . . . 18
  - Product Information . . . . . 19
    - Product Manageability . . . . . 19
    - Product Redundancy . . . . . 19
- Deployment Scenarios . . . . . 20
- Environment Scenarios . . . . . 21
  - Workgroup . . . . . 21
  - Domain . . . . . 21
- User Interfaces . . . . . 22
  - NAS Server Web-Based User Interface . . . . . 22
    - Menu Tabs . . . . . 23
      - Status . . . . . 23
      - Network . . . . . 23
      - Disks . . . . . 23
      - Users . . . . . 23
      - Shares . . . . . 23
      - Maintenance . . . . . 23
      - HP Utilities . . . . . 23
      - Cluster . . . . . 23
      - Help . . . . . 23

|   |           |
|---|-----------|
| Welcome Screen Contents . . . . .                                       | 24        |
| Quick Start Guide . . . . .   | 24        |
| Rapid Startup Wizard . . . . .  | 24        |
| Set Server Appliance Name . . . . .                                     | 24        |
| Set Administrator Password . . . . .                                    | 24        |
| Set Default Page . . . . .  | 24        |
| NAS Server Desktop . . . . .  | 24        |
| NAS Management Console . . . . .  | 25        |
| HP Network Teaming Setup . . . . .                                      | 25        |
| Install Data Copy . . . . .   | 25        |
| <b>2 Setup Completion and Basic Administrative Procedures . . . . .</b> | <b>27</b> |
| Setup Completion . . . . .  | 28        |
| Setting up Ethernet NIC Teams (Optional) . . . . .                      | 28        |
| Installing the HP Network Teaming Utility . . . . .                     | 28        |
| Opening the HP Network Teaming Utility . . . . .                        | 30        |
| Adding and Configuring NICs in a Team . . . . .                         | 30        |
| Fault Tolerance . . . . .   | 31        |
| Load Balancing . . . . .  | 32        |
| Configuring the NIC Team Properties . . . . .                           | 33        |
| Renaming the Teamed Connection . . . . .                                | 33        |
| Showing a Connection Icon on the Taskbar . . . . .                      | 34        |
| Configuring the TCP/IP Protocol on the New Team . . . . .               | 34        |
| Checking the Status of the Team . . . . .                               | 36        |
| NIC Teaming Troubleshooting . . . . .                                   | 37        |
| Using AutoPath or Secure Path . . . . .                                 | 37        |
| Clustering the NAS Server . . . . .                                     | 38        |
| Managing System Storage . . . . .                                       | 38        |
| Creating and Managing Users and Groups . . . . .                        | 38        |
| Creating and Managing File Shares . . . . .                             | 39        |
| Installing and Configuring Data Replication Software . . . . .          | 39        |
| Activating the iLO Port Using the License Key . . . . .                 | 41        |
| Basic Administrative Procedures . . . . .                               | 41        |
| Setting the System Date and Time . . . . .                              | 42        |
| Shutting Down or Restarting the Server . . . . .                        | 43        |
| Viewing and Maintaining Audit Logs . . . . .                            | 44        |
| Using Terminal Services . . . . .                                       | 45        |
| Improper Closure of Terminal Services . . . . .                         | 45        |
| Setting up E-mail Alerts . . . . .                                      | 46        |
| Updating the Software . . . . .   | 46        |
| Changing System Network Settings . . . . .                              | 46        |
| <b>3 Storage Management Overview . . . . .</b>                          | <b>49</b> |
| Storage Management Process . . . . .                                    | 49        |
| Storage Elements Overview . . . . .                                     | 51        |
| Physical Hard Drives . . . . .  | 51        |
| Arrays . . . . .  | 51        |
| Logical Drives (LUNs) . . . . .   | 52        |
| Fault-Tolerance Methods . . . . .                                       | 54        |
| RAID 0—Data Striping . . . . .  | 54        |



|  |           |
|--|-----------|
| Advantages . . . . .   | 54        |
| Disadvantages . . . . .  | 54        |
| RAID 1—Drive Mirroring . . . . .                                     | 55        |
| Advantages . . . . .   | 55        |
| Disadvantages . . . . .  | 56        |
| RAID 5—Distributed Data Guarding . . . . .                           | 56        |
| Advantages . . . . .   | 57        |
| Disadvantages . . . . .  | 57        |
| RAID ADG—Advanced Data Guarding and RAID 5DP—Double Parity . . . . . | 57        |
| Advantages . . . . .   | 58        |
| Disadvantage . . . . .   | 58        |
| Online Spares . . . . .  | 59        |
| Physical Storage Best Practices . . . . .                            | 59        |
| Logical Storage Elements Overview . . . . .                          | 59        |
| Partitions . . . . .   | 59        |
| Volumes . . . . .  | 60        |
| Utilizing LDM Storage Elements . . . . .                             | 60        |
| Persistent Storage Management Elements Overview . . . . .            | 61        |
| File System Elements . . . . .                                       | 61        |
| File-Sharing Elements . . . . .                                      | 62        |
| Clustered Server Elements . . . . .                                  | 62        |
| <b>4 Persistent Storage Manager . . . . .</b>                        | <b>63</b> |
| Operational Overview . . . . .                                       | 63        |
| Reading Snapshots . . . . .  | 64        |
| Creating Snapshots . . . . .   | 64        |
| PSM Snapshot Attributes . . . . .                                    | 64        |
| Read Only . . . . .  | 64        |
| Read/Write . . . . .   | 65        |
| Always Keep . . . . .  | 65        |
| Automated Snapshot Deletion . . . . .                                | 65        |
| Data Recovery . . . . .  | 65        |
| File/Folder/Volume Recovery . . . . .                                | 65        |
| Snapshots and Drive Defragmentation . . . . .                        | 65        |
| PSM and Backup . . . . .   | 66        |
| Snapshot Use with Veritas Backup Exec . . . . .                      | 67        |
| Backup Provisions . . . . .  | 67        |
| Restore Provisions . . . . .   | 67        |
| Snapshots Performance Impact . . . . .                               | 68        |
| Recovering Snapshots after a System Restore or System Loss . . . . . | 68        |
| Granule Size Update Utility . . . . .                                | 69        |
| Clearing the Cache File from the System . . . . .                    | 70        |
| Re-extending Volumes from Old Snapshots . . . . .                    | 70        |
| Volume Display in Persistent Storage Manager . . . . .               | 70        |
| Persistent Storage Manager Storage Limitations . . . . .             | 71        |
| Accessing Persistent Storage Manager . . . . .                       | 71        |
| Global Settings . . . . .  | 72        |
| Maximum Persistent Images . . . . .                                  | 72        |
| Inactive Period . . . . .  | 72        |
| Inactive time-out . . . . .  | 73        |

|  |           |
|--|-----------|
| Image directory . . . . .  | 73        |
| Restore Defaults . . . . .   | 73        |
| Volume Settings . . . . .  | 73        |
| Available Volume . . . . .   | 73        |
| Size . . . . .   | 74        |
| Free Space . . . . .   | 74        |
| Cache Size . . . . .   | 74        |
| Usage . . . . .  | 74        |
| Volume Configuration Settings . . . . .                              | 74        |
| Warning threshold reached when . . . . .                             | 75        |
| Begin deleting images when . . . . .                                 | 75        |
| Cache size . . . . .   | 75        |
| Schedules . . . . .  | 75        |
| Creating a New Schedule . . . . .                                    | 76        |
| Editing Persistent Image Schedule Properties . . . . .               | 77        |
| Deleting a Persistent Image Schedule . . . . .                       | 78        |
| Persistent Image and Group Information . . . . .                     | 79        |
| Image name and location on volume . . . . .                          | 79        |
| Persistent image group name . . . . .                                | 79        |
| Number of images in group . . . . .                                  | 79        |
| Volumes included in this image . . . . .                             | 79        |
| Image attributes . . . . .   | 79        |
| Retention weight . . . . .   | 80        |
| Most recent image in group . . . . .                                 | 80        |
| Oldest image in group . . . . .                                      | 80        |
| Next image in group to be deleted . . . . .                          | 80        |
| Managing Persistent Images . . . . .                                 | 80        |
| Creating a New Persistent Image . . . . .                            | 81        |
| Deleting a Persistent Image . . . . .                                | 82        |
| Editing Persistent Image Properties . . . . .                        | 83        |
| Undoing Persistent Image Changes . . . . .                           | 84        |
| Restoring an Image . . . . .   | 85        |
| Known Issues . . . . .   | 87        |
| Event log error at cache full . . . . .                              | 87        |
| Display Error on SAK . . . . .                                       | 87        |
| Always Keep error at cache file full . . . . .                       | 87        |
| Improper display of default Cache File Size . . . . .                | 87        |
| Page file setting . . . . .  | 87        |
| No Boot - No Revert . . . . .  | 87        |
| Reverting of System Drive Prohibited . . . . .                       | 87        |
| No support for mount points in UNIX, AppleTalk, or NetWare . . . . . | 87        |
| <b>5 User and Group Management . . . . .</b>                         | <b>89</b> |
| Domain Compared to Workgroup Environments . . . . .                  | 90        |
| User and Group Name Planning . . . . .                               | 90        |
| Managing User Names . . . . .  | 91        |
| Managing Group Names . . . . .                                       | 91        |
| Workgroup User and Group Management . . . . .                        | 92        |
| Managing Local Users . . . . .                                       | 92        |
| Adding a New User . . . . .  | 93        |

|  |            |
|--|------------|
| Deleting a User . . . . .  | 93         |
| Modifying a User Password . . . . .  | 94         |
| Modifying User Properties . . . . .  | 94         |
| Managing Local Groups. . . . .   | 95         |
| Adding a New Group. . . . .  | 96         |
| Deleting a Group . . . . .   | 96         |
| Modifying Group Properties . . . . .   | 97         |
| General Tab. . . . .   | 97         |
| Members Tab . . . . .  | 97         |
| Drive Quotas. . . . .  | 98         |
| Managing Quotas . . . . .  | 98         |
| Enabling and Disabling Quota Management . . . . .                                | 100        |
| Creating New Quota Entries for a User or Group . . . . .                         | 101        |
| Deleting Quota Entries for a User or Group . . . . .                             | 102        |
| Modifying Quota Entries for a User or Group. . . . .                             | 102        |
| <b>6 Folder and Share Management . . . . .</b>                                   | <b>105</b> |
| Folder Management . . . . .  | 106        |
| Navigating to a Specific Volume or Folder . . . . .                              | 107        |
| Creating a New Folder . . . . .  | 108        |
| Deleting a Folder . . . . .  | 109        |
| Modifying Folder Properties . . . . .  | 109        |
| Creating a New Share for a Volume or Folder . . . . .                            | 110        |
| Managing Shares for a Volume or Folder . . . . .                                 | 111        |
| Managing File Level Permissions . . . . .  | 112        |
| Share Management . . . . .   | 118        |
| Share Considerations . . . . .   | 118        |
| Defining Access Control Lists . . . . .  | 118        |
| Integrating Local File System Security into Windows Domain Environments. . . . . | 119        |
| Comparing Administrative (Hidden) and Standard Shares . . . . .                  | 119        |
| Planning for Compatibility between File Sharing Protocols . . . . .              | 119        |
| NFS Compatibility Issues . . . . .   | 119        |
| Managing Shares . . . . .  | 120        |
| Creating a New Share . . . . .   | 120        |
| Deleting a Share. . . . .  | 121        |
| Modifying Share Properties. . . . .  | 122        |
| CIFS Sharing. . . . .  | 122        |
| NFS Sharing . . . . .  | 124        |
| FTP Sharing . . . . .  | 124        |
| Web Sharing (HTTP) . . . . .   | 124        |
| NetWare Sharing (NCP) . . . . .  | 125        |
| AFP (AppleTalk) Sharing. . . . .   | 126        |
| Installing Services for AppleTalk. . . . .                                       | 126        |
| Installing Windows NT Services for Macintosh . . . . .                           | 126        |
| Protocol Parameter Settings . . . . .  | 127        |
| CIFS Protocol Settings. . . . .  | 128        |
| NFS Protocol Settings . . . . .  | 128        |
| FTP Protocol Settings . . . . .  | 128        |
| HTTP Protocol Settings . . . . .   | 128        |
| NCP (NetWare) Protocol Settings . . . . .  | 128        |

AFP (AppleTalk) Protocol Settings ..... 128

**7 UNIX File System Management ..... 129**

  Network File System ..... 130

  Server for NFS ..... 130

    Authenticating User Access ..... 131

    Indicating the Computer to Use for the NFS User Mapping Server ..... 131

    Logging Events ..... 132

    Installing NFS Authentication Software on the Domain Controllers and Active Directory Domain  
    Controllers ..... 133

  NFS File Shares ..... 134

    Creating a New Share ..... 134

    Deleting a Share ..... 135

    Modifying Share Properties ..... 135

    Encoding Types ..... 137

  NFS Protocol Properties Settings ..... 138

    NFS Async/Sync Settings ..... 139

    NFS Locks ..... 139

  NFS Client Groups ..... 141

    Adding a New Client Group ..... 142

    Deleting a Client Group ..... 142

    Editing Client Group Information ..... 143

  NFS User and Group Mappings ..... 144

    Types of Mappings ..... 144

      Explicit Mappings ..... 144

      Simple Mappings ..... 144

      Squashed Mappings ..... 145

    User Name Mapping Best Practices ..... 145

    Creating and Managing User and Group Mappings ..... 146

      General Information ..... 146

      Simple Mapping ..... 147

      Explicit User Mapping ..... 148

      Explicit Group Mapping ..... 149

    Backing up and Restoring Mappings ..... 150

      Backing up User Mappings ..... 151

      Restoring User Mappings ..... 151

  NFS File Sharing Tests ..... 152

  Terminal Services, Telnet Service, and Remote Shell Service ..... 153

    Using Terminal Services ..... 153

    Using Telnet Server ..... 153

    Using Remote Shell Service ..... 153

  Password Synchronization ..... 154

    Password Synchronization Best Practices ..... 155

    Password Synchronization Requirements ..... 155

    Implementing Password Synchronization ..... 155

    Configuring Advanced Settings ..... 156

    Installing Password Synchronization on Domain Controllers and Active Directory Domain Controllers ..... 156

    Customizing Password Synchronization ..... 157

|          |   |            |
|----------|---|------------|
| <b>8</b> | <b>NetWare File System Management</b> .....                             | <b>159</b> |
|          | Installing Services for NetWare .....                                   | 160        |
|          | Managing File and Print Services for NetWare .....                      | 161        |
|          | Creating and Managing NetWare Users .....                               | 162        |
|          | Adding Local NetWare Users .....  | 162        |
|          | Enabling Local NetWare User Accounts .....                              | 163        |
|          | Managing NCP Volumes (Shares) .....                                     | 164        |
|          | Creating and Managing NCP File Shares Using the WebUI .....             | 164        |
|          | Creating a New NCP Share .....  | 164        |
|          | Deleting an NCP Share .....   | 165        |
|          | Modifying NCP Share Properties .....                                    | 166        |
|          | Creating and Managing NCP Shares using the NAS Management Console ..... | 167        |
|          | Creating a New NCP Share using the NAS Management Console .....         | 168        |
|          | Modifying NCP Share Properties using the NAS Management Console .....   | 171        |
|          | NOTES: .....  | 172        |
| <b>9</b> | <b>Cluster Administration</b> .....                                     | <b>173</b> |
|          | Cluster Overview .....  | 175        |
|          | Cluster Terms and Components .....                                      | 176        |
|          | Nodes .....   | 176        |
|          | Resources .....   | 176        |
|          | Virtual Servers .....   | 176        |
|          | Failover .....  | 176        |
|          | Quorum Disk .....   | 177        |
|          | Cluster Concepts .....  | 177        |
|          | Sequence of Events for Cluster Resources .....                          | 177        |
|          | Hierarchy of Cluster Resource Components .....                          | 179        |
|          | Cluster Planning .....  | 180        |
|          | Storage Planning .....  | 180        |
|          | Network Planning .....  | 181        |
|          | Protocol Planning .....   | 182        |
|          | Before Beginning Installation .....                                     | 183        |
|          | HP StorageWorks NAS Software Updates .....                              | 183        |
|          | Checklist for Cluster Server Installation .....                         | 183        |
|          | Shared Disk Requirements .....  | 183        |
|          | Cluster Installation .....  | 184        |
|          | Installation Overview .....   | 184        |
|          | Setting Up Networks .....   | 184        |
|          | Configure the Private Network Adapter .....                             | 185        |
|          | Configure the Public Network Adapter .....                              | 185        |
|          | Rename the Local Area Network Icons .....                               | 185        |
|          | Verifying Connectivity and Name Resolution .....                        | 185        |
|          | Verifying Domain Membership .....                                       | 185        |
|          | Setting Up a Cluster User Account .....                                 | 186        |
|          | About the Quorum Disk .....   | 186        |
|          | Configuring Shared Disks .....  | 186        |
|          | Verifying Disk Access and Functionality .....                           | 186        |
|          | Install Cluster Service Software .....                                  | 187        |
|          | Configuring the First Node .....  | 187        |
|          | Configuring Cluster Disks .....   | 189        |

|   |            |
|---|------------|
| Validating the Cluster Installation .....                                       | 192        |
| Configuring the Second Node .....   | 192        |
| Verify Installation .....   | 193        |
| HP StorageWorks NAS Software Updates .....                                      | 194        |
| Configuring Physical Disk Resources after the Cluster Install is Complete ..... | 194        |
| Basic Cluster Administration Procedures .....                                   | 195        |
| Failing Over and Failing Back .....   | 195        |
| Restarting One Cluster Node .....   | 195        |
| Shutting Down One Cluster Node .....  | 196        |
| Powering Down Both Cluster Nodes .....  | 196        |
| Powering Up Both Cluster Nodes .....  | 196        |
| Cluster Groups and Resources, including File Shares .....                       | 197        |
| Cluster Group Overview .....  | 197        |
| Node Based Cluster Groups .....   | 198        |
| Load Balancing .....  | 198        |
| Cluster Resource Overview .....   | 198        |
| File Share Resource Planning Issues .....                                       | 199        |
| Resource Planning .....   | 199        |
| Permissions and Access Rights on Share Resources .....                          | 199        |
| NFS Cluster Specific Issues .....   | 200        |
| Persistent Storage Manager .....  | 200        |
| Non Cluster Aware File Sharing Protocols .....                                  | 201        |
| Pathing Software in a Clustered Deployment .....                                | 201        |
| <b>10 Remote Access Methods and Monitoring .....</b>                            | <b>203</b> |
| Web Based User Interface .....  | 204        |
| Terminal Services .....   | 204        |
| Integrated Lights-Out Port .....  | 204        |
| Features .....  | 205        |
| Security Features .....   | 205        |
| Manage Users Feature .....  | 205        |
| Manage Alerts Feature .....   | 206        |
| Integrated Lights-Out Port Configuration .....                                  | 206        |
| Using the Integrated Lights-Out Port to Access the NAS Server .....             | 207        |
| Telnet Server .....   | 208        |
| Enabling Telnet Server .....  | 208        |
| Configuring Telnet Server .....   | 209        |
| Authentication Information .....  | 209        |
| Auditing Information .....  | 209        |
| Server Settings .....   | 209        |
| Sessions Information .....  | 209        |
| Remote Shell Daemon .....   | 210        |
| HP Systems Management Version 7 .....   | 210        |
| <b>A PSM Error Codes .....</b>  | <b>211</b> |
| <b>Index .....</b>  | <b>219</b> |

## about this guide

This administration guide provides information to help:

- Plan storage configuration
- Set up physical storage
- Set up virtual storage
- Manage users and groups
- Manage folders and shares
- Manage a UNIX® file system
- Manage a NetWare file system
- Remotely access the NAS server

About this Guide topics include:

- [Overview](#), page 12
- [Conventions](#), page 13
- [Rack Stability](#), page 15
- [Getting Help](#), page 15

## 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, consider the items:

- Knowledge of Microsoft Windows NT or 2000 operating systems
- Knowledge of HP hardware
- Location of all documentation shipped with the 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   | <a href="#">Figure 1</a>  |
| Key and field names, menu items, buttons, and dialog box titles                     | <b>Bold</b>   |
| File names, application names, and text emphasis                                    | <i>Italics</i>  |
| User input, command and directory names, and system responses (output and messages) | Monospace font<br>COMMAND NAMES are uppercase monospace font unless they are case sensitive |
| Variables   | <monospace, italic font>  |
| Website addresses   | Underlined sans serif font text:<br><a href="http://www.hp.com">http://www.hp.com</a>       |

## 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 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 Web Based User Interface
  - NAS Desktop

---

**Note:** The NAS 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 b3000 v2 is a business class NAS solution that provides reliable performance, manageability, fault tolerance, and scalable storage.

The NAS e7000 v2 is an enterprise class NAS solution that provides reliable performance, manageability, continuous data availability, and scalable storage through the fusion of NAS and SAN.

## Server Hardware Features

See the *HP StorageWorks NAS b3000 v2 Getting Started Guide* and the *HP StorageWorks NAS e7000 v2 Getting Started Guide* for a listing of hardware specifications.

## Software Features

Advanced features included and supported by the NAS server include:

- Microsoft Windows Powered OS with Service Pack 3
- Microsoft Clustering Support—requires two servers attached to the same SAN Fiber fabric
- HP Array Configuration Utility (ACU) for MSA1000 only
- HP Secure Path 4.0 Workgroup Edition (licensing included)
- HP AutoPath (installed, but licensing is required from HP)
- HP StorageWorks Platform Kits
- HP Insight Manager 7
- Microsoft Services for Macintosh
- Microsoft Services for NetWare
- Microsoft Services for UNIX (SFU)
- NAS Web Based User Interface (WebUI)
  - Rapid Startup Wizard
- HP StorageWorks NAS Data Copy (Trail Version)
- 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 server 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 server 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.
- HP Systems Management 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. HP Systems Management 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.

## Product Redundancy

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

The clustering ability of the NAS device further ensures continuous data availability, because data being processed by one server head will transition over to the other server head in a failover situation.

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

The server contains dual 36 GB 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.

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 embedded Integrated Lights-Out (iLO) 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 makes the port independent of the server and the operating system.

## Deployment Scenarios

Various deployment scenarios are possible. See the *HP StorageWorks NAS b3000 v2 Getting Started Guide* and the *HP StorageWorks NAS e7000 v2 Getting Started Guide* for configurations. Typical application of NAS devices 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 server 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 server 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 server 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 server, Microsoft Terminal Services, and other remote administration methods to configure and administer all aspects of the NAS server.



---

## Environment Scenarios

The NAS server is deployed into one of two security modes:

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

The NAS server uses standard Windows user and group administration methods in each of these environments. For procedural instructions on managing users and groups, see Chapter 5 of this guide.

Regardless of the deployment, the NAS server 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.

---

**Note:** In a clustered deployment, the clusters must be members of a domain. Therefore, workgroup environments are supported only in non-clustered deployments.

---

## Domain

When operating in a Windows NT or Active Directory domain environment, the NAS server 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 server obtains user account information from the domain controller when deployed in a domain environment. The NAS server itself cannot act as a domain controller.

## User Interfaces

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

- NAS server WebUI
- NAS server 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 Server Web-Based User Interface

The WebUI provides for system administration, including user and group management, share management, and local storage management.

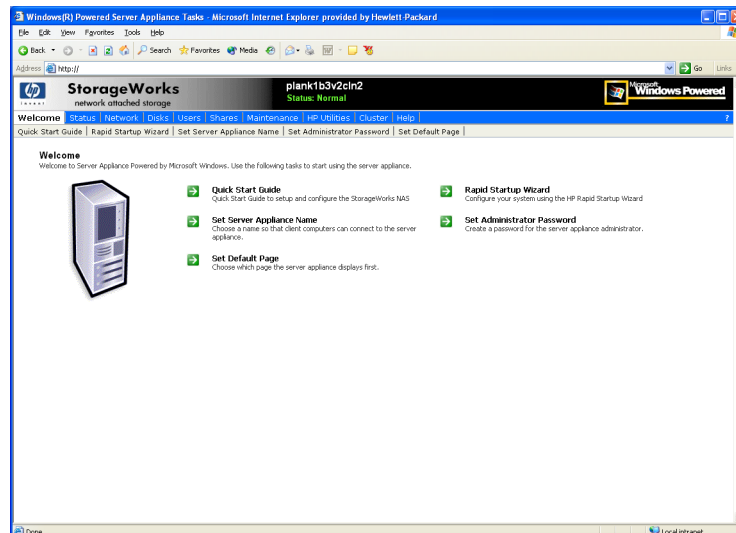
See the *HP StorageWorks NAS b3000 v2 Getting Started Guide* or the *HP StorageWorks NAS e7000 v2 Getting Started Guide* for detailed information on using the Rapid Startup Utility for initial setup.

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/
```

The default user name is: Administrator. The password field should be left blank. 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:

## Menu Tabs

### 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 5.

### 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 6 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.

### Cluster

Use this option to configure and manage the cluster.

---

**Note:** This option is only available on systems that have already established a cluster.

---

### Help

This option contains help information for the WebUI. Page-specific help is also available throughout the WebUI by clicking on the ? in the upper right corner. See [Figure 1](#).

## Welcome Screen Contents

### Quick Start Guide

Use to setup and configure the NAS server.

### 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 Server Desktop


The NAS server desktop can be accessed by:

- Directly connecting a keyboard and mouse
- Using the WebUI Maintenance tab and selecting **Terminal Services**
- Using the embedded iLO port

---

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

---



**Figure 2: NAS server desktop**

The following icons are available from the Desktop:

- NAS Management Console
- HP Network Teaming 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 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 and NFS file shares are managed through this folder.
- **System** contains system summary information.

### **HP Network Teaming Setup**

Click this icon to install the HP 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.

---

**Note:** In a cluster installation, the cluster must be established prior to installing Data Copy.

---



# Setup Completion and Basic Administrative Procedures

## 2

This chapter continues the process of setting up the system that was started using the 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)
  - Using AutoPath or Secure Path
  - Clustering the server
  - 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)
- Clustering the server
- 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 server is equipped with the HP Network Teaming and Configuration utility. The 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 server does not ship with NIC teaming configured.

---

---

**Note:** Installing NIC teaming requires a restart of the server.

---

Procedures include:

- Installing the HP Network Teaming utility
- Opening the HP Network Teaming utility
- Adding and configuring NICs in a team
- Configuring the NIC team properties
- Checking the status of the team
- NIC teaming troubleshooting

### Installing the HP Network Teaming Utility

Before using the HP Network Teaming utility, it must be installed.

---

**Note:** Installing and configuring NIC teaming should always be performed via iLO port or the console using a direct attached keyboard, monitor, and mouse since IP connections could be reset during the configuration process. Do not use Terminal Services.

---



To install the HP Network Teaming utility:

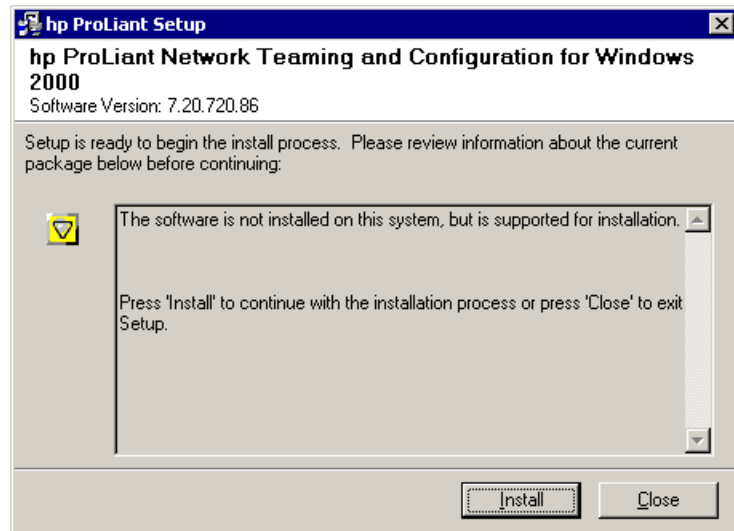
1. From the WebUI, use Terminal Services to go to the NAS server desktop. Double-click the **HP Network Teaming Setup** icon on the desktop.

If the icon is not displayed, enter the following command after selecting Start/Run:

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

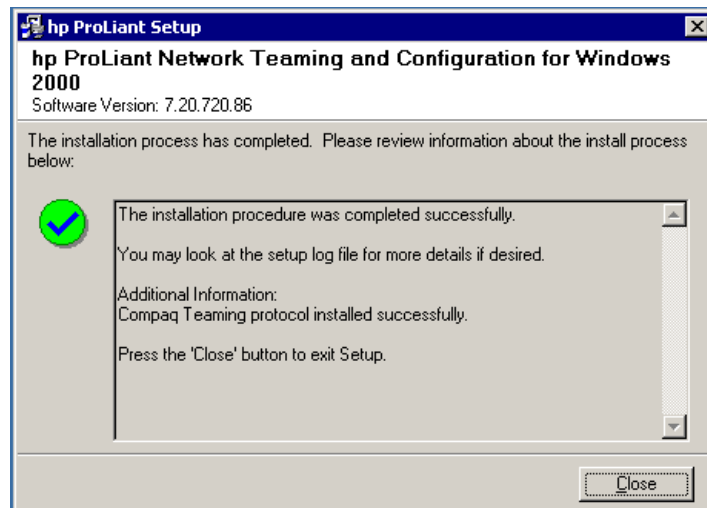
```
c:\winnt\bin\nicteam\JP\cpqsetup.exe (Japanese)
```

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, the server must be restarted at this time.

## Opening the HP Network Teaming Utility

The HP Network Teaming utility is now accessible from the Windows toolbar at the bottom of the NAS server desktop. To open the utility, click the **HP Network Teaming utility** icon.




Figure 5: HP Network Teaming 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 HP Network Teaming utility. The **Network Teaming and Configuration Properties** dialog box is displayed. The type of NIC and the slot and port used is shown.

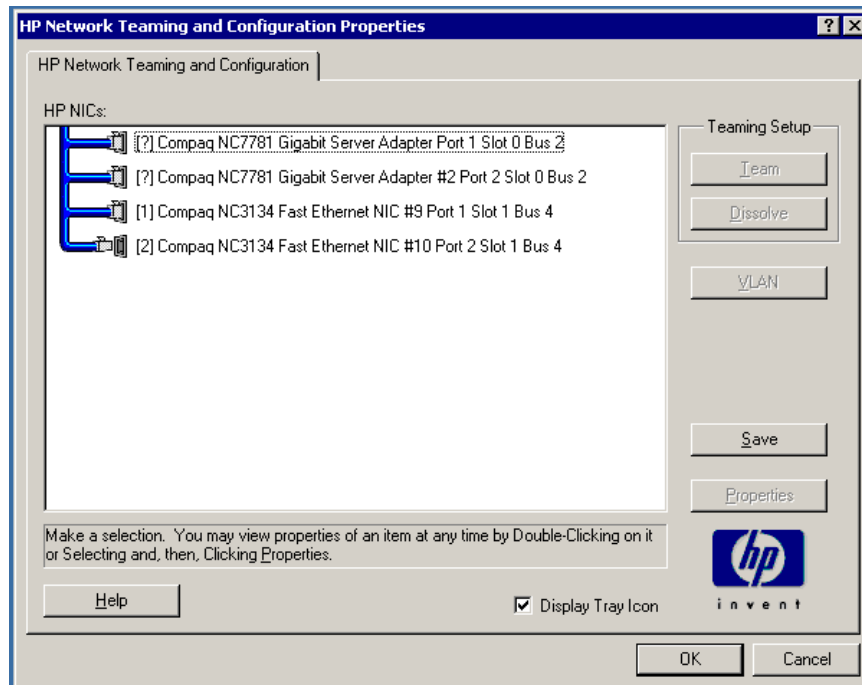
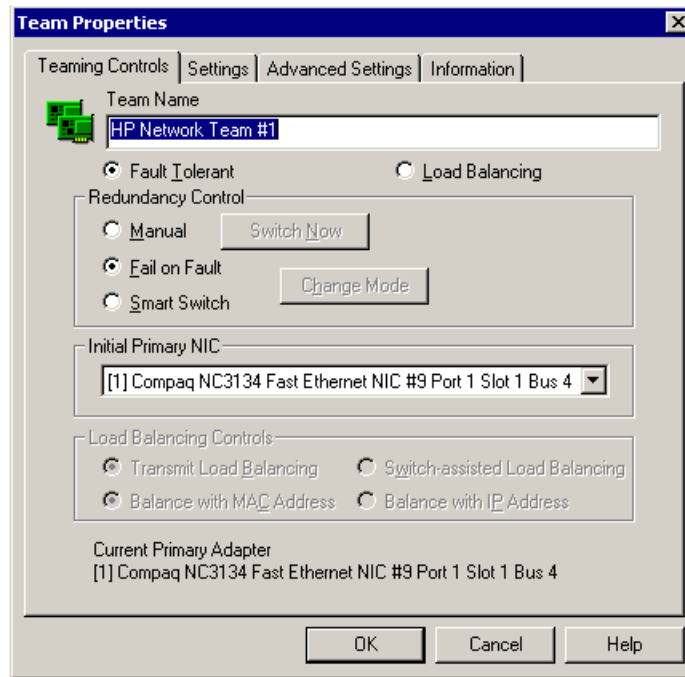


Figure 6: HP Network Teaming Properties dialog box

2. Highlight the NICs to team.
3. 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**

4. 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 allows change from a Primary NIC to a Secondary NIC only when **Switch Now** is clicked.

---

**Note:** The **Switch Now** option is disabled until **Manual** is selected and then **OK** is clicked.

---

- **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 HP Network Teaming 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**

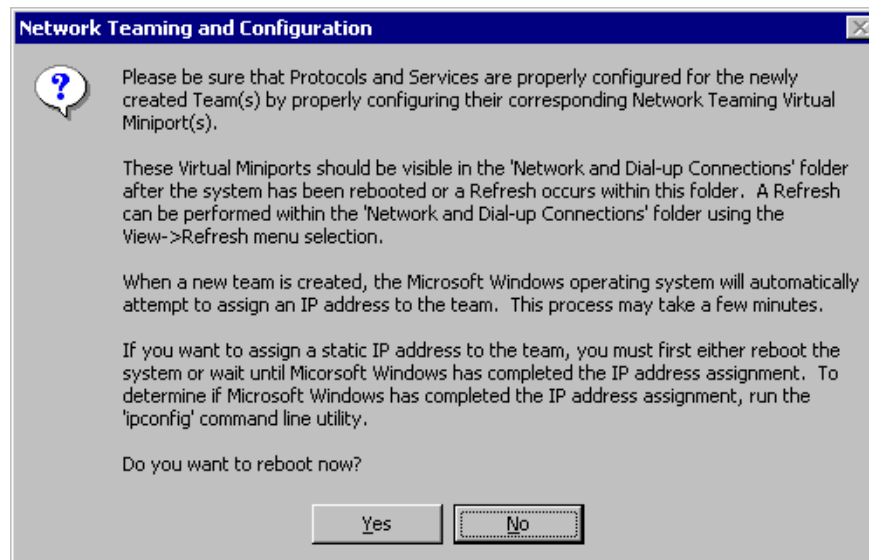
Detailed information about these four load balancing teaming options can be found in the HP Network Teaming Help.

- **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 **OK** to accept the team properties.
6. Click **OK** in the HP Network Teaming and Configuration Properties Screen to apply the changes.
7. Click **Yes** when prompted to apply all configuration changes. Wait while the adapters are configured. This process could take several seconds.
8. The following screen is displayed, indicating that there are additional procedures to perform in the NIC teaming process. Click **Yes** to reboot now.



**Figure 9: HP Network Teaming dialog box**

## 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 **HP Network Teaming Virtual Miniport**.
3. Right-click the connection icon for **HP 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.

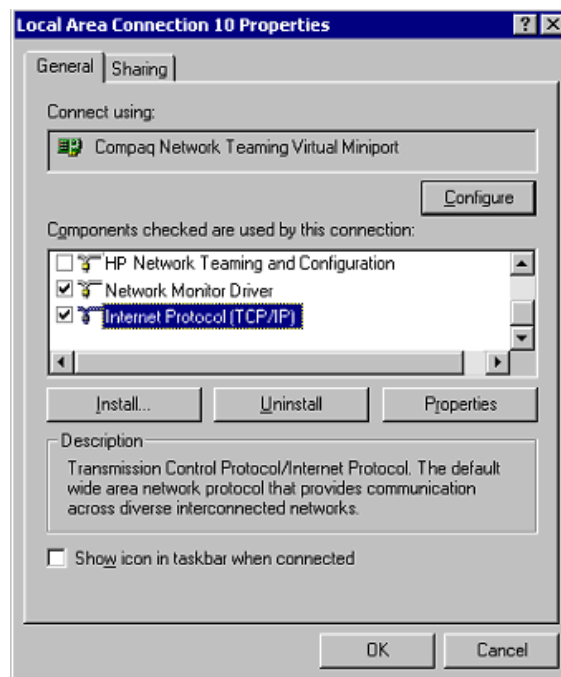
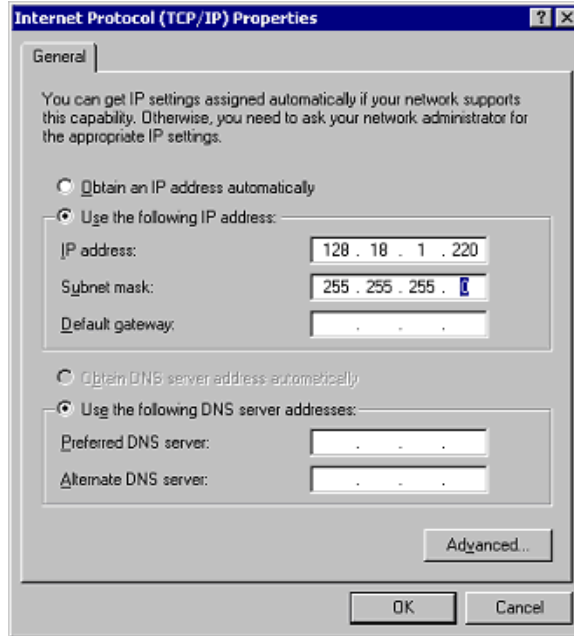


Figure 10: NIC Team Properties dialog box

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 a 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 HP Network Teaming utility. The **Configuration Properties** screen is displayed, showing the teamed NICs.

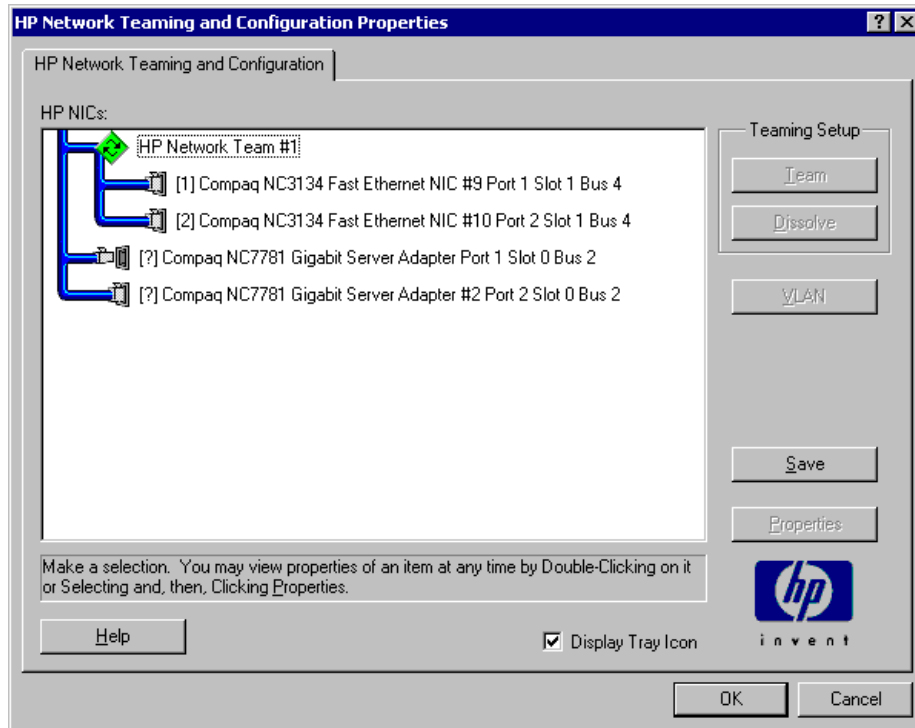









Figure 12: NIC Teaming status



## NIC Teaming Troubleshooting

Problems with the NIC teaming feature are diagnosed by the connection icons displayed in the **HP 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 an 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.<br><b>Note:</b> 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 HP Network Teaming utility.

## Using AutoPath or Secure Path

Pathing software is required in configurations where multipathing to the storage is desired or required. For clustered products it is highly recommended to maintain two paths to the data as path software allows for datapath failure to occur without forcing a node failover. Two products ship with the HP StorageWorks NAS devices; Secure Path and AutoPath. Secure Path is fully licensed and is contained in the shipping product. AutoPath is included as well but a license is required from HP. Both products are installed using the SAN Connection Tool found in the HP Utilities tab of the WebUI.

## Clustering the NAS Server

Many aspects of configuring a NAS device in a clustered configuration are unique to that environment. The cluster management chapter later in this guide provides the details behind this specific configuration and the steps necessary to form a cluster. Throughout the remaining chapters, references to the cluster management chapter are made when special consideration must be applied when utilizing a cluster configuration. Such items include:

- Logical Disk Support
- Lowest Common Unit of Failover
- File Share Protocol Support
- Users and Group Management
- Domain Considerations
- NFS File Share Support
- Persistent Storage Manager

## Managing System Storage

The NAS server 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. See the following chapters for more detailed information on managing system storage:

- Chapter 3 discusses storage management planning in detail.
- Chapter 4 discusses snapshot management procedures.
- Chapter 6 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 5.

The following information is included in Chapter 5:

- 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 6 for complete information on managing file shares.

The following information is included in Chapter 6:

- 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

HP 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 server 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 server, including standalone and clustered device deployments. In clustered environments, NAS Data Copy must be properly configured on each node of the cluster.

See the NAS Data Copy online documentation for configuration and administration of NAS Data Copy.

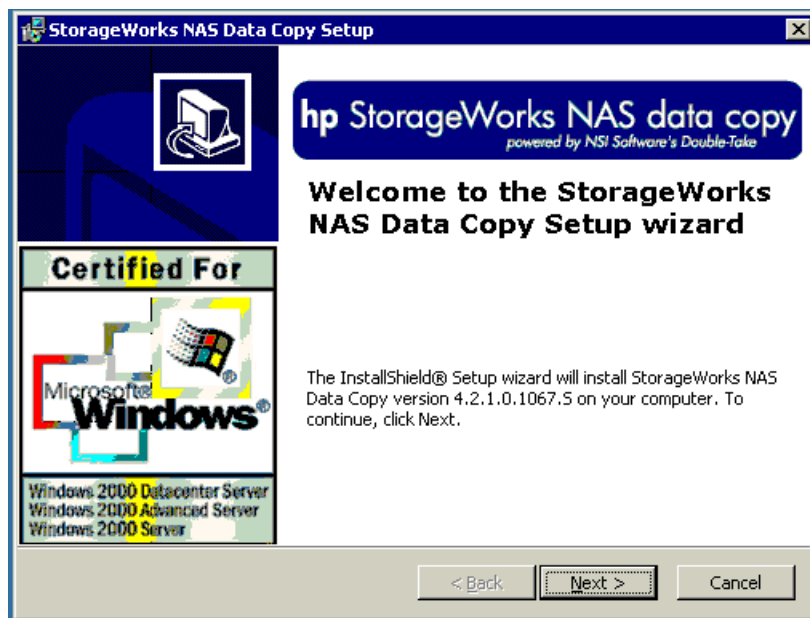
---

**Note:** Establish the cluster before installing data copy.

---

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 server 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 found in the country kit and follow the enclosed instructions.

To access the iLO port, click on the **Remote Management** link in the **HP Utilities** tab located in the WebUI.

## 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.

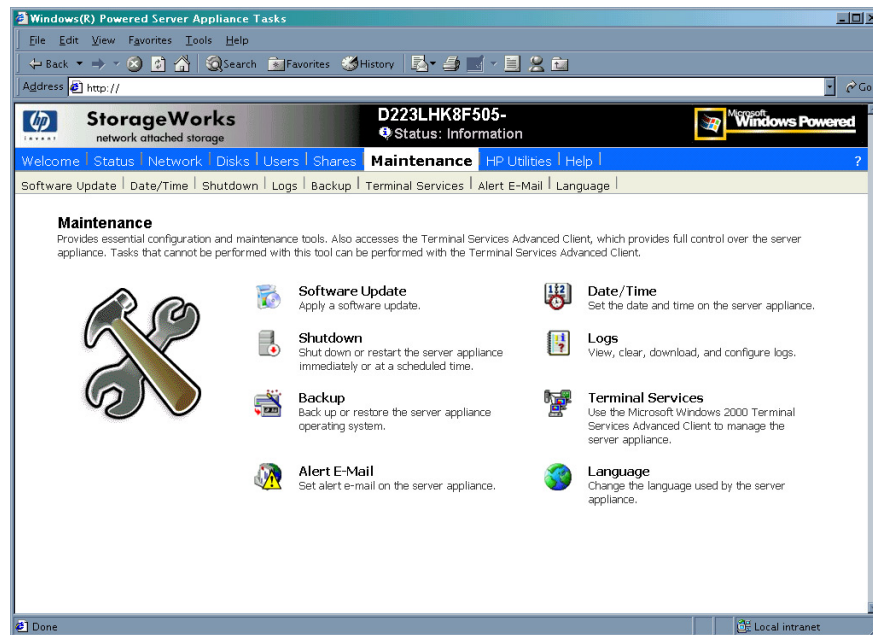


Figure 14: Maintenance menu

## 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.

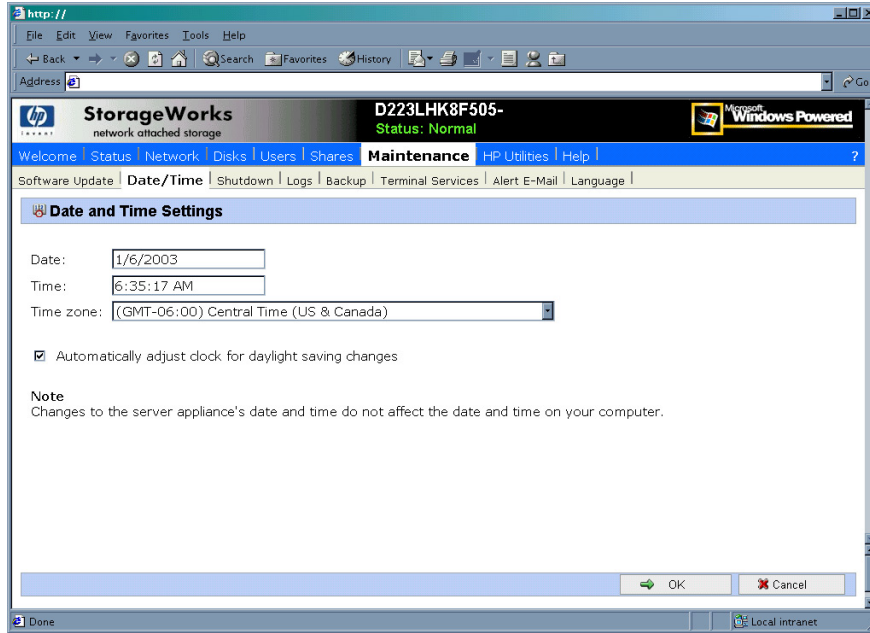


Figure 15: Date and Time dialog box

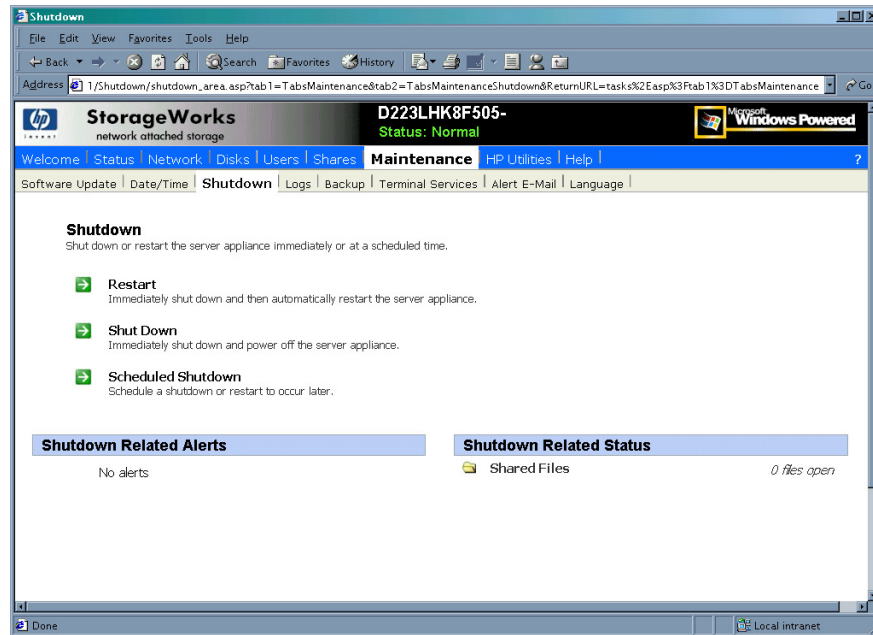
**Note:** In a clustered deployment, be sure to synchronize the time on the nodes.

## 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 server WebUI, select **Maintenance, Shutdown**. Several options are displayed: **Restart**, **Shut Down**, and **Scheduled Shutdown**.



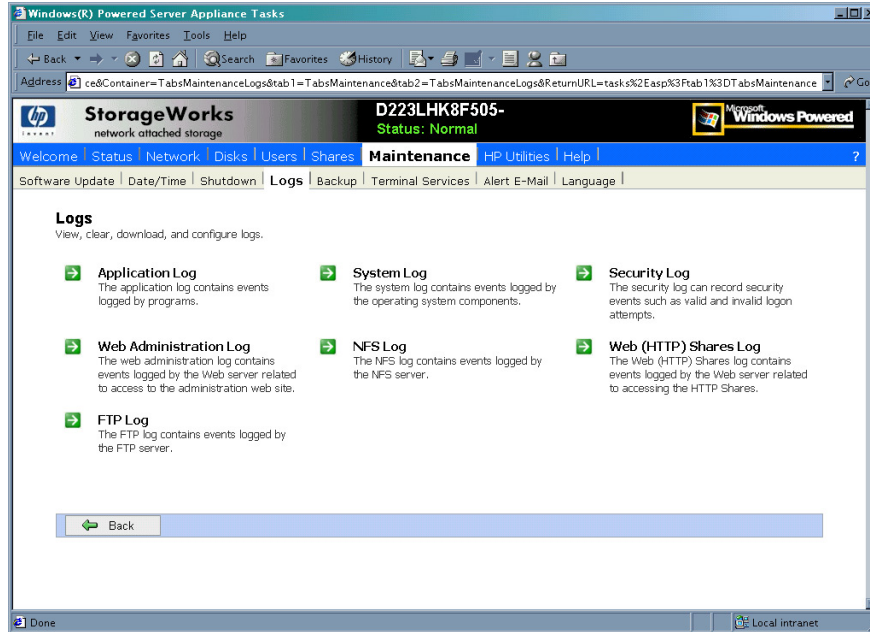
**Figure 16: Shutdown menu**

- 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**.

## Viewing and Maintaining Audit Logs

A variety of audit logs are provided on the NAS server. 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.

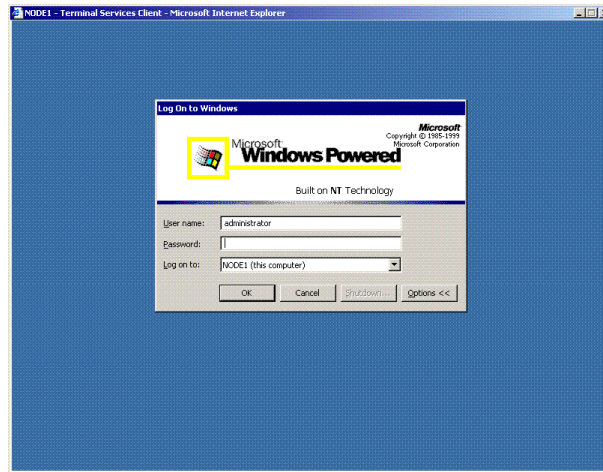


## Using Terminal Services

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 (✕) to close that session of Terminal Services. Click on **Start/Log Off Administrator** to exit Terminal Services.

## Improper Closure of Terminal Services

Certain operations such as drive management via ACU can leave the utilities running if the browser is closed versus exiting from the program via the application menu or logging off the terminal server session. In the case of ACU, the drive lights will remain blinking until one of the following occurs:

1. The ACU is accessed again via the link under **Disk** in the Web User Interface and the application is closed properly.
2. The orphaned Terminal Server session is closed by Terminal Server and the orphaned application is closed. The default timeout has been set to 15 minutes but it may require up to 30 minutes for ACU to exit. This value may be adjusted in the **Terminal Services Configuration** tab under properties of the connection in the HP StorageWorks NAS Management Console accessed either through the desktop or via a Terminal Server session.

Other applications may become orphaned in this manner when the Terminal Server Session is exited improperly. A maximum of two Terminal Server sessions may be used at any given time. Improper exit from a session can result in the sessions becoming consumed. Sessions and processes may be terminated via the Terminal Services Manager via Start -> Programs -> Administrator Tools.

## 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. Use the Software Update Wizard as a guide to select, verify, and update 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 server to a domain.

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

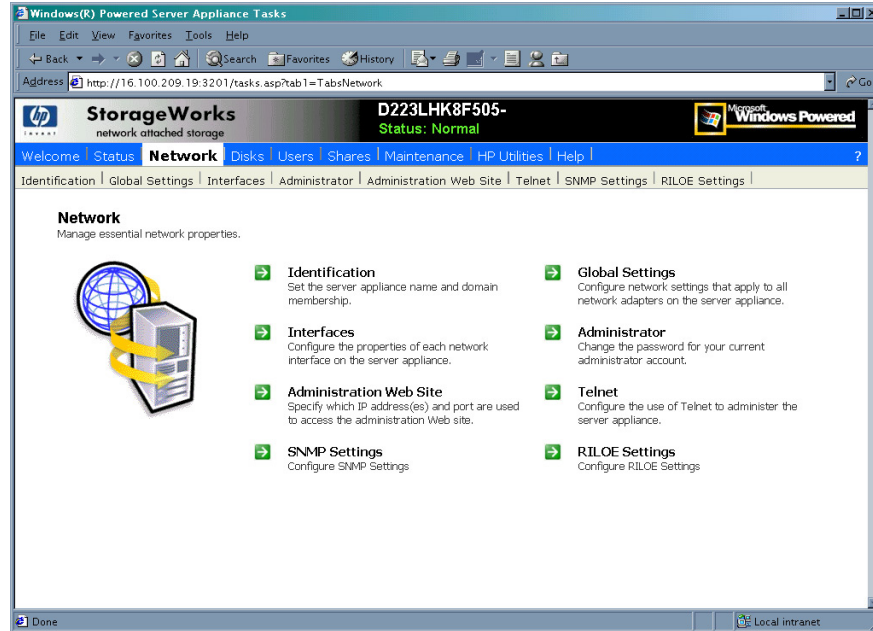


Figure 19: Network menu



# Storage Management Overview

## 3

The NAS server 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.

This chapter defines and discusses physical, logical, and snapshot storage concepts including:

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

Additional storage management information is included in the following chapters:

- Chapter 4 discusses snapshot management procedures.
- Chapter 6 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 Logical 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. Cluster resources use similar constructs to form fault tolerant shares. Lastly, 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. See [Figure 20](#) for an illustration of these storage management elements.

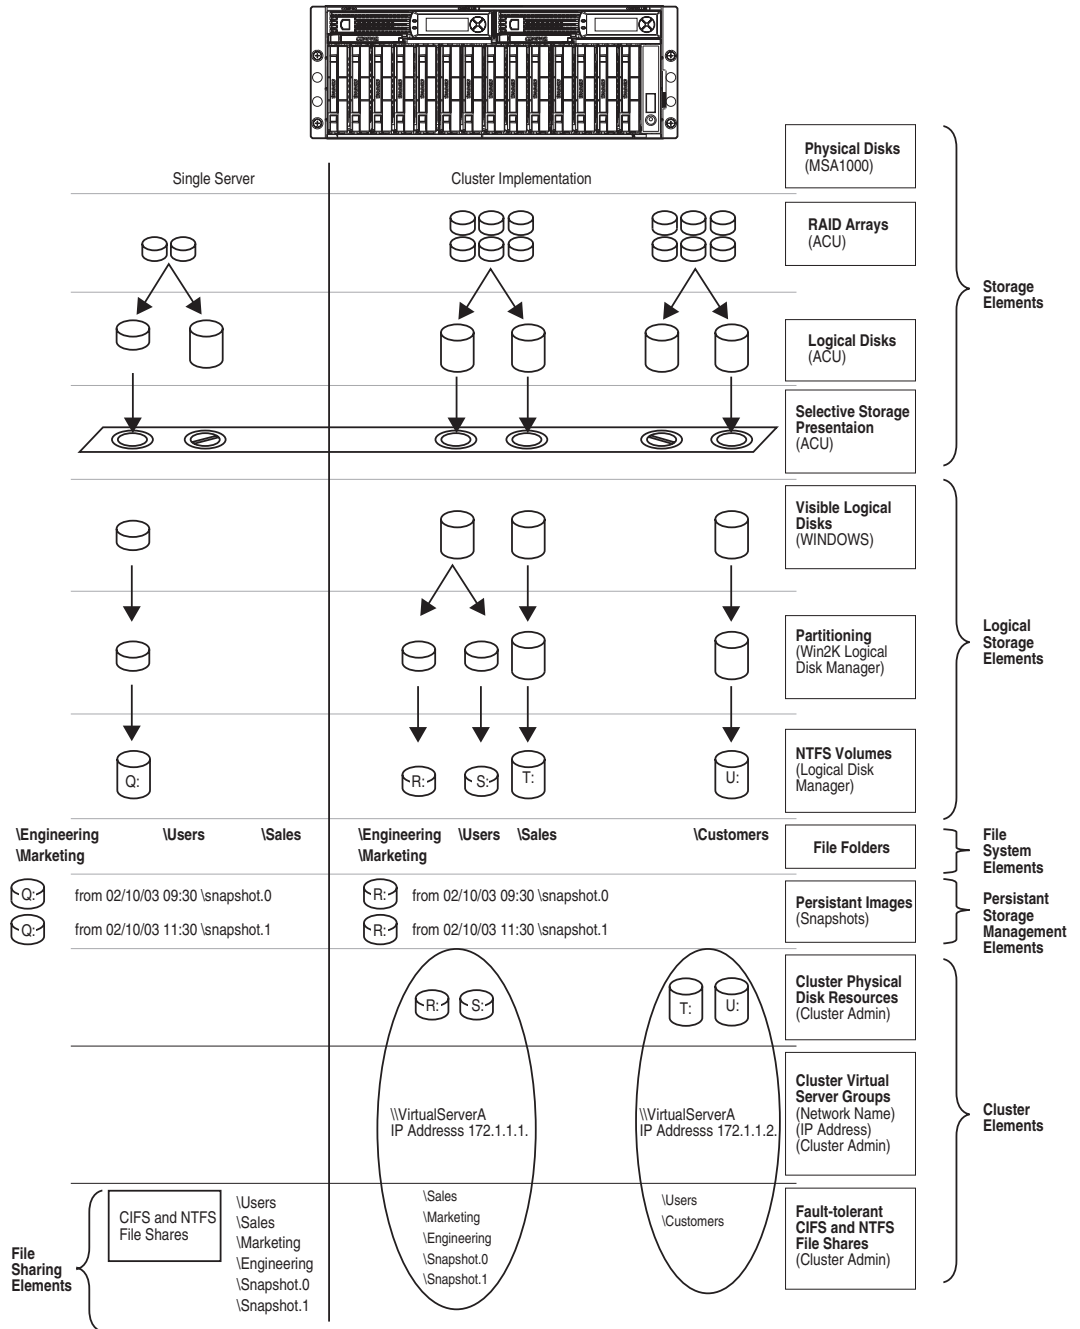


Figure 20: Storage Management process

## Storage Elements Overview

The NAS server offers optimized performance for a growing environment. Storage capacity can increase as a business grows without downtime or compromised performance. Storage limitations are based on the type of SAN the NAS server is connected to. See the individual SAN documentation for limitations of Windows 2000 Advanced Server. 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

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.

## 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. 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 expansion of a LUN. 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 for working with arrays vary from SAN storage to SAN storage system. The individual documentation included with each storage system should be reviewed. HP recommends creating one logical drive from the array.

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

---

**Note:** This type of configuration may not apply to all supported SANs and serves only as an example.

---

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 server 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) See Note below.

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 MSA1000. RAID 5DP is available only with HP Virtual Arrays and is equivalent to ADG.

---

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. 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 media

## 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 on the NAS device 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.

## 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:

- Lower write performance than RAID 0 or RAID 1.
- Increased possibility of data loss if a second drive fails before data from the first failed drive has been rebuilt.

## RAID ADG—Advanced Data Guarding and RAID 5DP—Double Parity


RAID ADG and RAID 5DP are similar to RAID 5 in that parity information is generated (and stored) to protect against data loss caused by drive failure. With RAID ADG and RAID 5DP, 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 27](#), 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 or RAID 5DP is less than when configured with RAID 5.

---

**Note:** The ADG feature is available only with the MSA1000. RAID 5DP is available only with HP Virtual Arrays and is equivalent to ADG.

---



**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 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 disadvantages of RAID ADG are a lower write performance (lower than RAID 5), due to the need for two sets of parity data and less usable space relative to RAID 5.

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<br>Striping (no<br>fault<br>tolerance) | RAID 1 /<br>RAID 1+0<br>Mirroring                                 | RAID 5<br>Distributed<br>Data<br>Guarding | RAID ADG<br>Advanced<br>Data<br>Guarding |
|---|---|---|---|--|
| Usable drive space formula                            | $n$   | $n/2$   | $(n-1)/n$                                 | $(n-2)/n$                                |
| Minimum number of hard drives                         | 1   | 2   | 3   | 4  |
| 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                                      |

## Online Spares

Further protection against data loss can be achieved by assigning an online spare (or hot spare) to any configuration except RAID 0. 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. However, unless RAID ADG is being used that can support two drive failures in an array, in the unlikely event that another drive in the array should fail while data is being rewritten to the spare, the logical drive will still fail.

## 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.

## 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 server 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 and Microsoft Clusters do 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 and Microsoft Clusters do 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.

---

**Note:** Dynamic disks cannot be used for clustered configurations because Microsoft Cluster only supports basic disks.

---

## 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) or Microsoft Cluster. 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 4 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 6 of this guide.

## File-Sharing Elements

The NAS server 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 6 of this guide.

## Clustered Server Elements

The HP StorageWorks NAS b3000 v2 and NAS e7000 v2 support several file-sharing protocols, including CIFS, NFS, FTP, HTTP, NCP, and AppleTalk. Only CIFS, NFS, and FTP are cluster aware protocols. NCP, HTTP, and AppleTalk can be installed on each node but the protocols can not be setup through cluster administrator nor will they failover during a node failure.



**Caution:** AppleTalk shares should not be created on clustered resources as this is not supported by Microsoft Clustering and data loss may occur.

---

Network names and IP address resources for the clustered file-share resource may also be established 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.

---

**Note:** Persistent storage elements can only be viewed from Windows clients and are only supported under file share resources (CIFS).

---

# Persistent Storage Manager

## 4

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

---

**Note:** The NAS b3000 v2 and e7000 v2 servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses using Persistent Storage Manager in a non-clustered environment. Please review the Cluster Chapter of this guide for additional information regarding PSM in a cluster.

---

## 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 “usable 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 reviewing 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.

## Snapshot Use with Veritas Backup Exec

### Backup Provisions

By default PSM snapshot files are registered as files to be ignored by Backup Exec. To back up files using snapshot directories a registry change is required after installation of Backup Exec.

To change the registry:

1. Go to a command prompt and run regedit.
2. Locate the registry entry:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\
FilesNotToBackup
```

3. Remove the entry from the list of files not to backed up which reads Persistent Storage Manager (Images).
4. Back up of snapshots may now be performed. Please keep in mind that a particular instance of a single snapshot should be used in volume backup, for example snapshot.0. Backing up the root snapshot can result in the backup of multiple copies of the volume in use causing the backup to dramatically increase in size.



**Caution:** Using Registry Editor incorrectly can cause serious problems that may require reinstallation of the operating system. Backup the Registry before making any changes. Use Registry Editor at your own risk. For information about how to backup, edit, and restore the Registry in Windows 2000, see Microsoft Knowledge Base Article Q322755.

### Restore Provisions

Special procedures must be followed to enable Veritas Backup Exec to restore the files from a snapshot that was used for backup.

To enable this functionality:

1. Select the files or directories for restore using the Veritas Backup Exec Interface. Record what the snapshot name is from the set that is being restored.
2. Select Redirect to allow the files to be redirected. Select a new directory for restoration. (for example: *e:\vp*)

3. Under the Windows 2000 tab of the restore properties page, locate the Junction Points Control, select the radio button **Preserve Existing Junction Points and restore files and directories from media**.
4. In the directory where the restored files are to be placed, create the directory with the same name as the snapshot name (e:\hp\snapshot.1). The exact directory must be created otherwise Backup Exec will create the directory and make it a junction point that points to the original snapshot. Consequently, the restore operation will fail.
5. Begin the restore job.

## 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 server 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 server 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 4: 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

A potential problem exists with restoring a volume from a snapshot. Dynamic disks may be extended and made larger with LDM. If an extended volume is restored from a smaller snapshot (one created before the volume was extended), the extra space will be unavailable after the snapshot restore. To reclaim the space, run the `reextend.exe` utility after restoring the snapshot.

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 extends 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 server is currently designed to work with 10 TB of storage with the ability to take 250 snapshots. However, the server is capable of addressing greater than 10 TB using a fully populated storage 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.

## Accessing Persistent Storage Manager

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

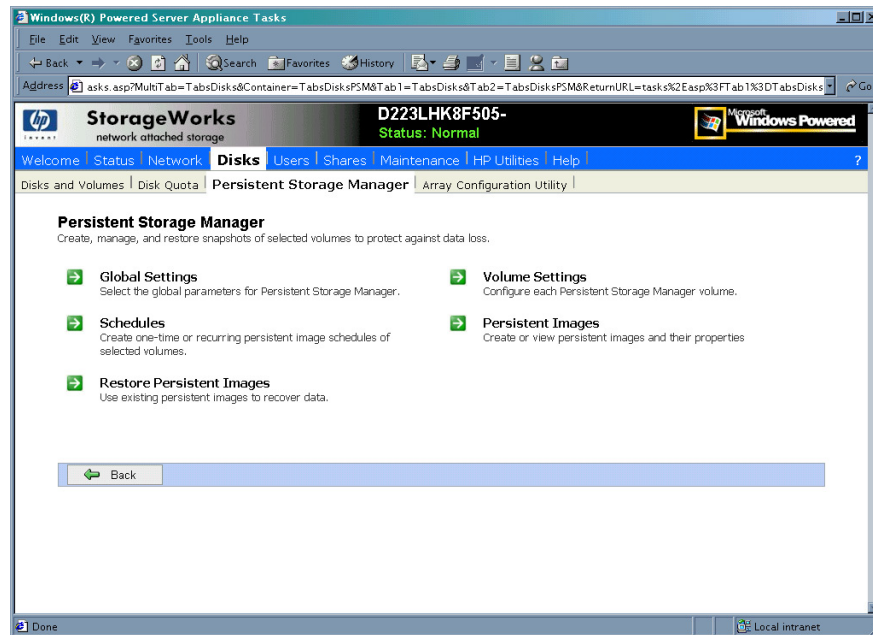


Figure 28: Persistent Storage Manager screen

## Global Settings

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

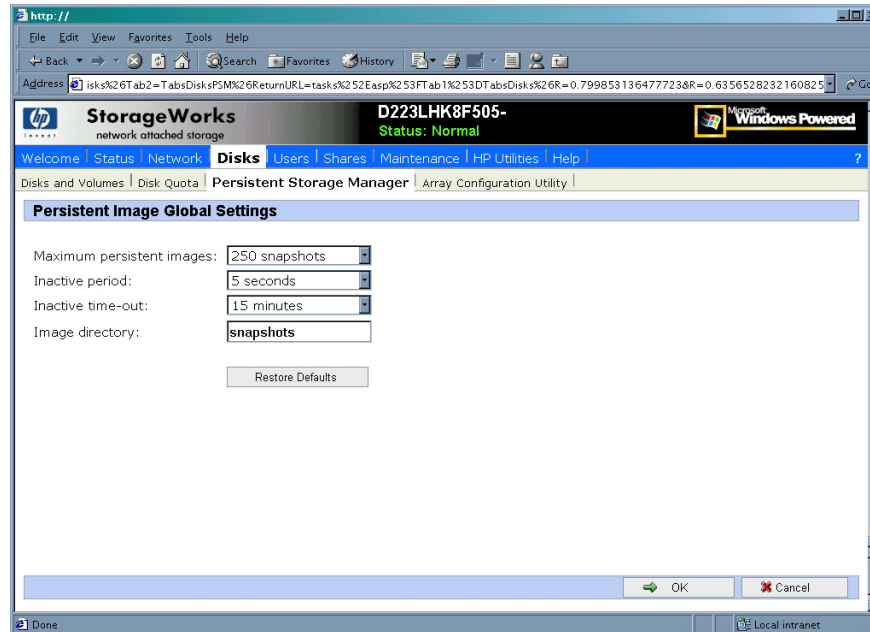


Figure 29: Global settings

## 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 allows snapshots to be created 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 it is possible to view the Persistent Storage Manager attributes for each volume and change volume settings using the **Configure** button in the **Tasks** list.

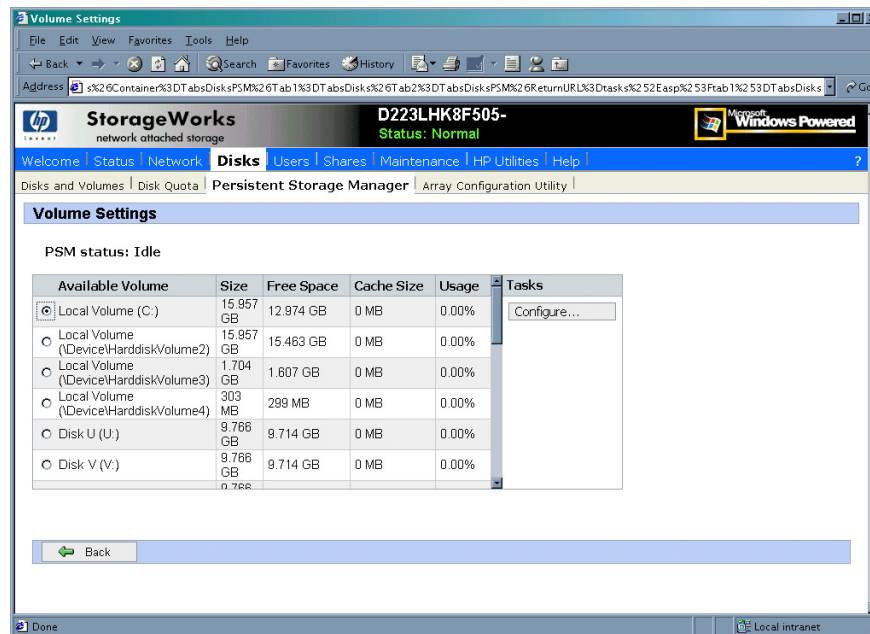


Figure 30: Volume settings

## Available Volume

This field lists all of the volumes that can support snapshots, allowing selection of the desired volume to configure.

## Size

This column displays the size of the volume.

## Free Space

This column displays the available storage size of the volume.

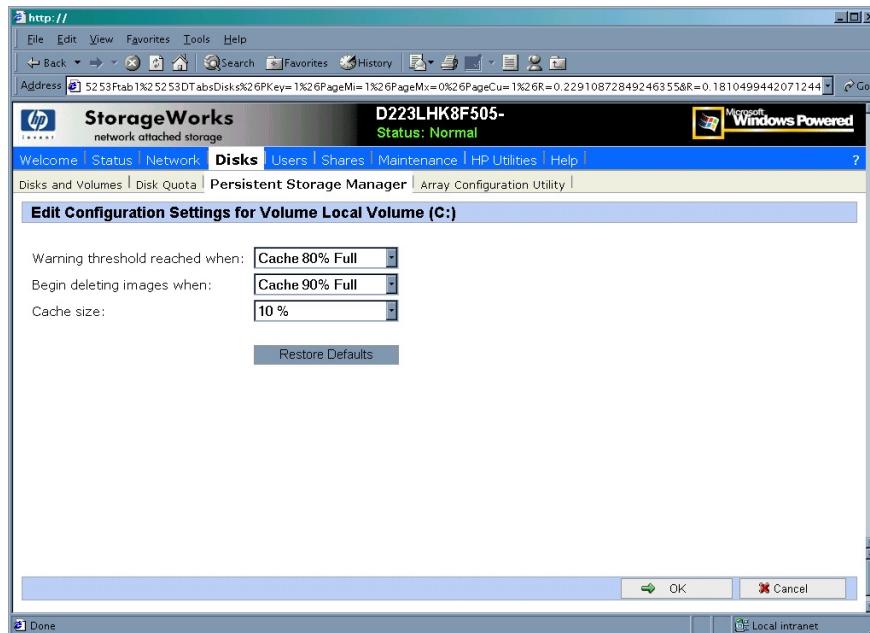
## Cache Size

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 31: Volume configuration settings**

Click **Configure** from the Volume Settings to modify 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. To remove the cache files all together use *CLEANVOL.EXE*; see the section on clearing the cache file. Also see the section on granular size in this chapter, prior to updating the percent reserved for cache size.

**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.

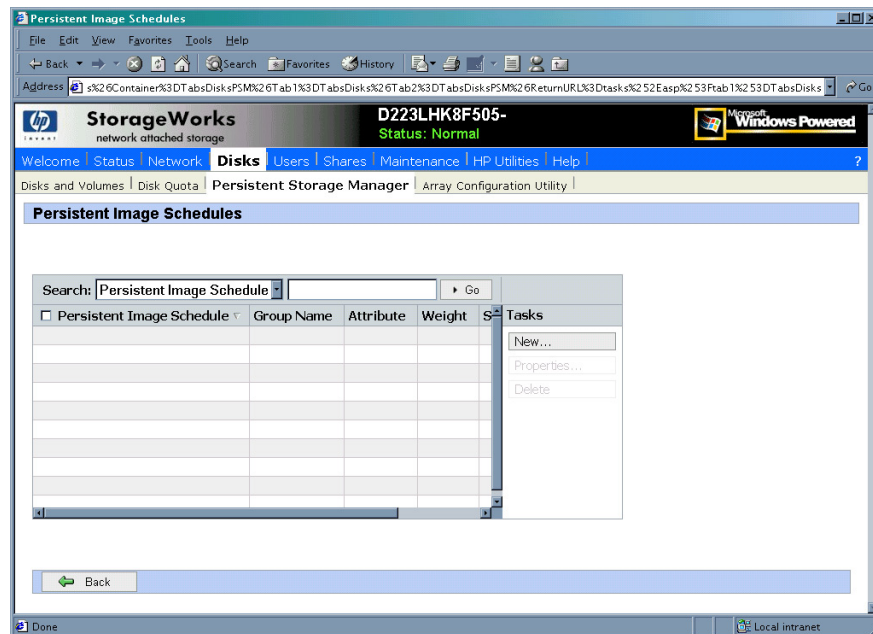
## Begin deleting images when

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. The default value is 10 percent.

## Schedules



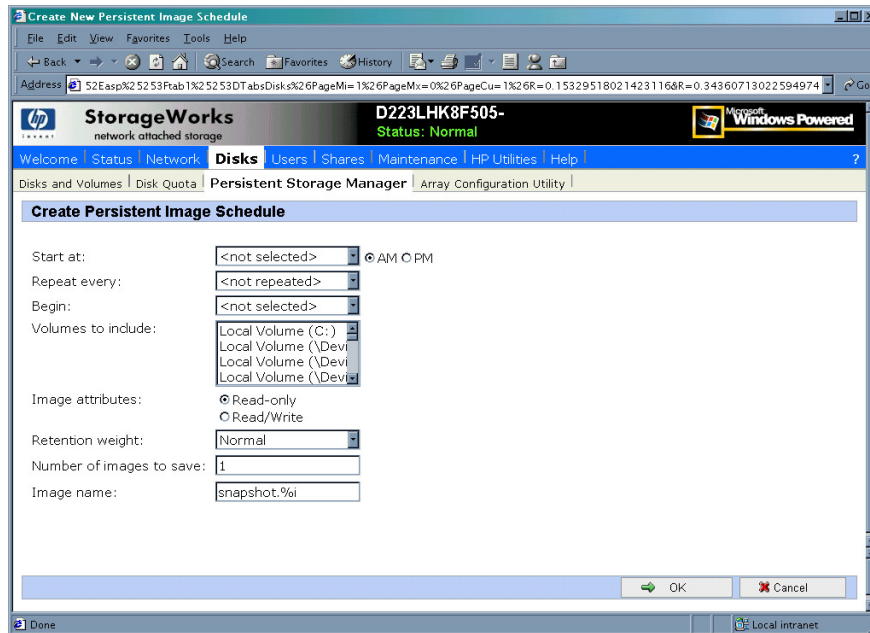
**Figure 32: Persistent image schedules**

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.

Access the Schedules screen to create new schedules, delete existing schedules, and edit schedule properties.

## Creating a New Schedule



**Figure 33: Creating persistent image schedule**

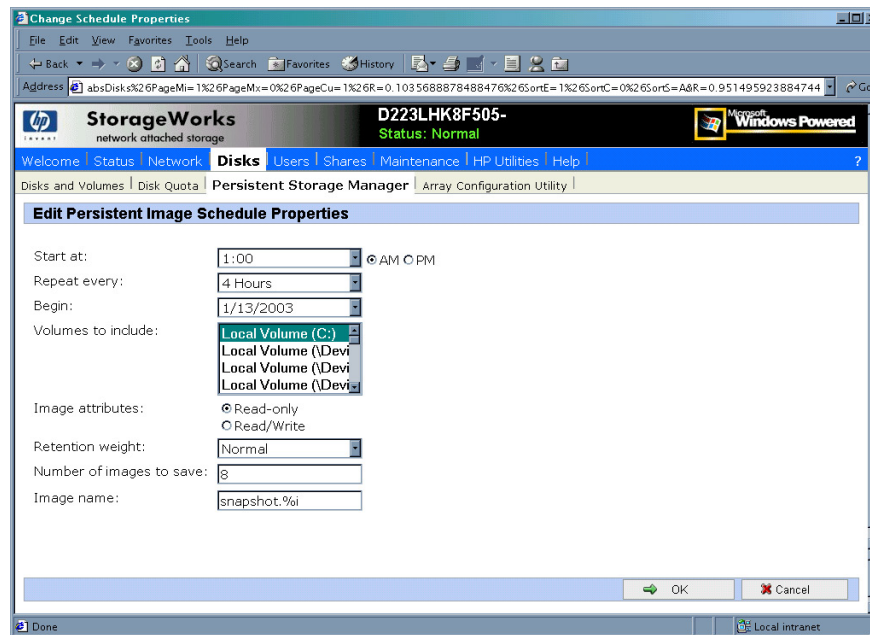
To create a new schedule, a starting time, repeat period, starting day, volume, and the number of snapshots to make available to users must be supplied.

To add a snapshot to the schedule:

1. Select **Schedules** from the **Persistent Storage Manager** screen.
2. In the **Tasks** list, select **New**.
3. Select the parameters for the schedule.
4. Click **OK**.



## Editing Persistent Image Schedule Properties

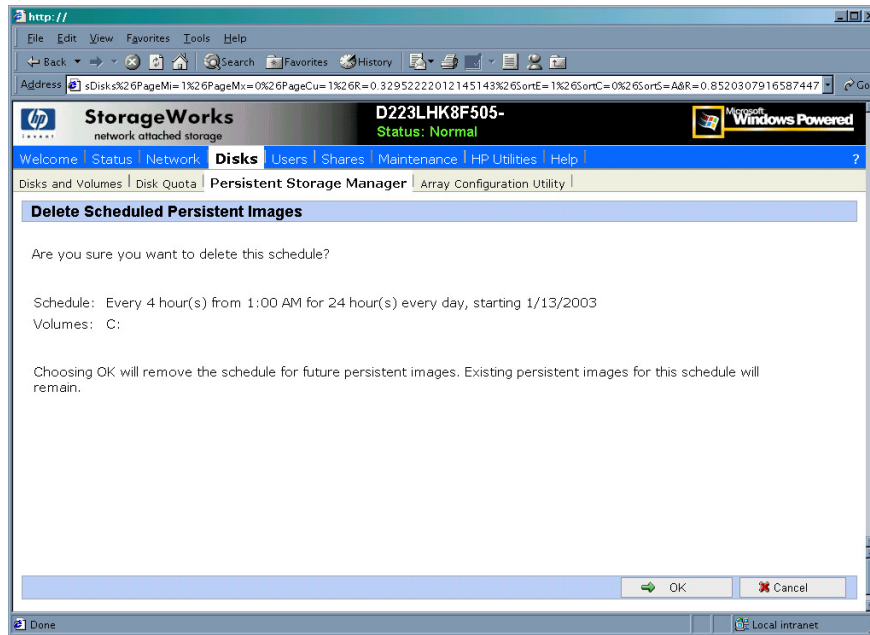


**Figure 34: Editing schedule properties**

To edit persistent image schedule properties:

1. Select **Schedules** from the **Persistent Storage Manager** screen.
2. In the **Tasks** list, select **Properties**.
3. Select the desired schedule changes.
4. Click **OK**.

## Deleting a Persistent Image Schedule

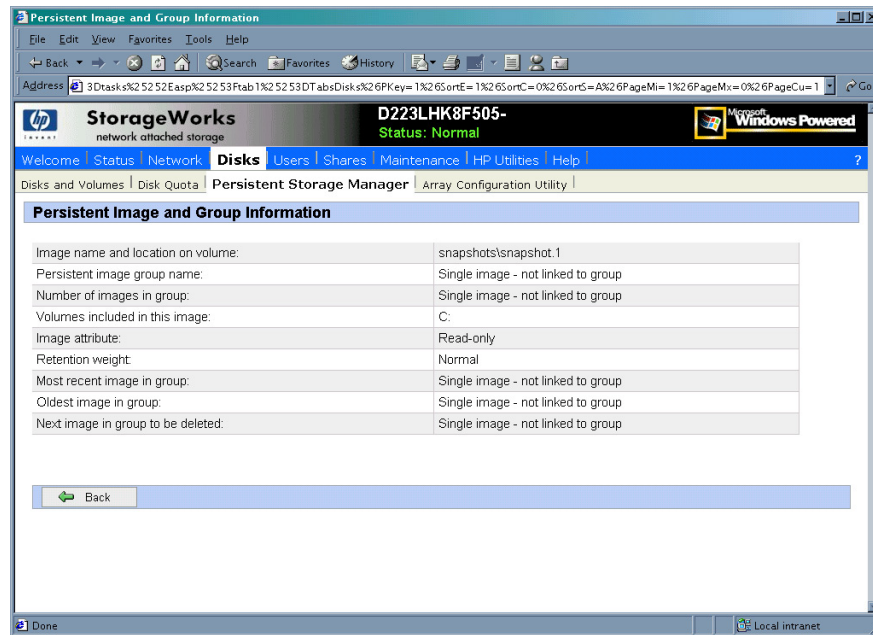


**Figure 35: Deleting scheduled images**

To delete a persistent image schedule:

1. Select **Schedules** from the **Persistent Storage Manager** screen.
2. Select the schedule to delete.
3. In the **Tasks** list, select **Delete**.
4. Click **OK**.

## Persistent Image and Group Information



**Figure 36: Persistent image and group information**

After a snapshot is created from the specified schedule, 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.

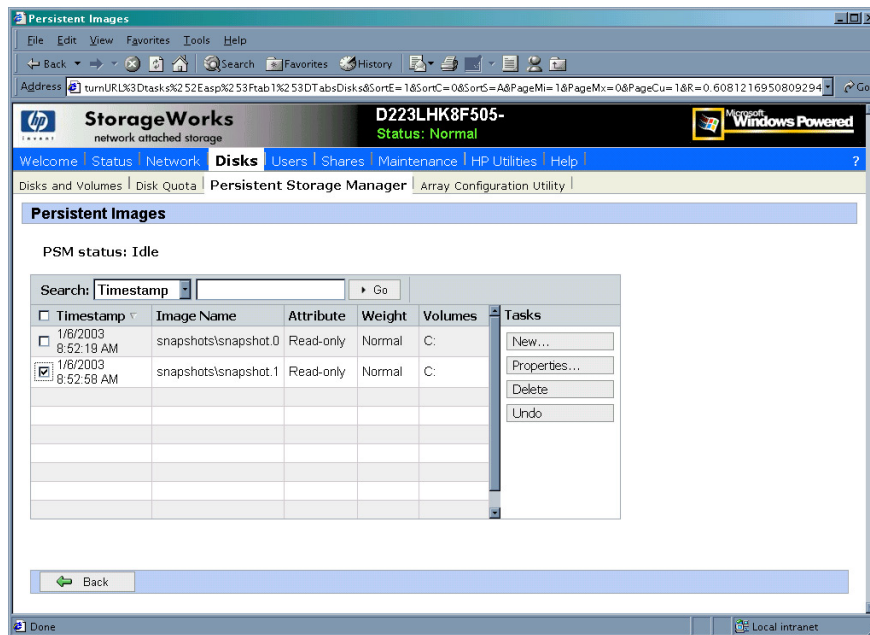
## Oldest image in group

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 37: Managing persistent images**

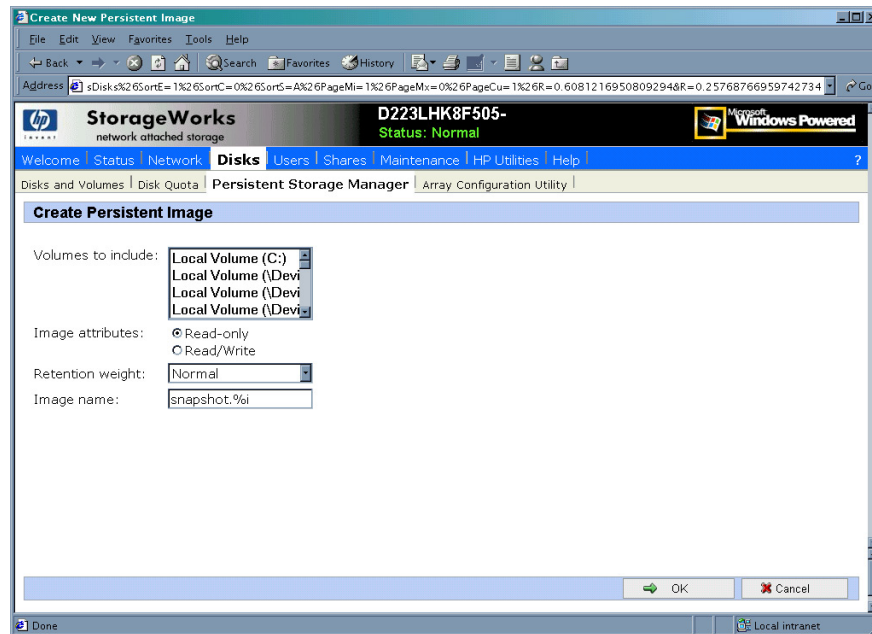
The **Persistent Images** page displays active persistent images. Each entry identifies the date and time the snapshot was created, the read-only or read/write attribute, the preservation weight, and the volume it preserves.

To manage snapshots:

1. From the **Persistent Storage Manager** screen select **Persistent Images**.
2. Select the desired snapshot.

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

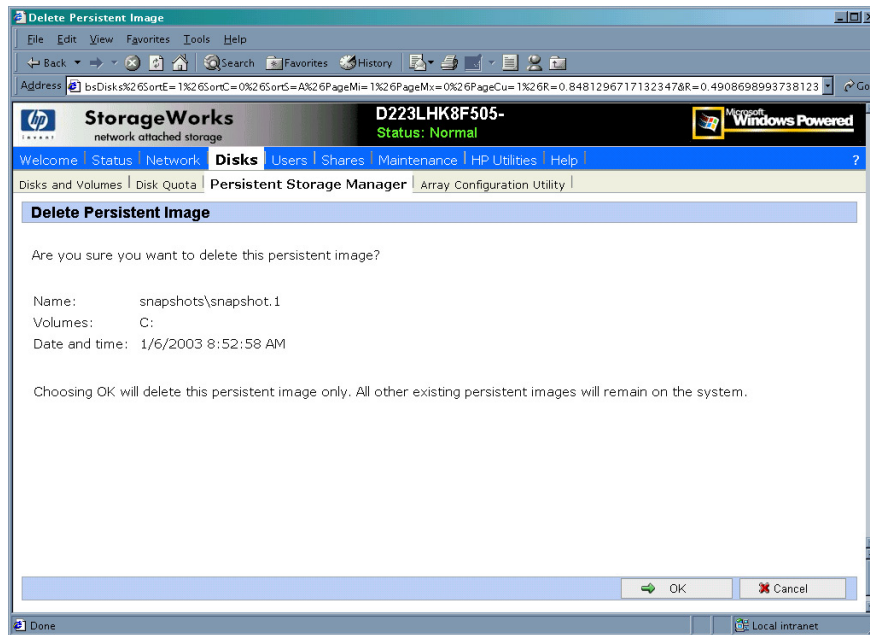


**Figure 38: Creating new persistent image**

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

1. From the **Persistent Storage Manager** 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

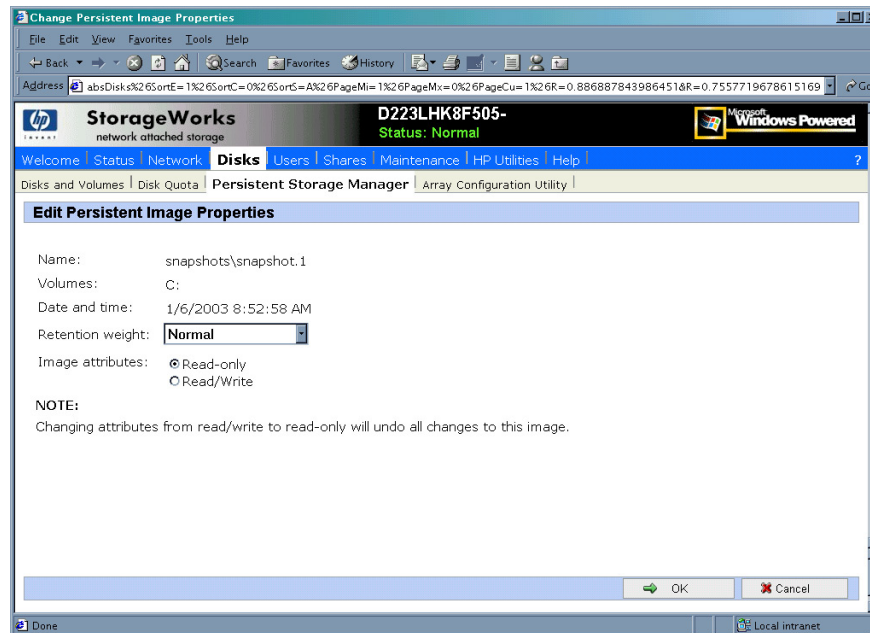


**Figure 39: Deleting verification**

To delete a persistent image:

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

## Editing Persistent Image Properties



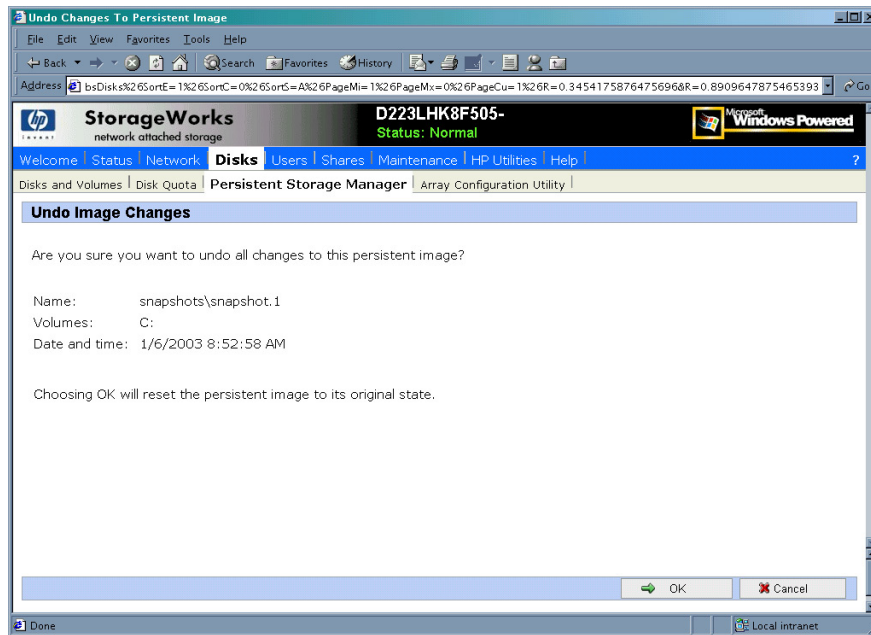
**Figure 40: Editing persistent image properties**

Properties such as the read-only attribute or preservation weight of an image can be changed.

To edit persistent image properties:

1. From the **Persistent Storage Manager** 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**.

## Undoing Persistent Image Changes



**Figure 41: Undoing image changes**

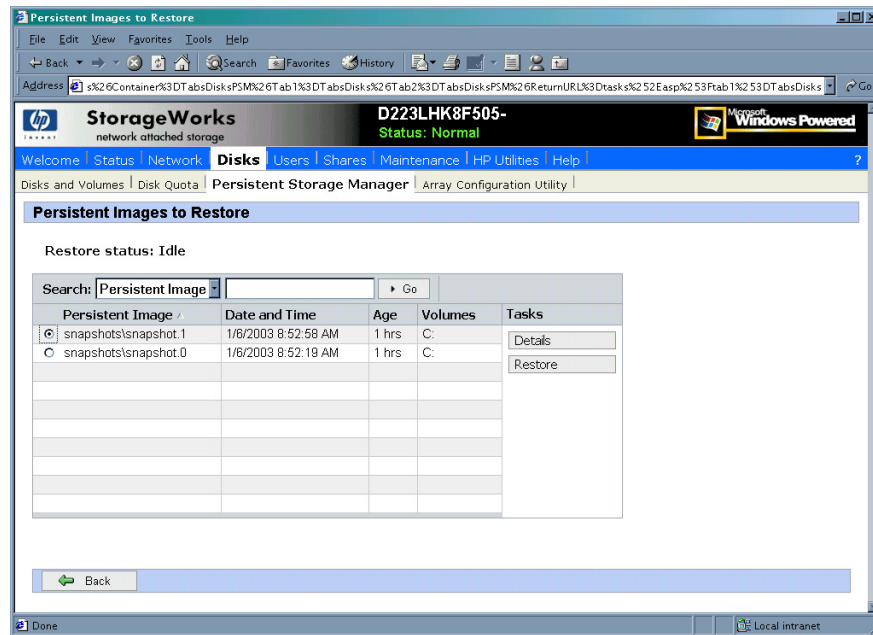
After creating a read/write snapshot, changes can be made to the image, for example, modifying files in the image, adding new files, or deleting existing files. If changes are made to an existing image and later needs to be reverted to the original file contents, use the following procedure to restore the original snapshot.

To undo snapshot changes:

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



## Restoring an Image

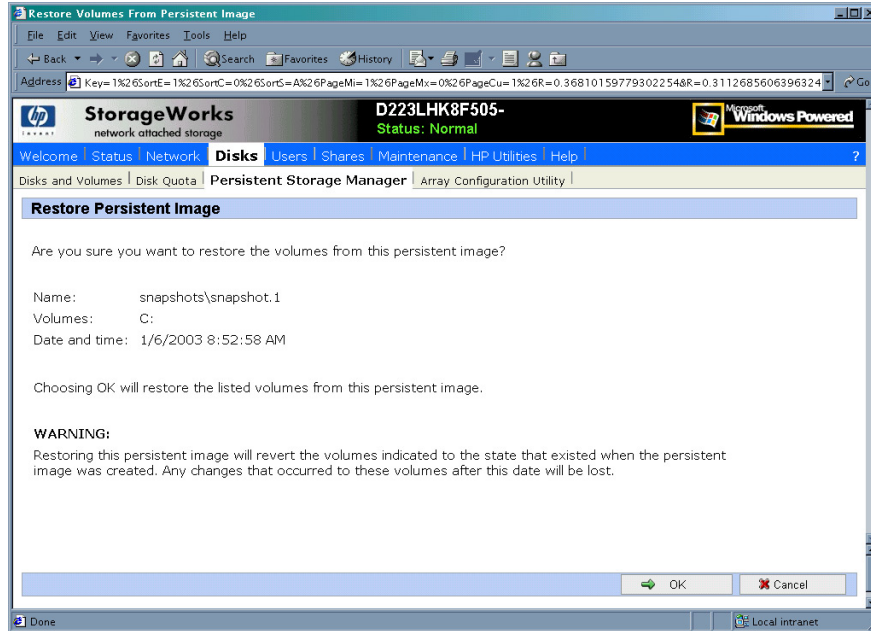


**Figure 42: Images available to restore**

The **Persistent Images to Restore** page displays a list of all snapshots. It is possible to view an image or restore a server appliance to a previously created image.

To restore a snapshot:

1. On the **Persistent Storage Manager** screen select **Restore Persistent Images**.
2. Select the snapshot to restore.
3. Choose **Restore**.



**Figure 43: Restoring confirmation screen**

4. After selecting **Restore**, the **Are you sure** screen opens.
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 server 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 snapshots on C:\ are tagged as **Always Keep** and the cache file fills up, the system may experience a BSOD at reboot.

HP recommends that you do not 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

All snapshots must be deleted 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

## 5

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 chapter 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 included with Windows 2000 Advanced Server.

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 server 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 server. 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 server 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 server 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 server 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 server, 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 server 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 5](#) provides examples of group names.

**Table 5: 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 server. Within the Users option, there are two choices:

- Managing local users
- Managing local groups

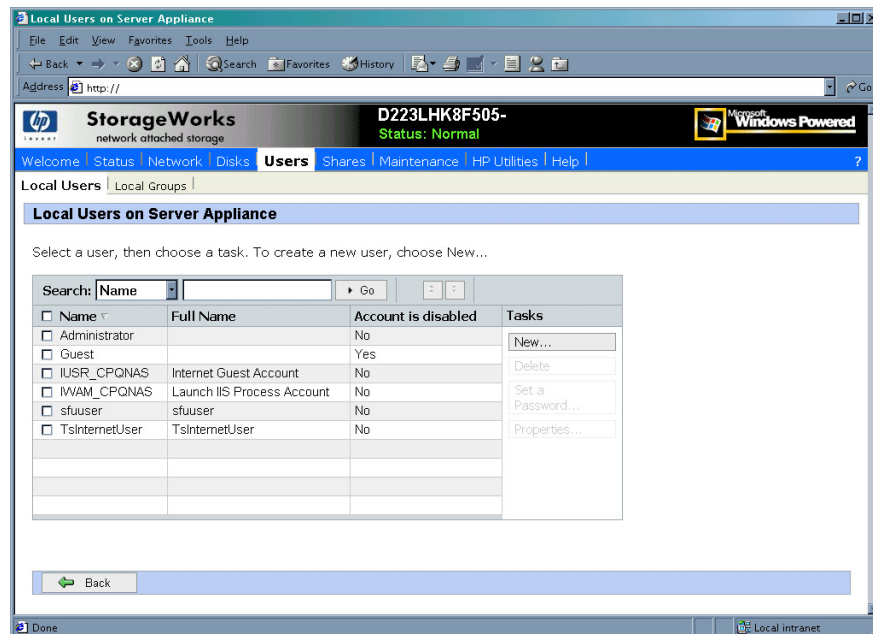
User and group administrative tasks include adding, deleting, and modifying user and group information. Managing local users and managing local groups are discussed in the following paragraphs.

### 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 44: 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.



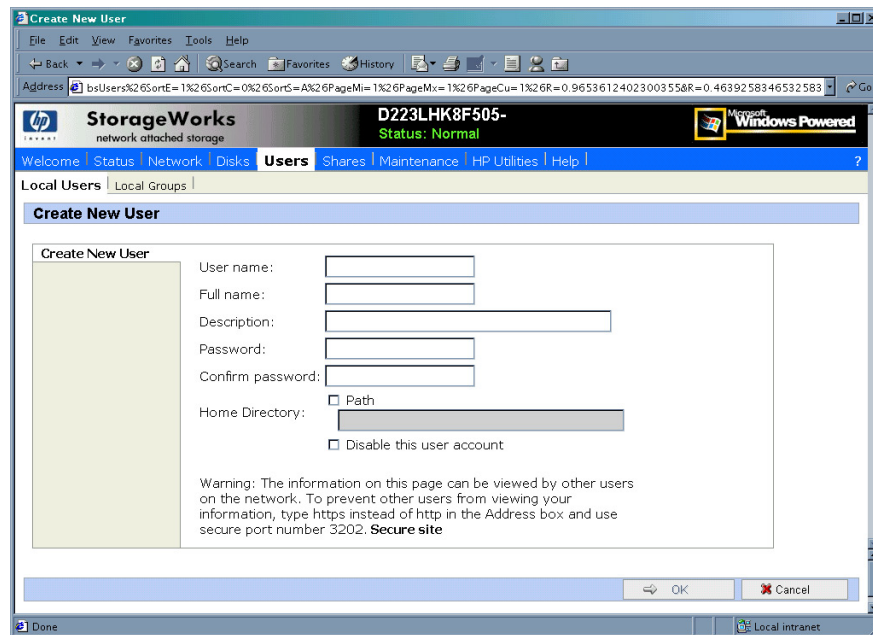
Existing user records can be retrieved in one of two ways:

- 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 45: 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.

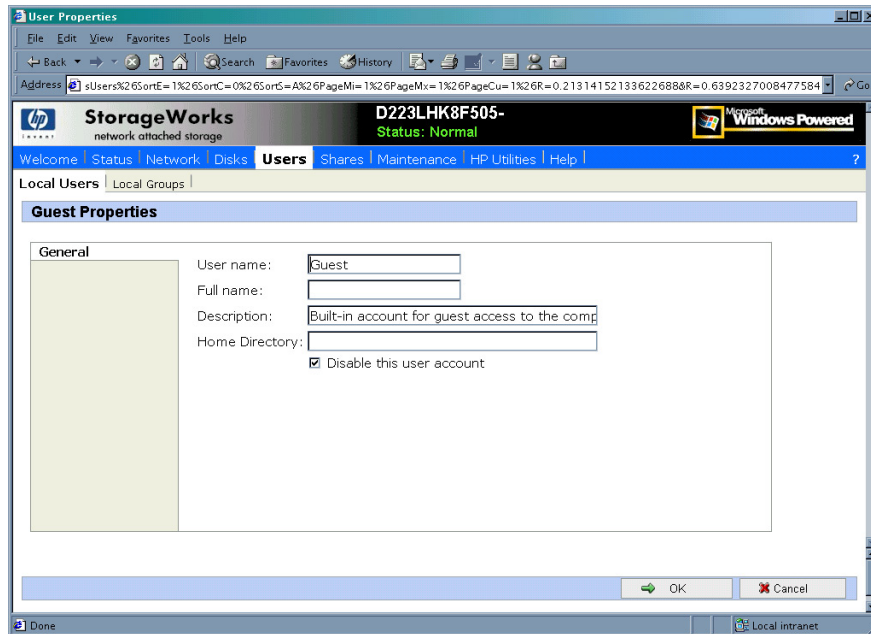
2. Enter the password and click **OK**. The Local Users dialog box is displayed again.

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



**Figure 46: 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.

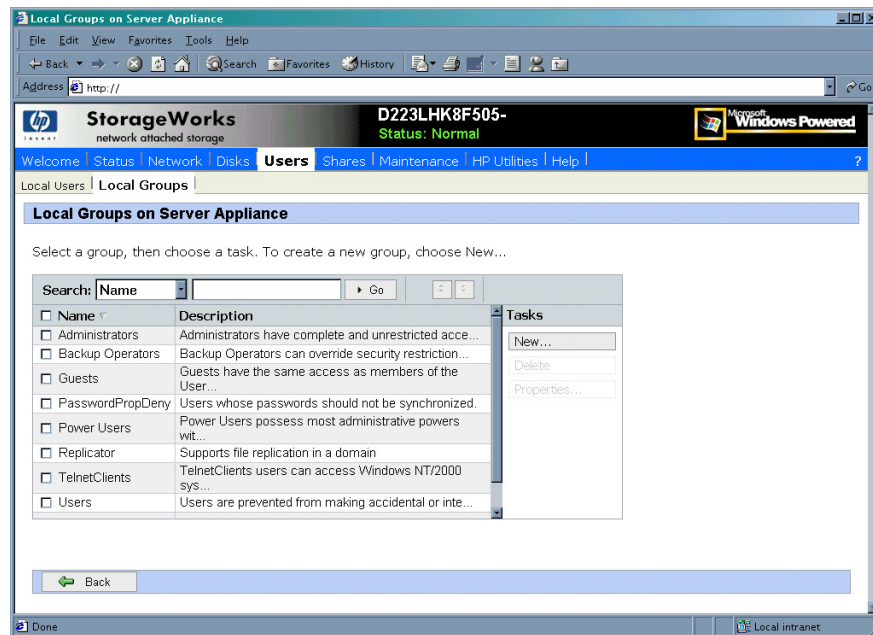
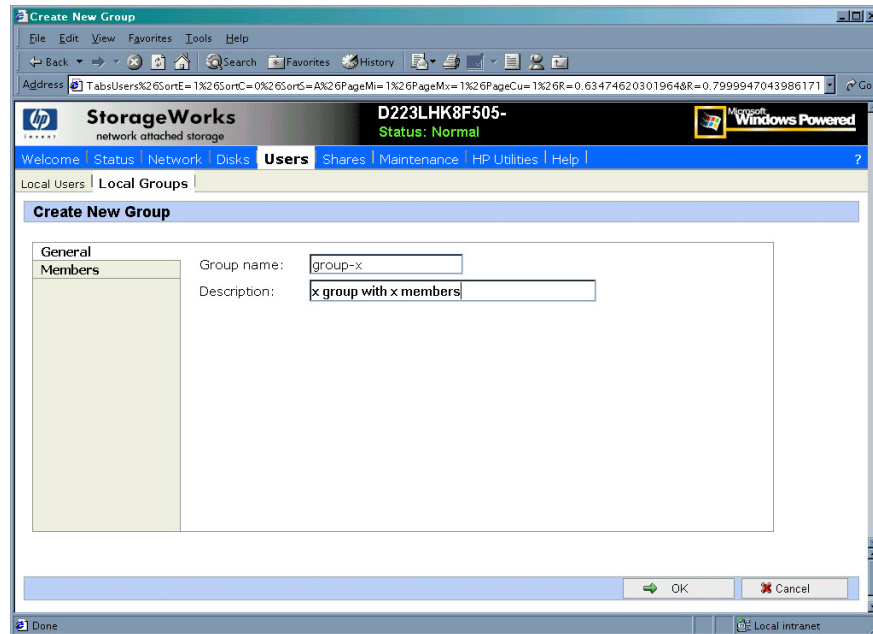


Figure 47: Local Groups dialog box

## 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 48: 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

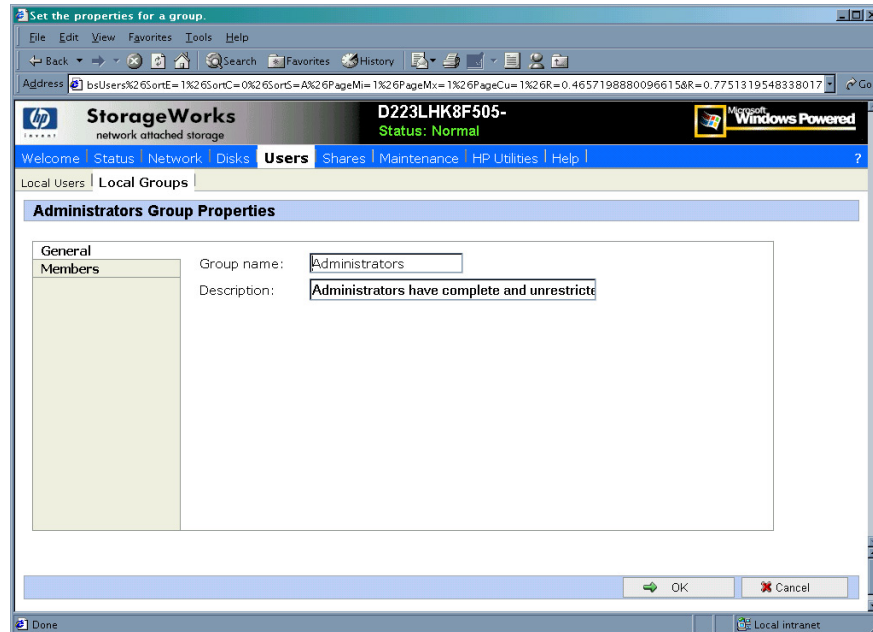
To delete a group:

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 49: 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 50 is an example of the **Members** tab.

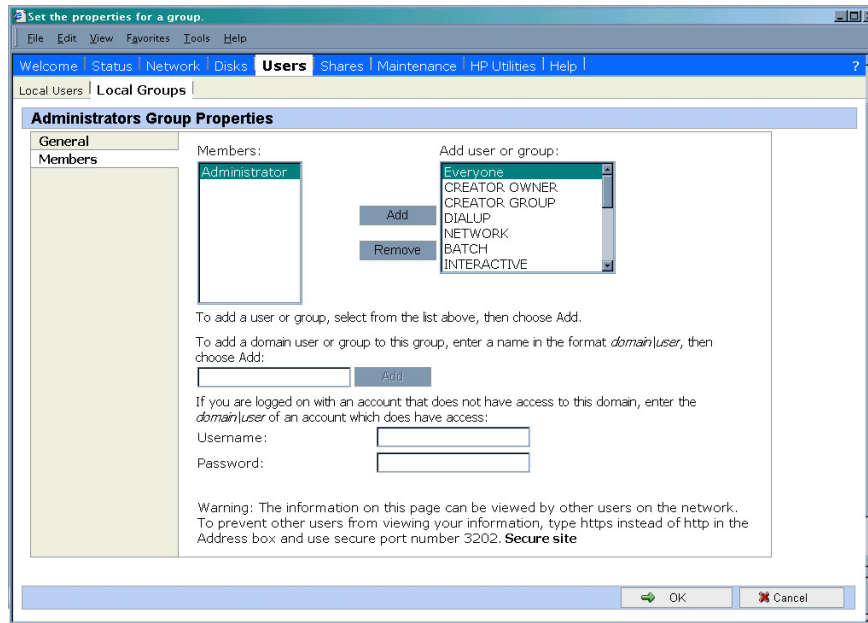


Figure 50: Group Properties dialog box, Members tab

## 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:

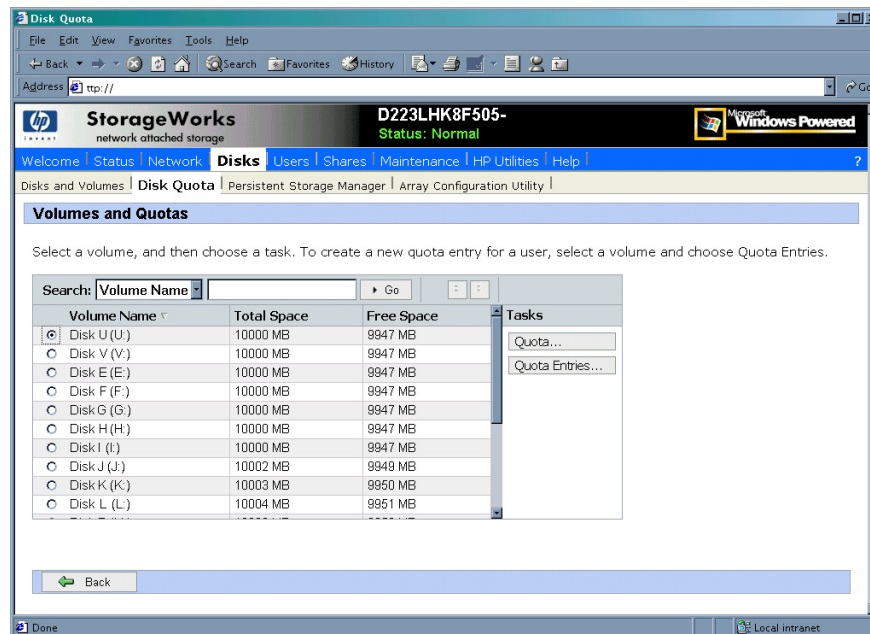
- Enabling and disabling quota management
- Creating new quota entries for a user or group
- Deleting quota entries for a user or group
- Modifying quota entries for a user or group

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 51](#) 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.

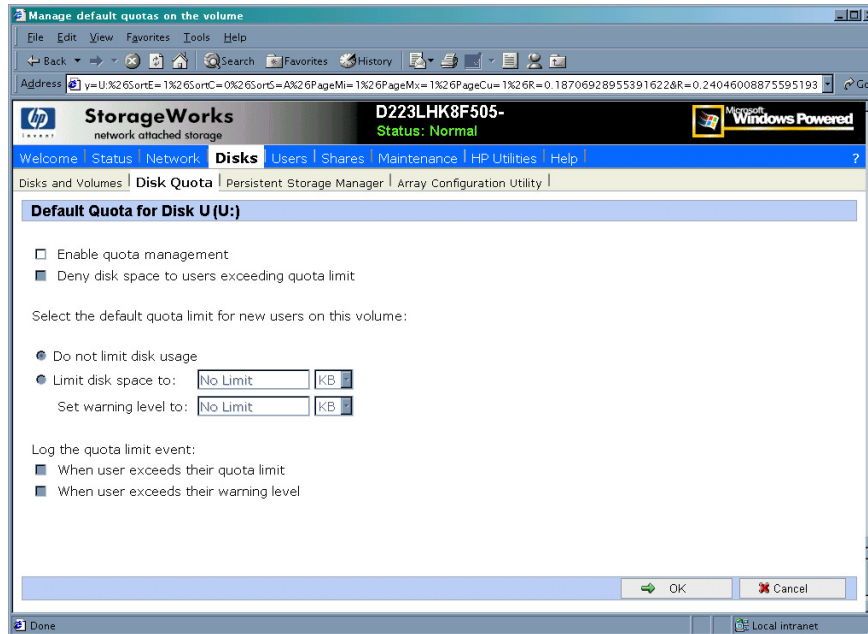


**Figure 51: Disk Quota dialog box**

## 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 52: Default Quota dialog box**

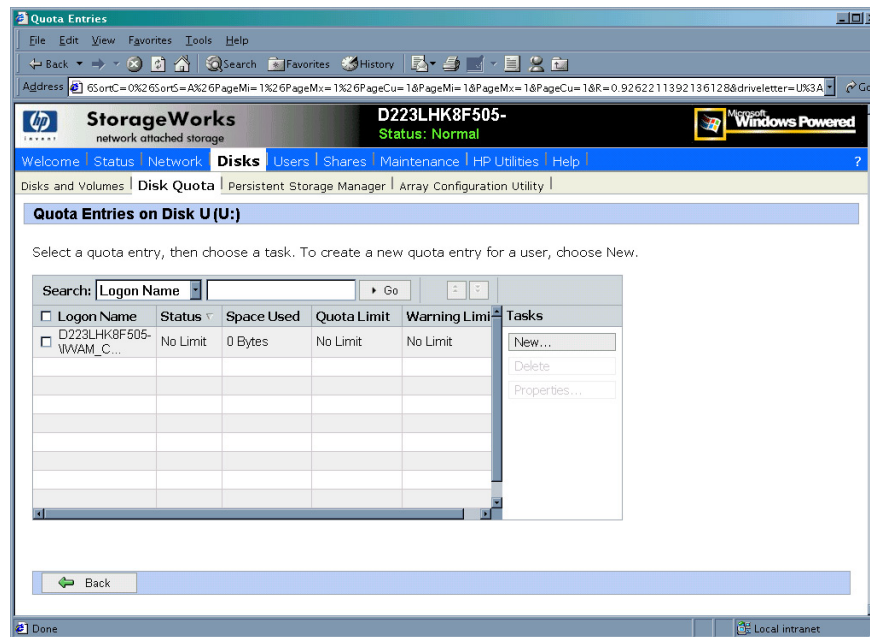
2. To enable quotas on the selected disk, select **Enable quota management**. Complete the additional data fields on the screen, including disk space and warning level limits and auditing settings.
3. To disable quotas on the selected disk, de-select **Enable quota management**.
4. After completed all field entries, click **OK**. The **Volume and Quotas** dialog box is displayed again.



## Creating New Quota Entries for a User or Group

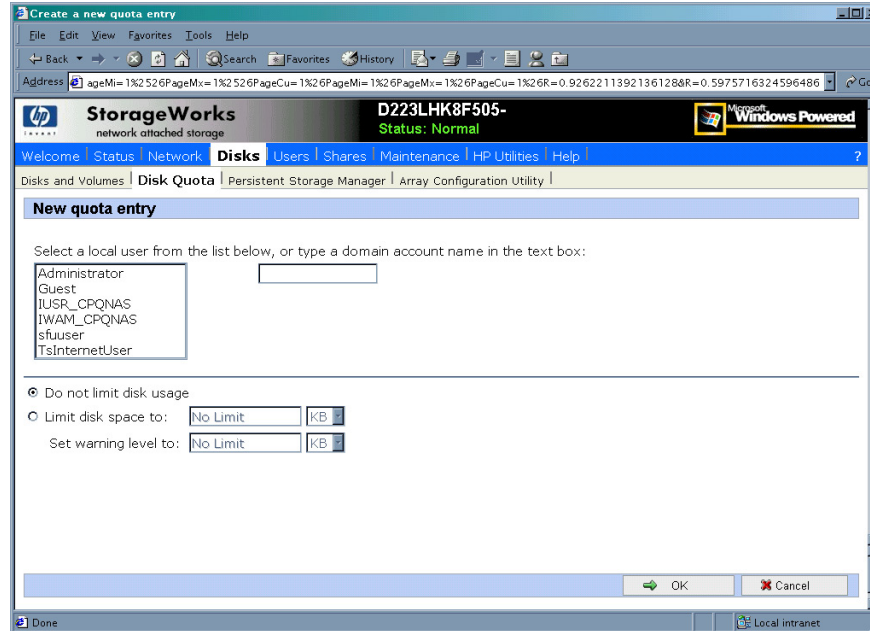
To create new quotas for a user or group:

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 53: 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 54: 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

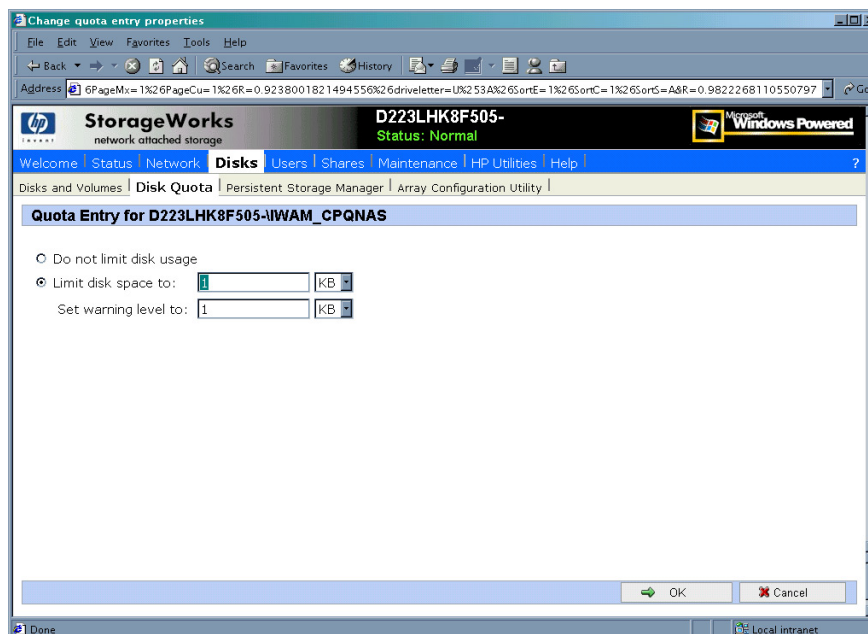
To delete quotas for a user or group:

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 55: 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

## 6

The HP StorageWorks NAS server supports several file sharing protocols, including CIFS, NFS, FTP, HTTP, NCP, and AFP (AppleTalk), however only CIFS, NFS and FTP are cluster-aware protocols. NCP, HTTP, and AFP can be installed on each node but the protocol cannot be set up with Cluster Administrator, nor will they failover if there is a node failure. In addition, AppleTalk is not supported on clustered resources as data loss can occur due to local memory use. 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/](http://www.microsoft.com/)

---

**Note:** The NAS b3000 v2 and e7000 v2 servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses share setup for a non-clustered deployment. For information on managing file shares in a cluster, see the “Cluster Administration” chapter.

---

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, 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 server, 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

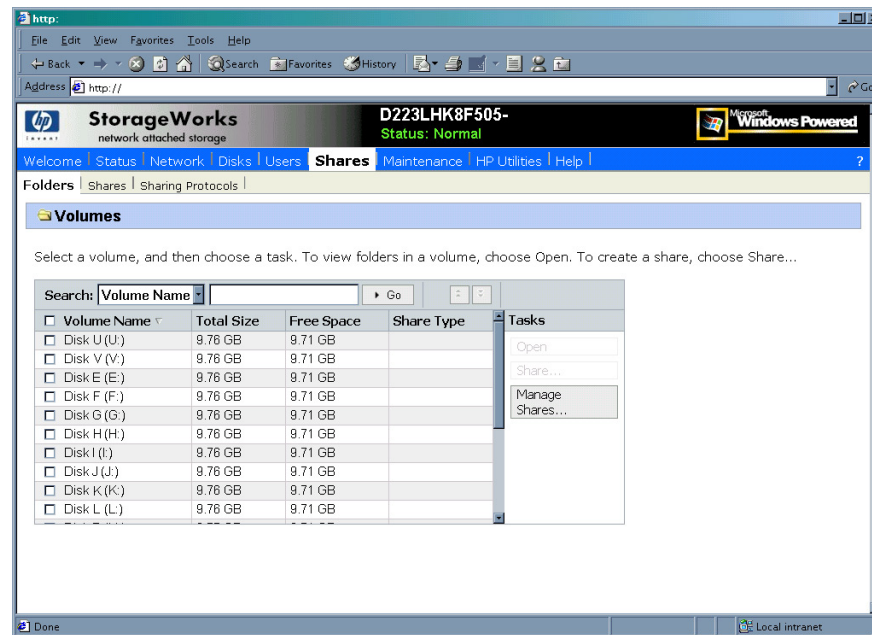
## Navigating to a Specific Volume or Folder

When working 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:

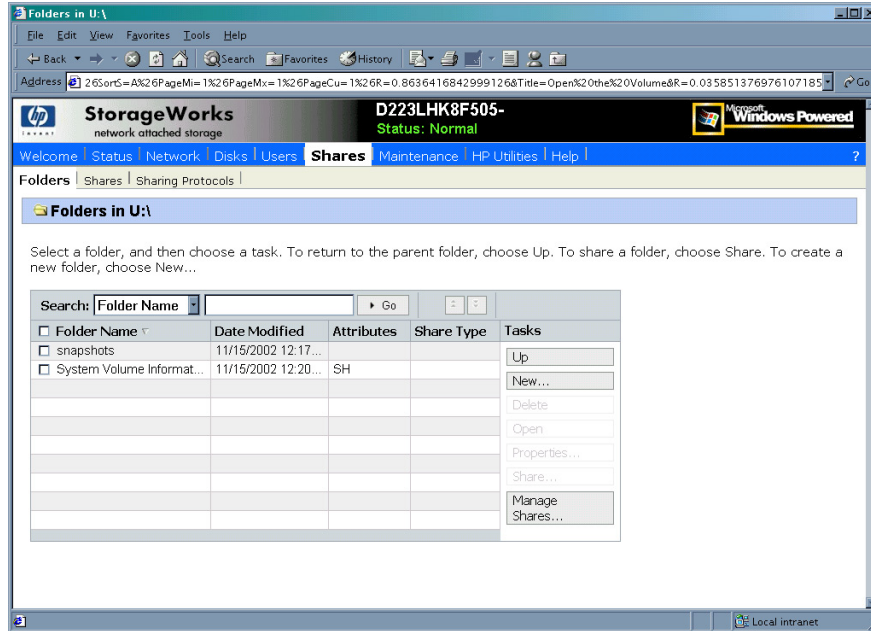
1. To navigate to a specific volume or folder, from the WebUI, select **Shares** and then **Folders**. Initially, the **Volumes** dialog box is displayed.

This initial dialog box displays all system volumes.



**Figure 56: 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 57](#) for an example of **Folders** dialog box.



**Figure 57: 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

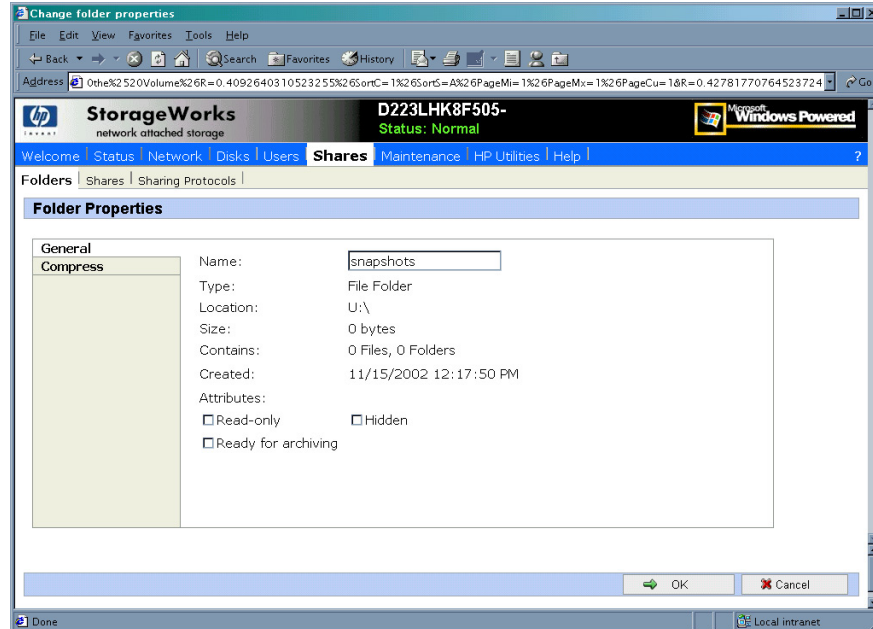
To create a new folder:

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 58: 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.

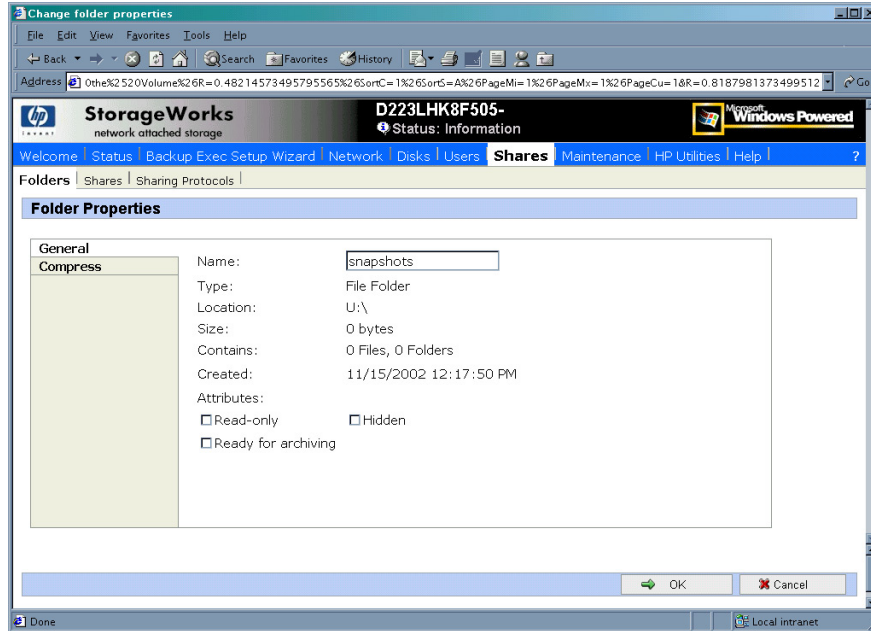
---

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 59: 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.

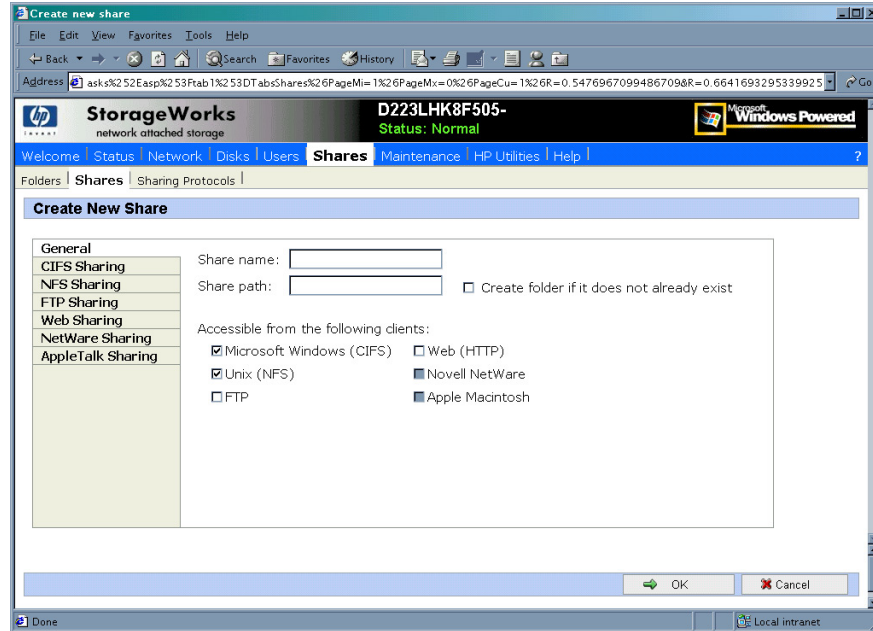
---

**Note:** This function does not operate in a cluster. Use Cluster Administrator to create shares for a cluster.

---

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 60: 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 “Modifying Share Properties” in the “Share Management” section for detailed procedural instructions on modifying shares.

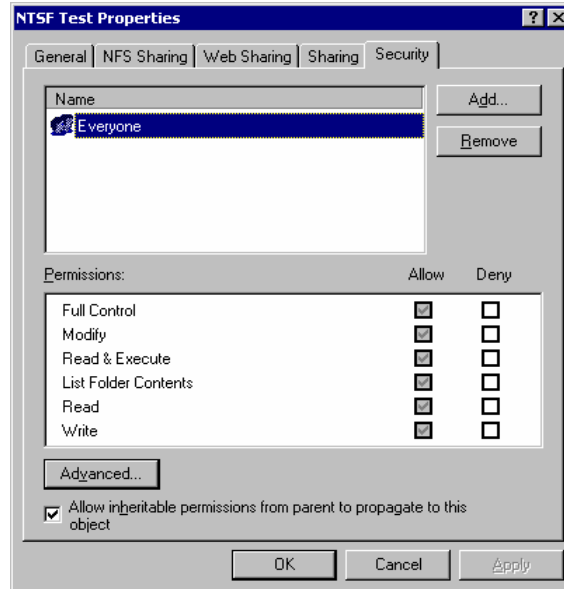
## Managing File Level Permissions

The WebUI of the NAS server 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 server. To access the NAS server 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 61](#) illustrates the properties available on the **Security** tab.



**Figure 61: Security Properties dialog box for folder name NTFS Test**

Several options are available in the **Security** tab dialog box:

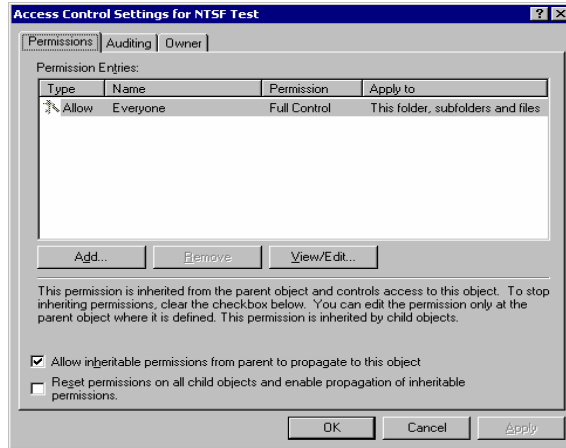
- To add users and groups to the permissions list, click **Add**. Then follow the dialog box instructions.
- To remove users and groups from the permissions list, highlight the desired user or group and then click **Remove**.
- If the **Allow inheritable permissions from parent to propagate to this object** 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.

---

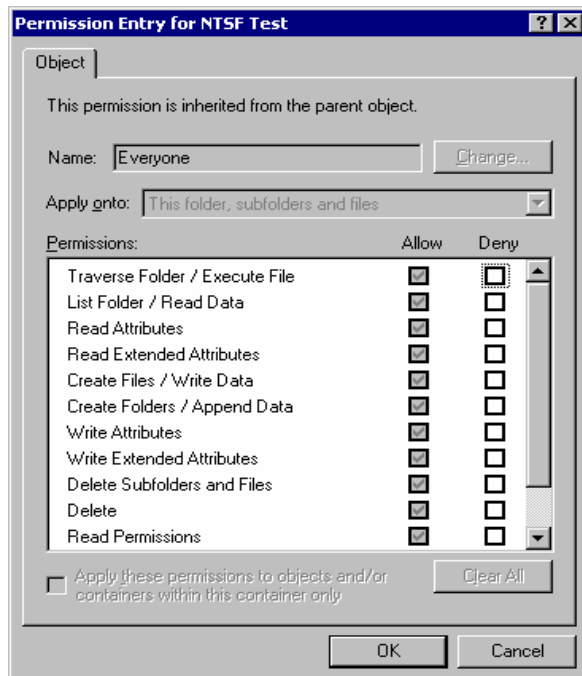
- To modify ownership of files or to modify individual file access level permissions, click **Advanced**.



**Figure 62: 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 to enable, and clear the permissions 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 63](#) illustrates the **View/Edit** screen and some of the permissions.




**Figure 63: 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 63](#) 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 setting of 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 64](#).



**Figure 64: Access Control Settings, Auditing tab dialog box for folder name NTSF Test**

[Figure 65](#) illustrates the screen that is displayed when a user or group to be audited is added.

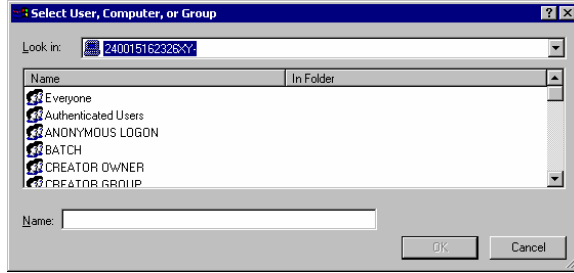
4. Select the appropriate domain or machine name from the **Look in:** drop-down list box at the top of the screen.

---

**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.

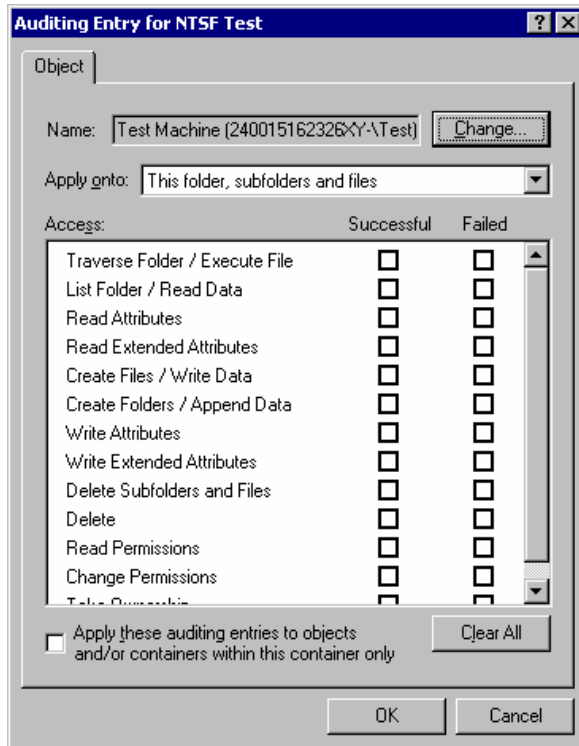
---

5. Select the user or group.



**Figure 65: Select User, Computer, or Group dialog box**

6. Click **OK**. [Figure 66](#) illustrates the **Auditing Entry** screen that is displayed.



**Figure 66: Auditing Entry dialog box for folder name NTFS Test**

7. Select the desired **Successful** and **Failed** audits for the user or group as shown in [Figure 66](#).
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 server.

---



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, the administrator gains access to the files and can then manually apply the appropriate security configurations. [Figure 67](#) illustrates the **Owner** tab.



**Figure 67: 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 server 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 server 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 server can be given access permissions to shares managed by the device. The domain name of the NAS server 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 server 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 server, 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 server, 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 server. 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.

To use the same name when sharing a folder through CIFS, and then export 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:

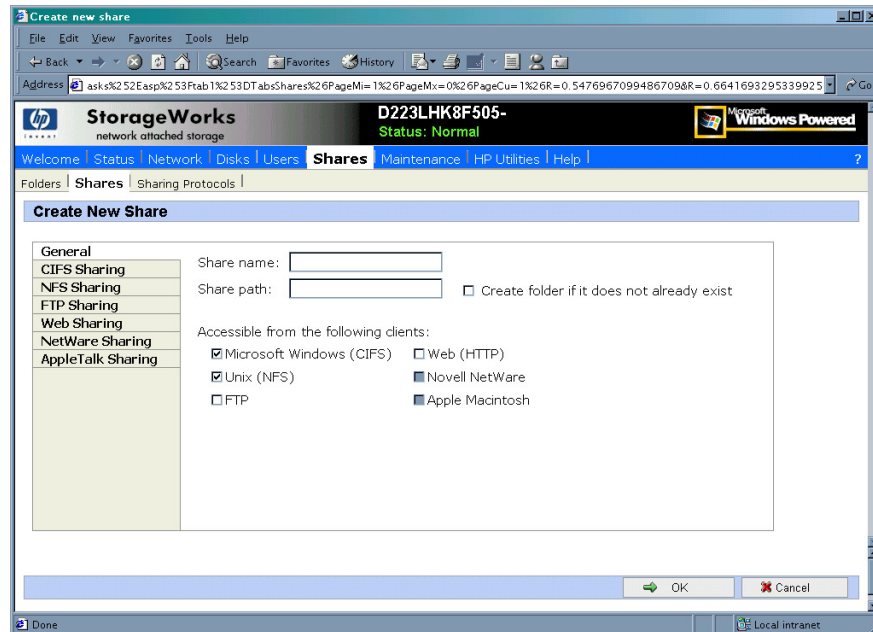
- Creating a new share
- Deleting a share
- Modifying share properties

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 68: 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:

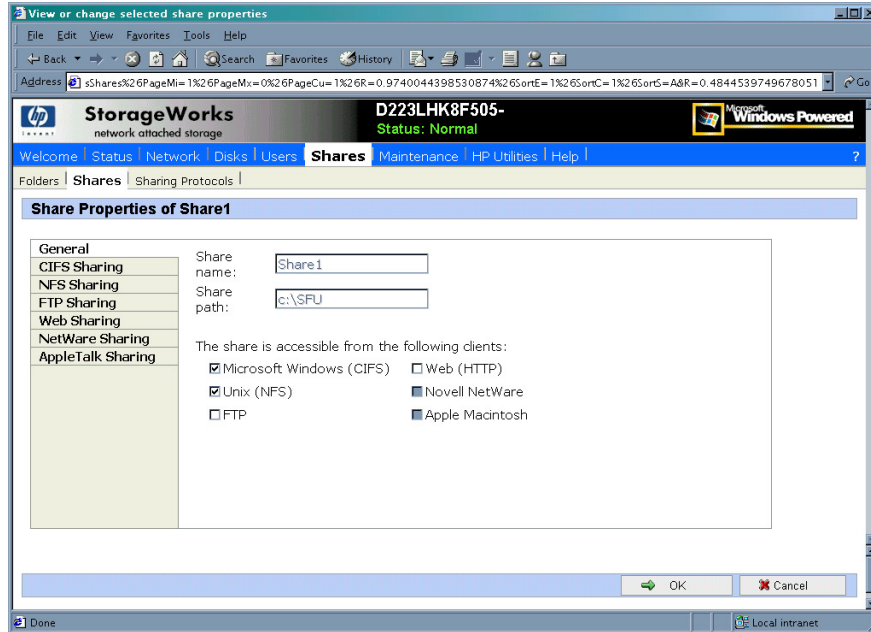
1. From the **Shares** menu, click **Delete**.
2. Verify that this is the correct share, and then click **OK**.

**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 69: 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.

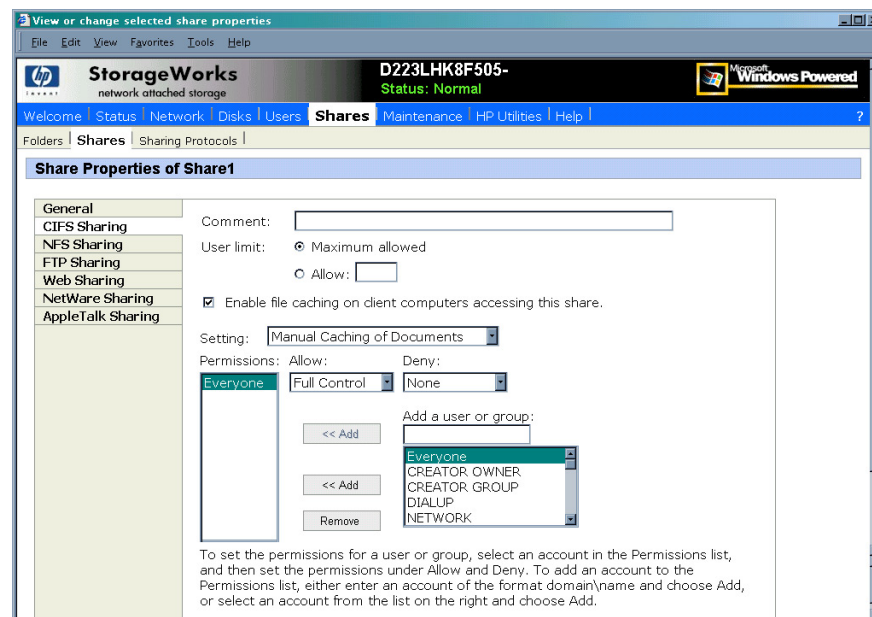
### CIFS Sharing

From the **CIFS Sharing** tab of the **Share Properties** dialog box:

1. Enter a descriptive **Comment**, and the **User limit** (optional).  
See [Figure 70](#) 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 70: Share Properties dialog box, CIFS Sharing tab**

3. Enter Permissions information:

The **Permissions** box lists the currently approved users for this share.

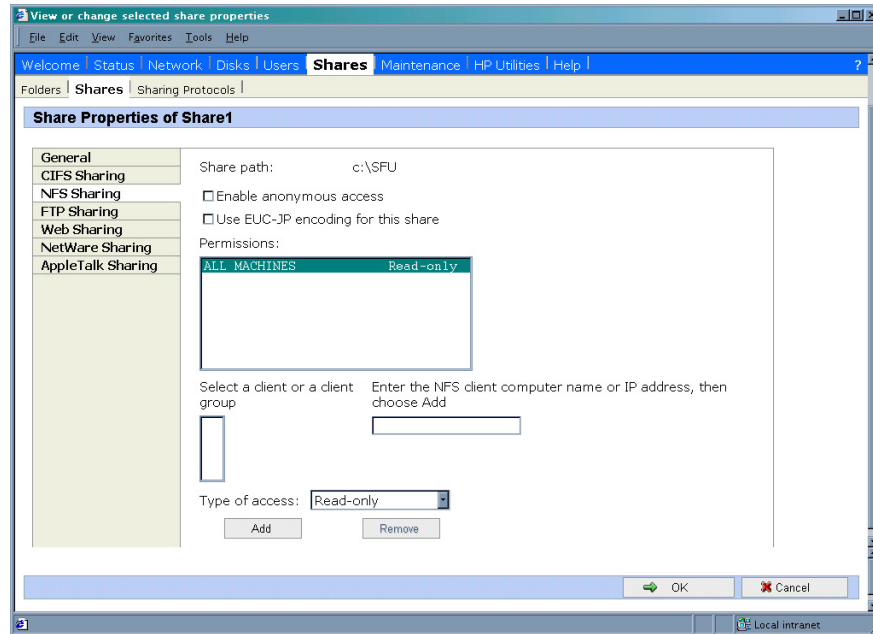
- *To add a new user or group*, either select a user or group from the box at the bottom right of the screen or manually enter the user or group name in the Add a user or group box and then click **Add**. That user or group is added to the Permissions box.
  - *To remove access to a currently approved 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 71: 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)

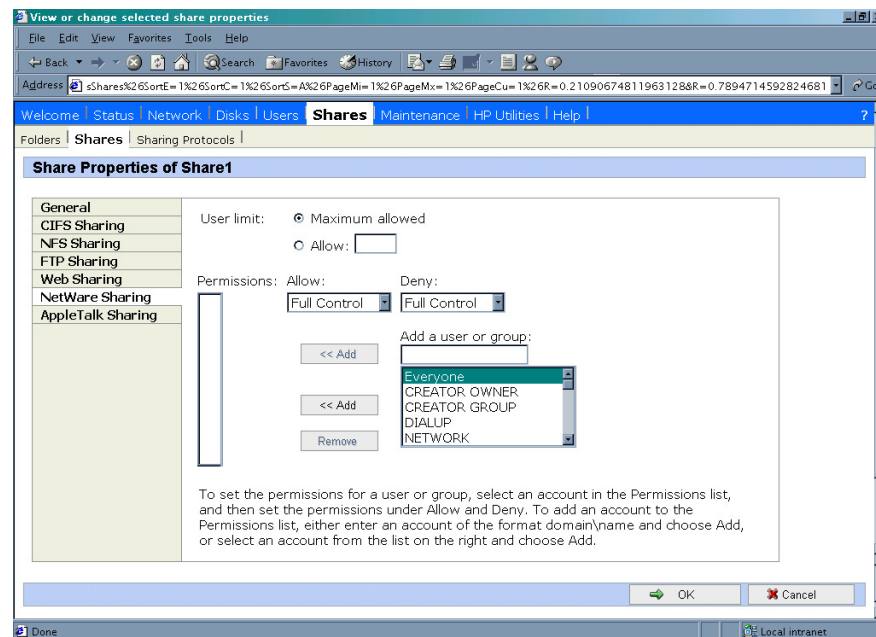
**Note:** NCP shares can be set up only after Microsoft Services for NetWare (SFN) has been installed on the NAS server. Procedures for installing SFN are included in the “NetWare File System Management” chapter.

From the **NetWare Sharing** tab, as illustrated in [Figure 72](#), 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 72: 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 server.

---

**Note:** AppleTalk shares should not be created on clustered resources as data loss can occur due to local memory use.

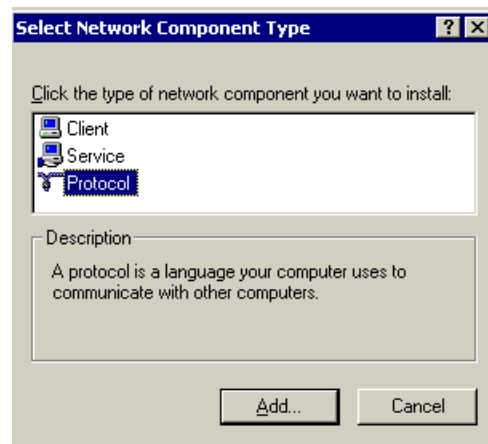
---

#### *Installing Services for AppleTalk*

To install Services for AppleTalk:

1. From the desktop of the NAS server, 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 73 is an example of the **Select Network Component Type** dialog box.



**Figure 73: 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

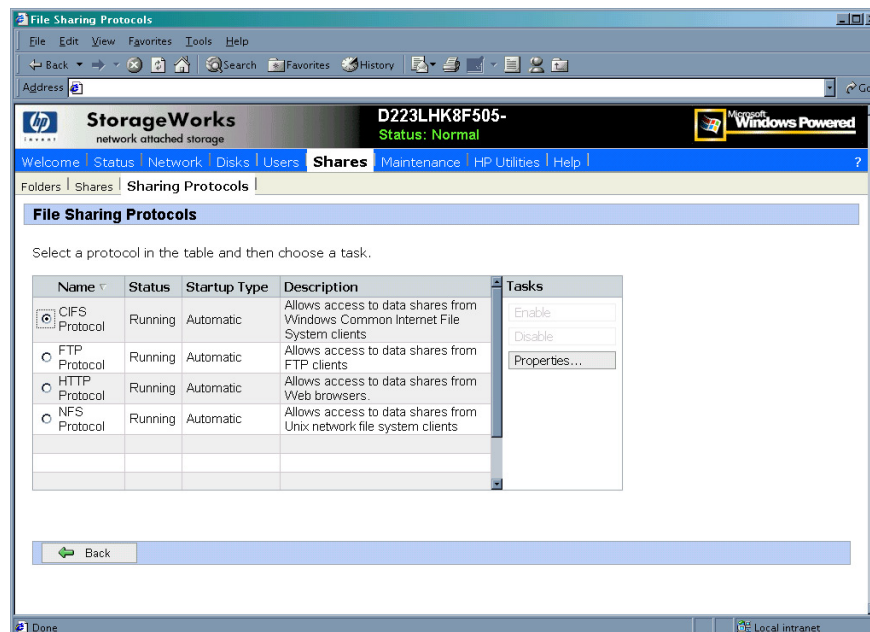
As previously mentioned, the NAS server supports the following protocols:

- 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 74: 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

## 7

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 server:

- Server for NFS
- User Name Mapping
- Telnet and Remote Shell Services
- Password Synchronization

---

**Note:** SFU can be implemented in both clustered and non-clustered environments. This chapter discusses SFU in a non-clustered deployment. For additional information that is specific to a cluster, see the “Cluster Administration” chapter.

---

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.

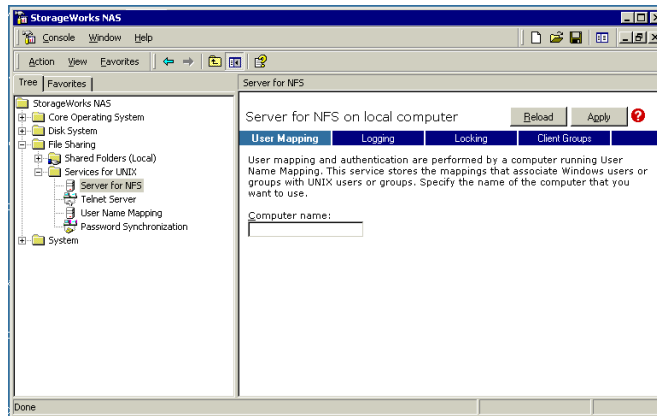
## Indicating the Computer to Use for the NFS User Mapping Server

During the processes of starting and installing the NAS server, the name localhost is assigned by default to the computer. It is assumed that the NAS server 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 75](#) 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 server. 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 75: 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\SFUNlog\nfssvr.log*) and log file size (7-MB default). The default log file is created when the changes are applied.



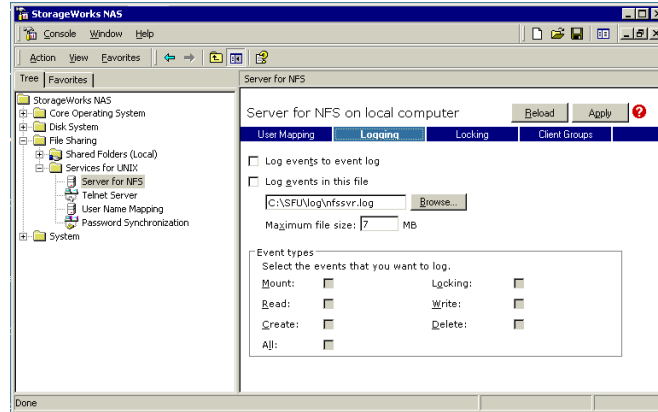


Figure 76: NAS Management Console Server for NFS screen, Logging tab

## 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 server.
2. Share out the *SFU* directory on the NAS server.
3. On the domain controller where the service is being installed, using Windows Explorer:
  - a. Connect to the SFU share on the NAS server.
  - 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:

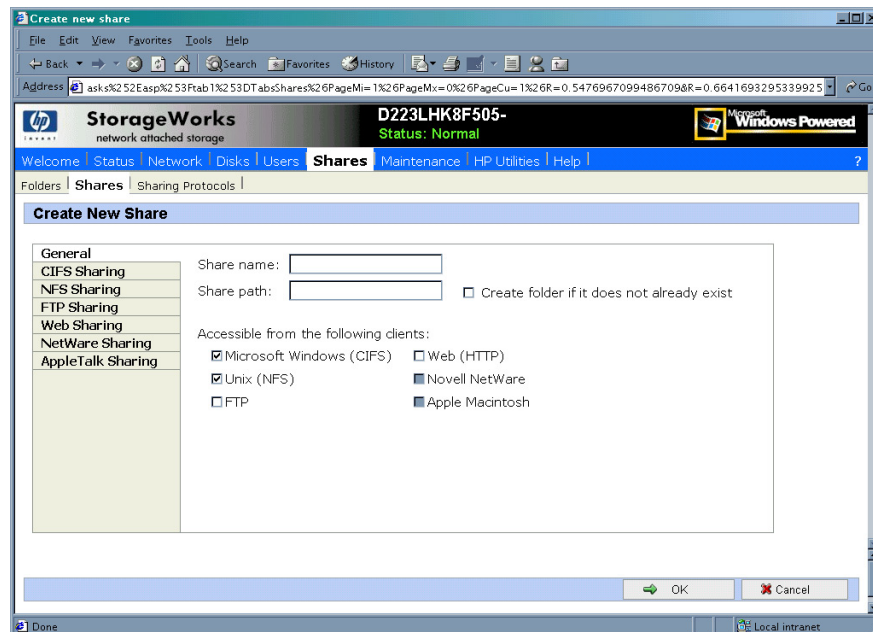
- Creating a new share
- Deleting a share
- Modifying share properties

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 77: 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. To use the same name when sharing a folder through CIFS, and then export 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.

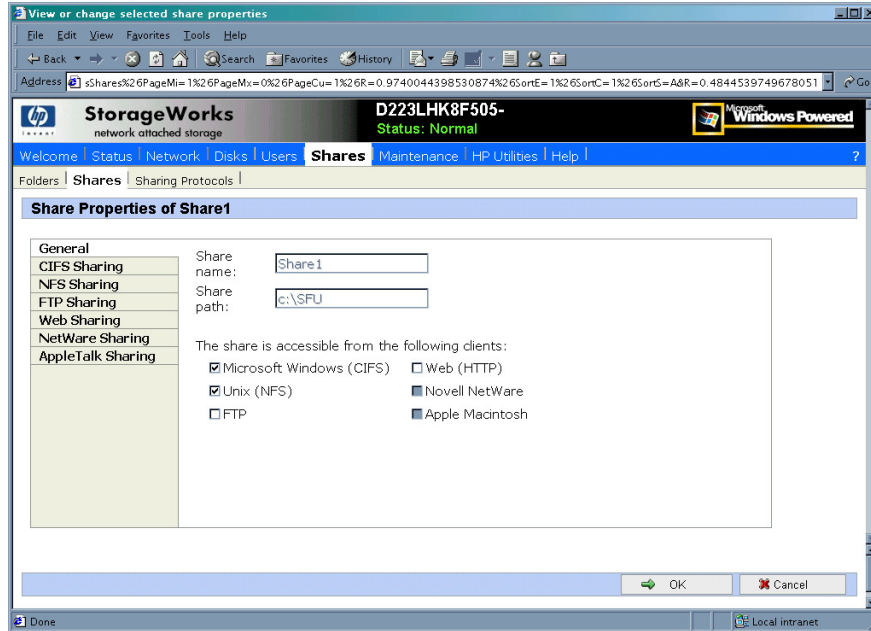


Figure 78: Share Properties dialog box, General tab

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.

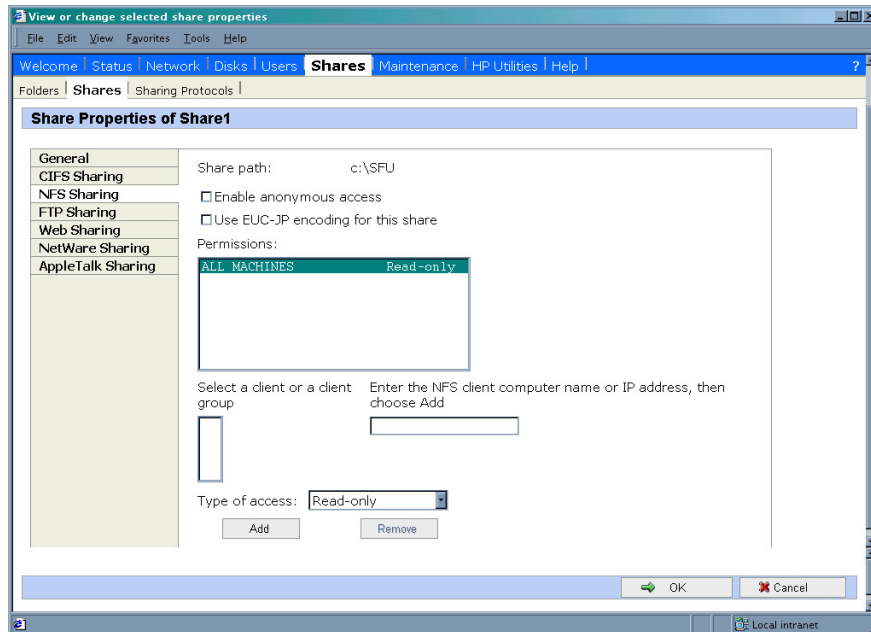


Figure 79: NFS Sharing tab

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)

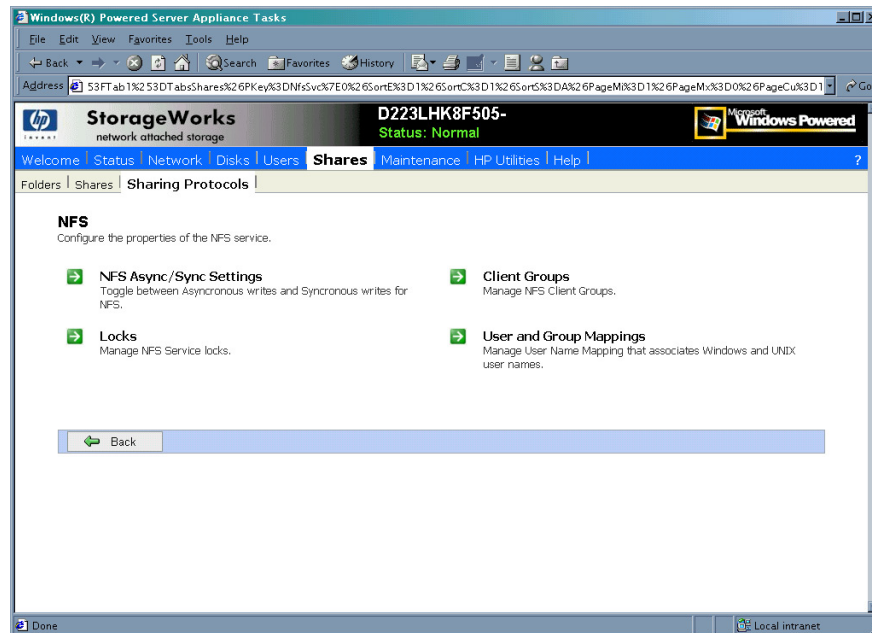
If the option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:

- 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 80: NFS Sharing Protocols menu**

NFS properties include:

- Async/Async 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.

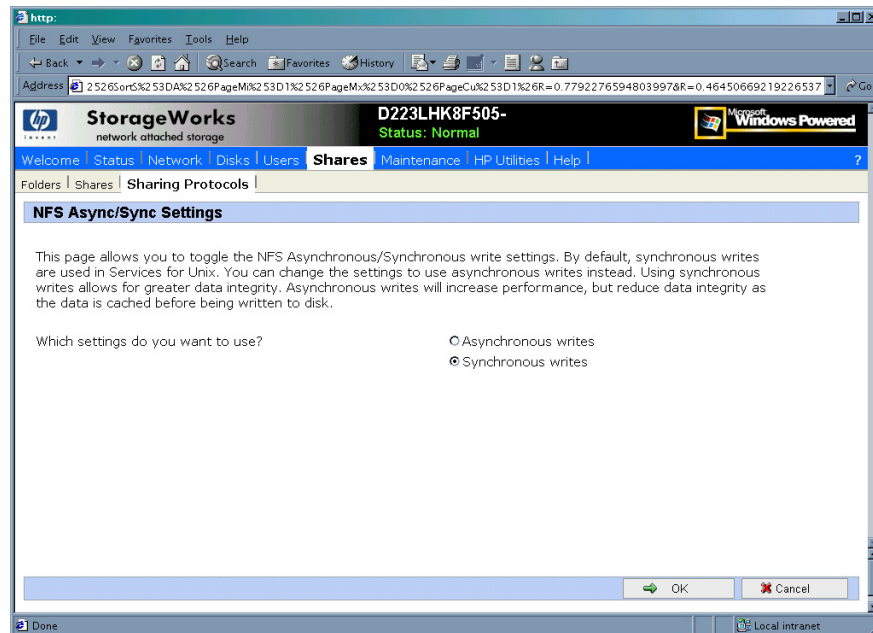
To indicate whether to use asynchronous or synchronous write settings:

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 81: 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**.

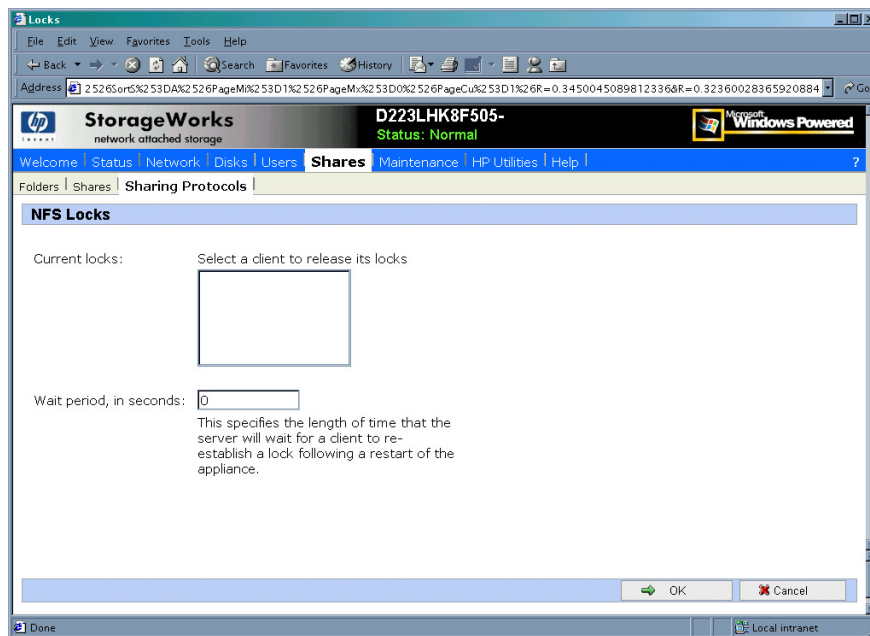
The **NFS Properties** menu is displayed.

2. In the **NFS Properties** menu, select **Locks**. The **NFS Locks** dialog box is displayed. **Figure 82** 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 server 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 82: NFS Locks dialog box**



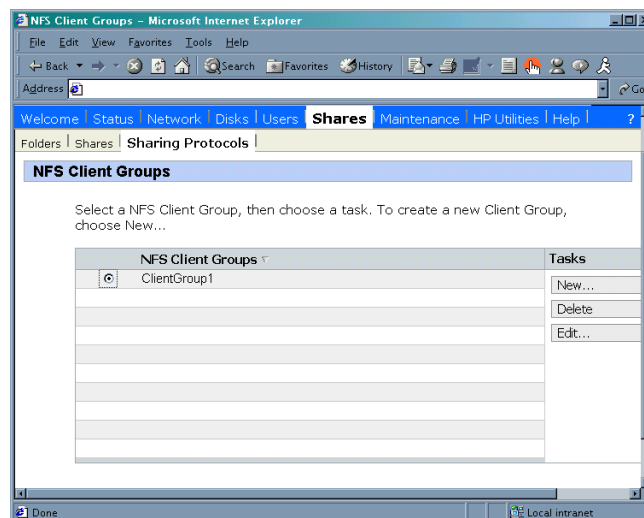
## NFS Client Groups

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 83: NFS Client Groups dialog box**

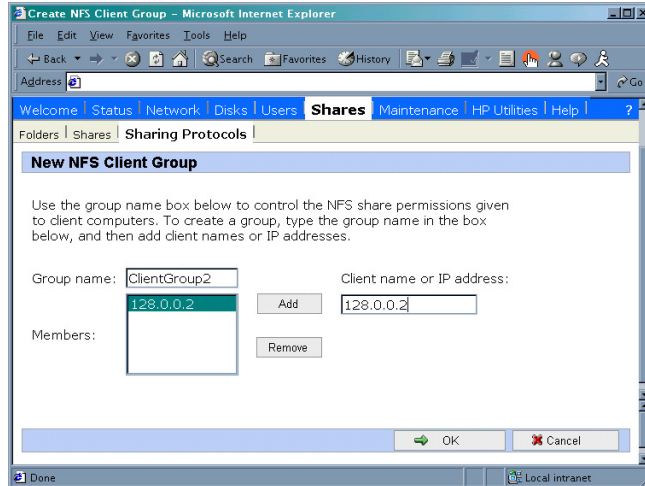
The following tasks are available:

- 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 84: 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.

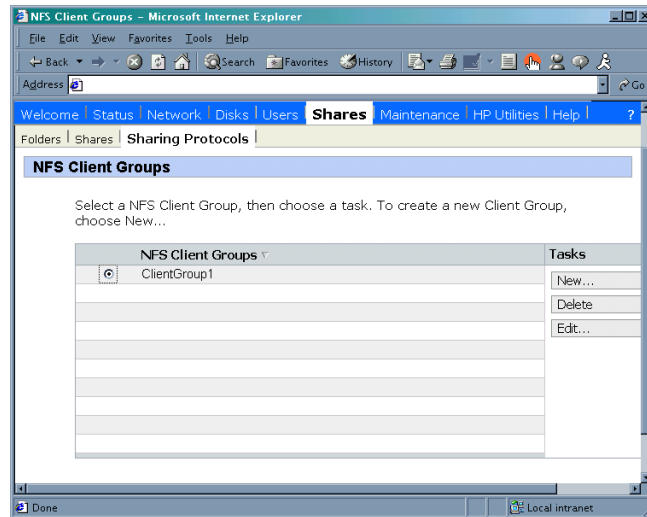


Figure 85: Client Groups dialog box

## 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.

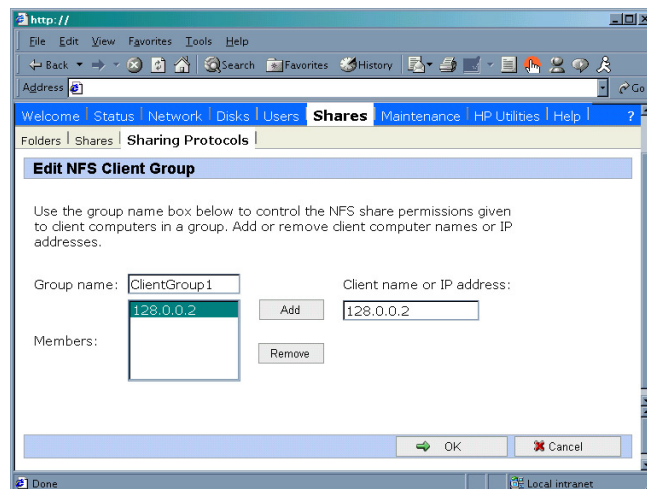


Figure 86: Edit NFS Client Groups dialog box

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 server 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 server 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.

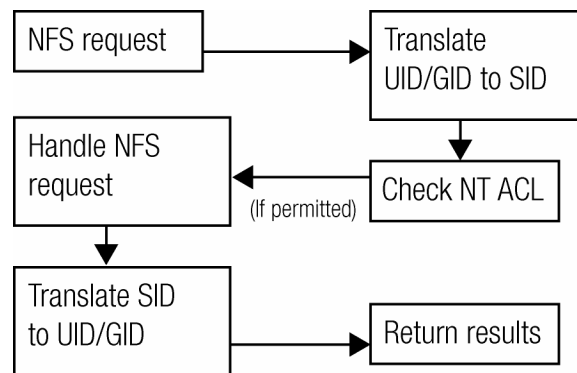
### Simple Mappings

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 system 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.

Figure 87 is a diagram showing an example of how the mapping server works for an `ls -al` command.



**Figure 87: Mapping Server “ls -al” Command example**

A double translation, as illustrated in Figure 87, 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.

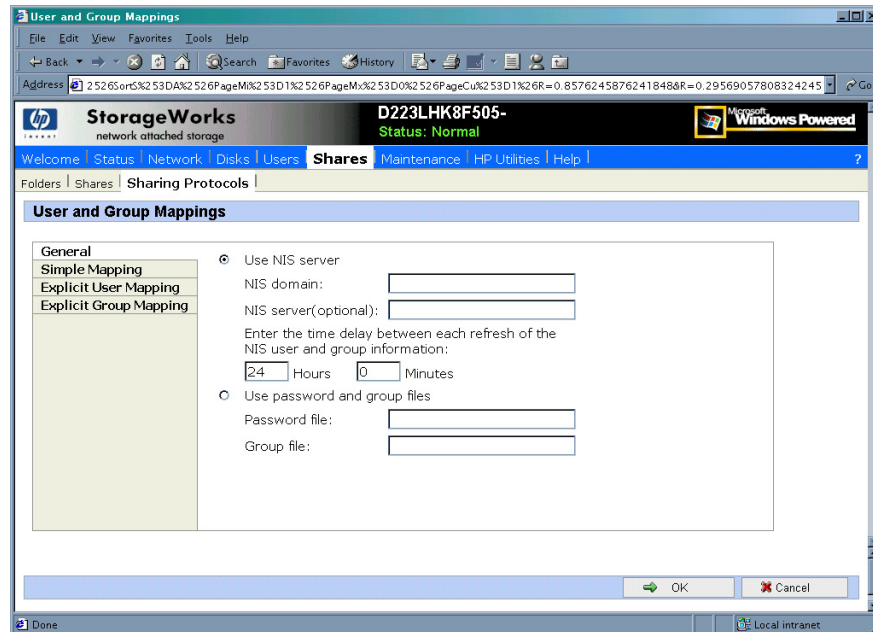
Each of these tabs is discussed in the following sections.

3. Enter mapping information on the appropriate tabs, then click **OK**.

### General Information

The NAS server 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 88: 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.

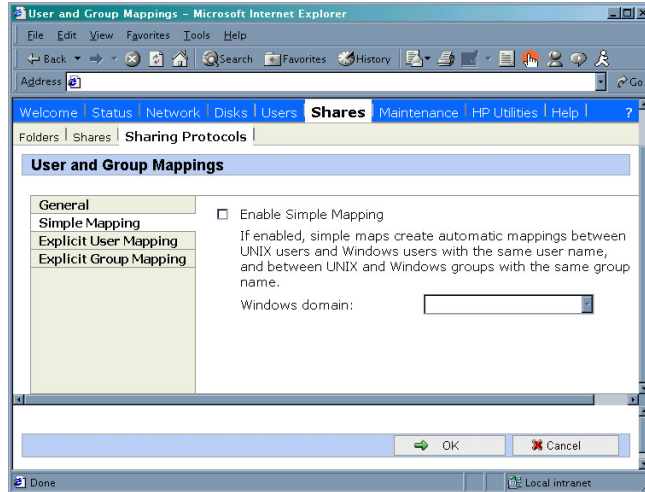


Figure 89: User and Group Mappings dialog box, Simple Mapping tab

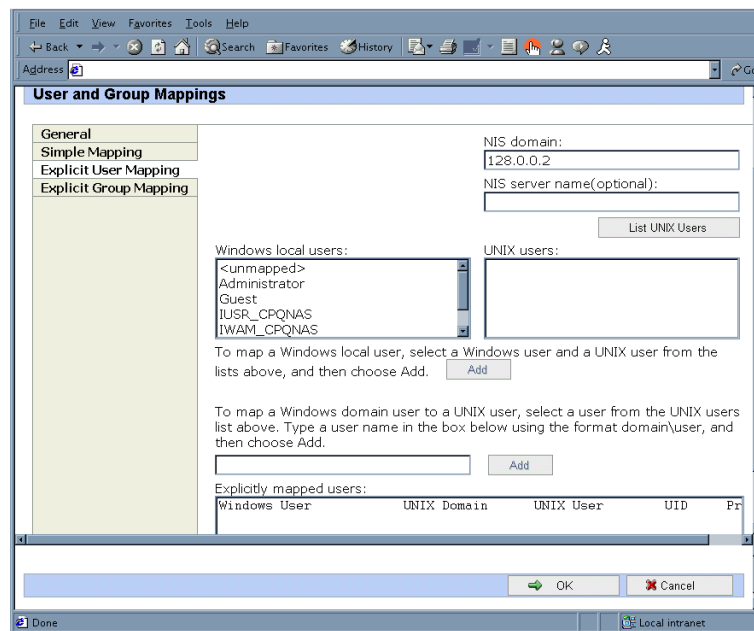
## Explicit User Mapping

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





**Figure 90: User and Group Mappings dialog box, Explicit User Mapping tab**

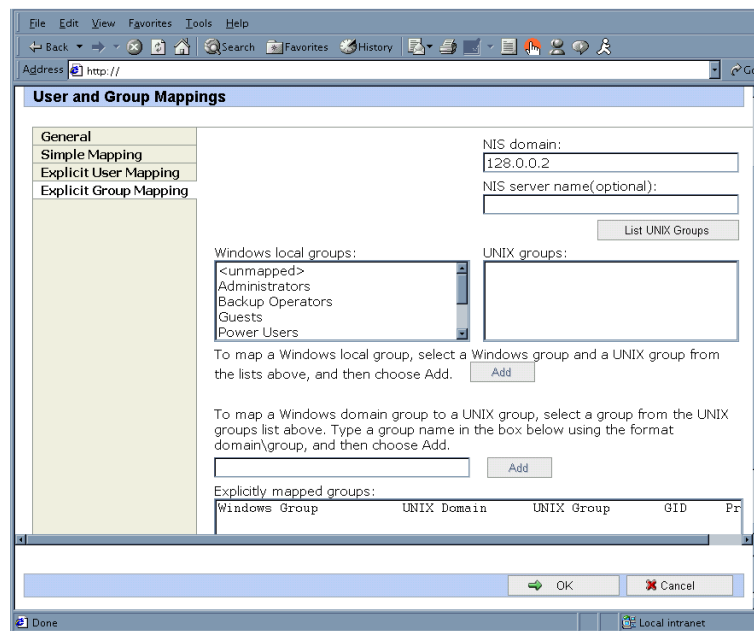
To create explicit user mappings:

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 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 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 91](#) 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 91: User and Group Mappings dialog box, Explicit Group Mapping tab**

To create explicit group mappings:

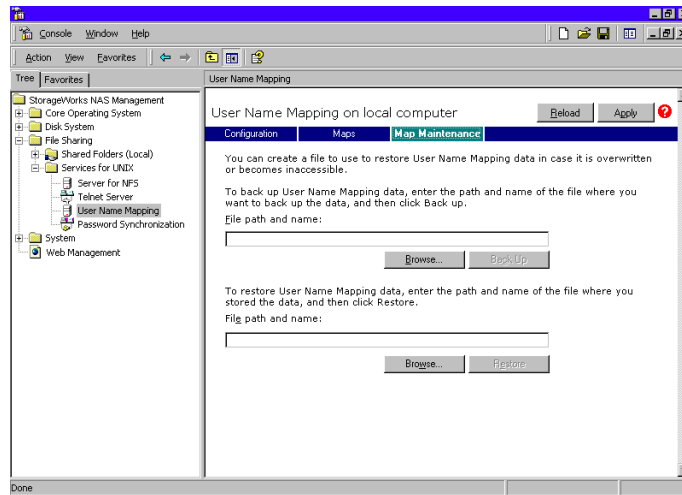
1. Click the **List UNIX Groups** button to populate the **UNIX Groups** box.
2. To map a local Windows group to a UNIX group, highlight the Windows group in the Windows local groups box and highlight the UNIX group to map, and then click **Add**. The **Explicitly mapped groups** box at the bottom of 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 92](#).

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



**Figure 92: 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 server 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 server 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 server:  
`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 6](#) for a list of commonly used commands.

### Using Terminal Services

Microsoft Terminal Services can be used to remotely access the NAS server 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 server. 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 server, but, by default, it is not activated. To use Telnet services, see the information in the “Remote Access Methods and Monitoring” chapter.

---

**Note:** Telnet is not cluster-aware.

---

### 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.

---

**Note:** Remote Shell Service is not cluster-aware.

---

Table 6 describes some common SFU commands.

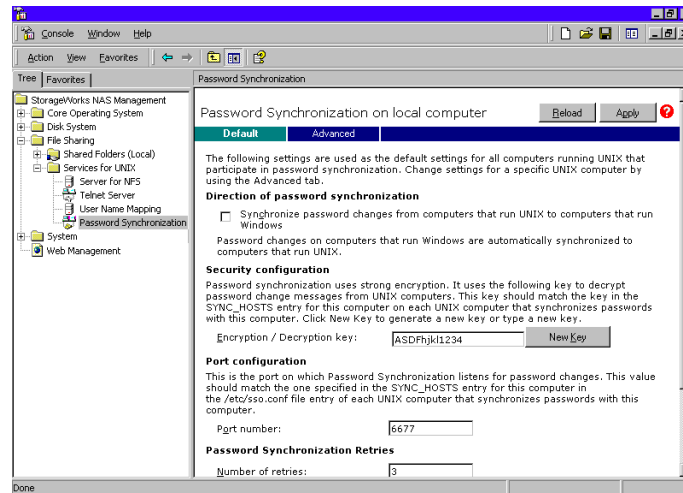
**Table 6: Command Line Interface Command Prompts**

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

## Password Synchronization

Password synchronization is an optional service that automatically synchronizes Windows passwords with UNIX passwords across multiple machines or environments. This service is included on the NAS server, but it is not activated.

**Note:** Password synchronization is not cluster-aware. Password synchronization may not occur during cluster failover conditions.



**Figure 93: 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.

For 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 a user is added 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 server 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 94: 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 server 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 server:
 

```
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 server 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 server:
 

```
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

## 8

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 server emulate a NetWare file and print server. FPNW eases the addition of the NAS server 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 server or through an existing NDS (Novell Directory Services) account.

---

**Note:** NetWare is not a clusterable protocol. With NetWare on both nodes of the cluster, the shares will not failover as the protocol is not cluster-aware.

---

---

**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 server 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](http://www.microsoft.com/WINDOWS2000/guide/server/solutions/NetWare.asp)

---

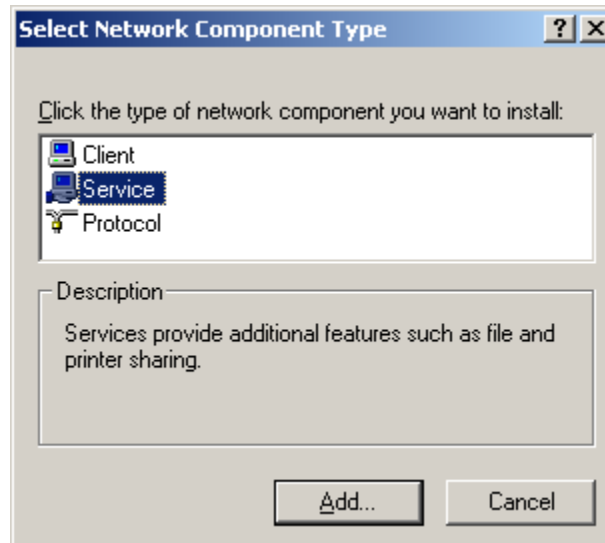
**Note:** The printing capabilities of File and Print Services for NetWare are not supported on the NAS server.

---

To install Services for NetWare:

1. From the desktop of the NAS server, 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 95** is an example of the **Select Network Component Type** dialog box.



**Figure 95: 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**.

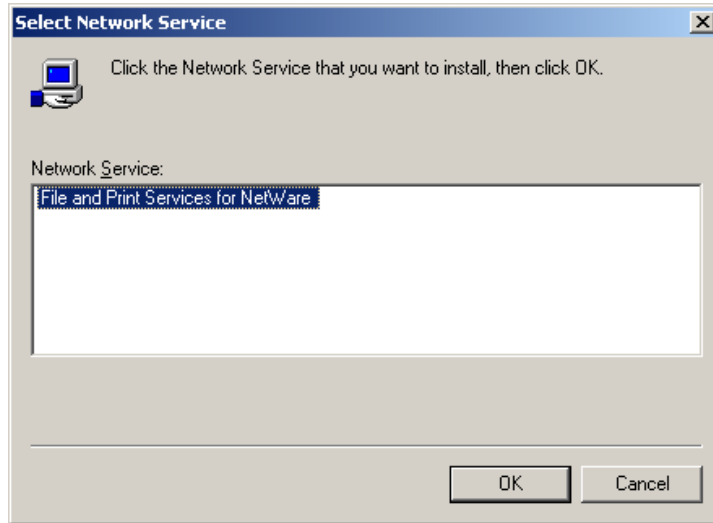


Figure 96: Installing File and Print Services for NetWare

## Managing File and Print Services for NetWare

To access FPNW:

1. From the desktop of the NAS server, click **Start, Settings, Control Panel**, and then double-click **FPNW**.




Figure 97: 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 an existing name is changed, the new name will not be effective until **File and Print Services for NetWare** is stopped and restarted. For example, in [Figure 97](#) 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 server.

### Adding Local NetWare Users

1. From the NAS server 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 98: 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 99: 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 server 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

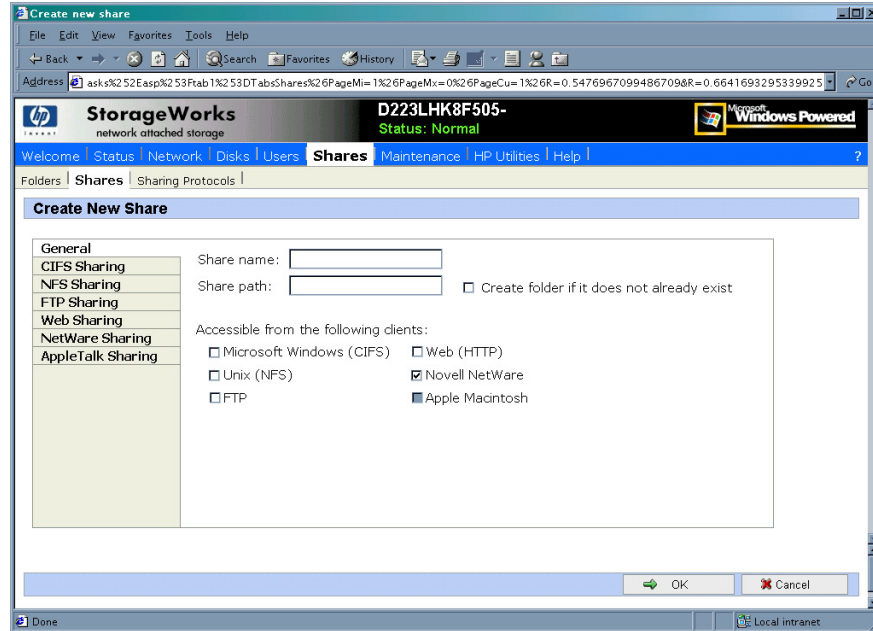
Each of these tasks is discussed in this section.

#### 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 100: 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:

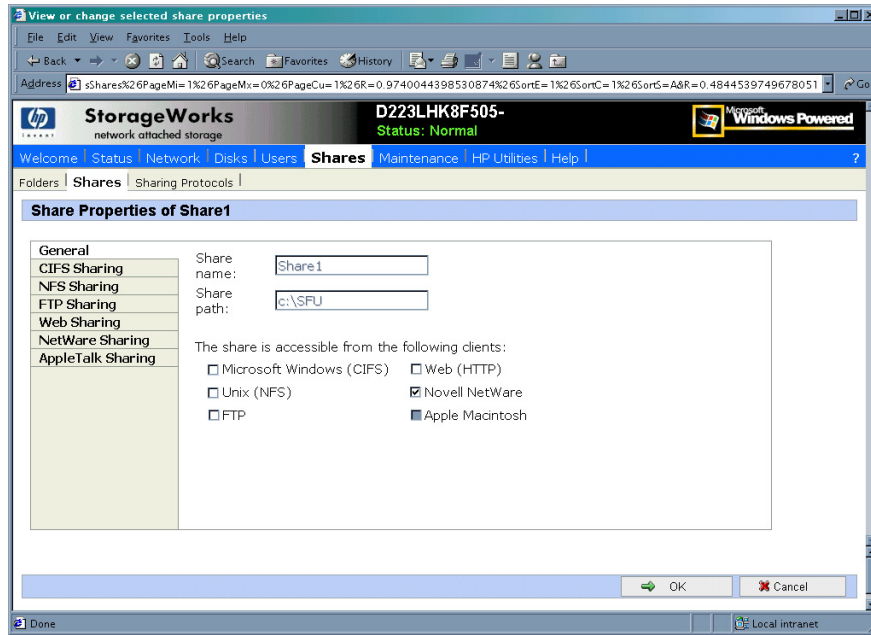
1. From the **Shares** menu, click **Delete**.
2. Verify that this is the correct share and click **OK**.

## Modifying NCP Share Properties

To change share settings:

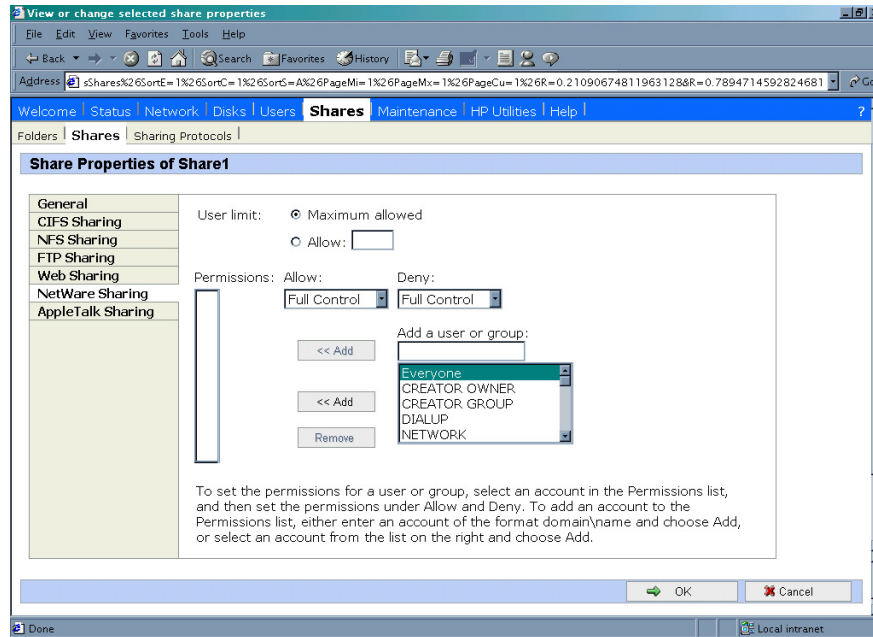
1. From the **Shares** menu, select the share to modify and then click **Properties**. The **General** tab of the **Share Properties** dialog box is displayed.

The name and path of the selected share are displayed.



**Figure 101: 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 102: 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 server, shares can be managed through the NAS Management Console. Tasks include:

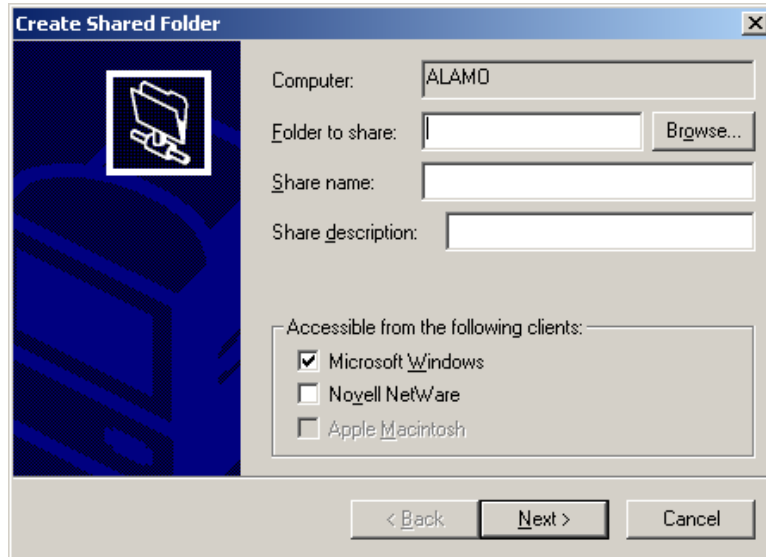
- Creating a new share
- Modifying share properties

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:

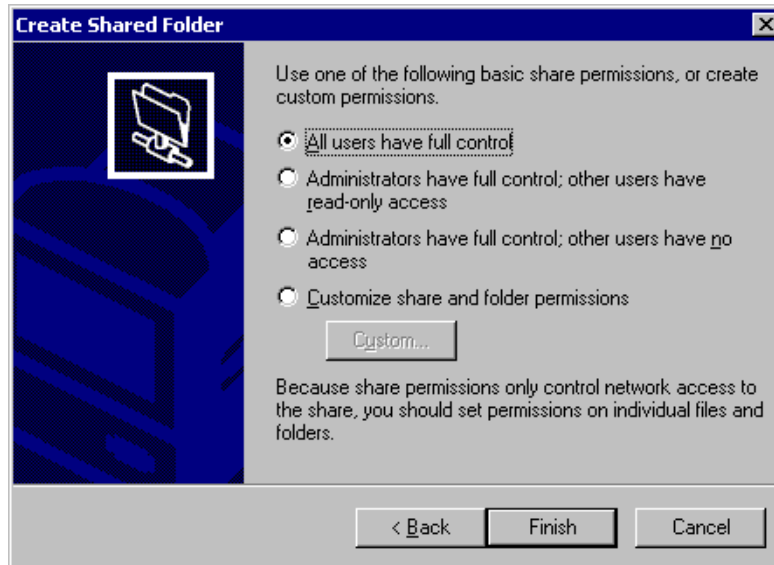
1. From the NAS server desktop, click the **NAS Management Console** icon, click **File Sharing, Shared Folders**, and then **Shares**.
2. Right-click **Shares**, and then click **New File Share**. The **Create Shared Folder** dialog box is displayed.



**Figure 103: Create Shared Folder dialog box**

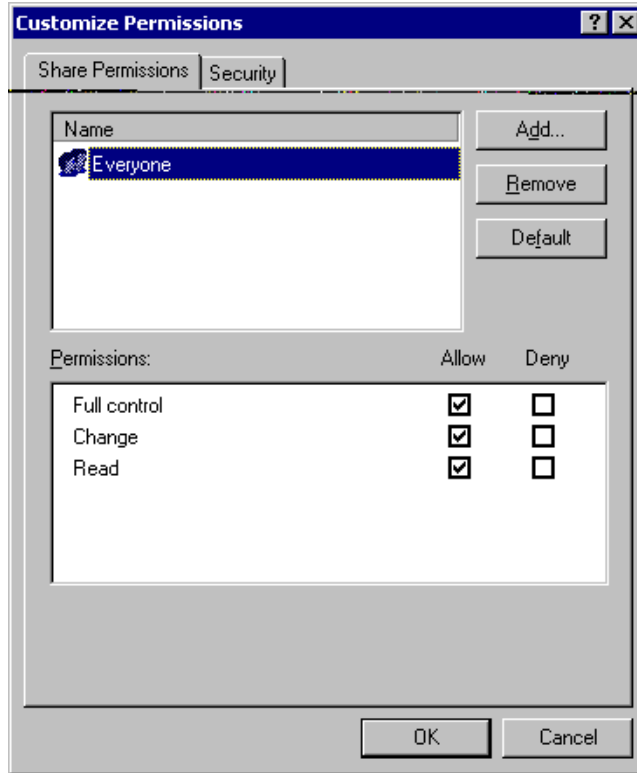
3. In **Folder to Share**, type the path of the directory to be shared.
4. In **Share Name**, type the name of the share. Users will see this name.
5. In **Share Description**, type a description for the share.

6. Select the **Novell NetWare** checkbox and then click **Next**. The dialog box illustrated in [Figure 104](#) is displayed.



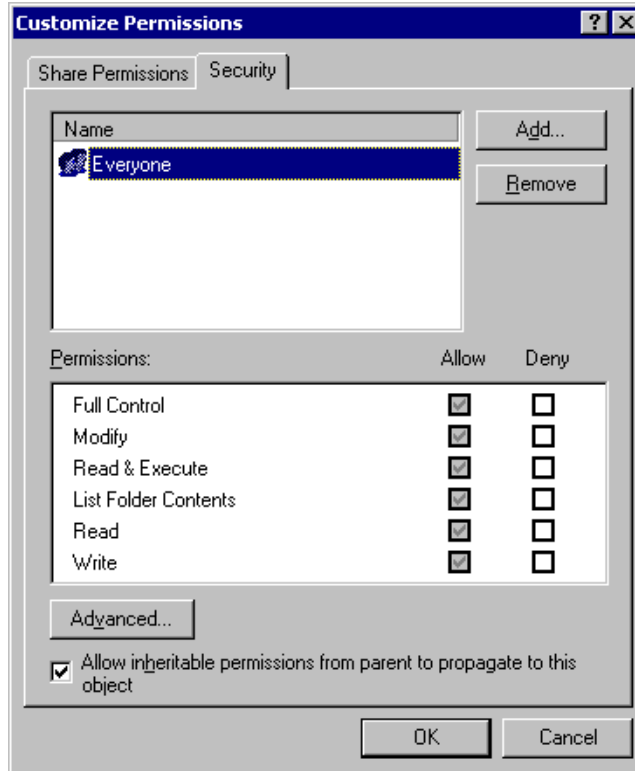
**Figure 104: NetWare Basic Share Permissions dialog box**

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



**Figure 105: 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 106: 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 server 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 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.



# Cluster Administration

## 9

One important feature of the HP StorageWorks NAS server is that it can operate as a single node or as a cluster. This chapter discusses cluster installation and cluster management issues. Some of these topics are discussed or mentioned elsewhere in this guide. The discussion in this chapter is more detailed than other references and addresses the unique administration procedures for operating in a clustered environment.

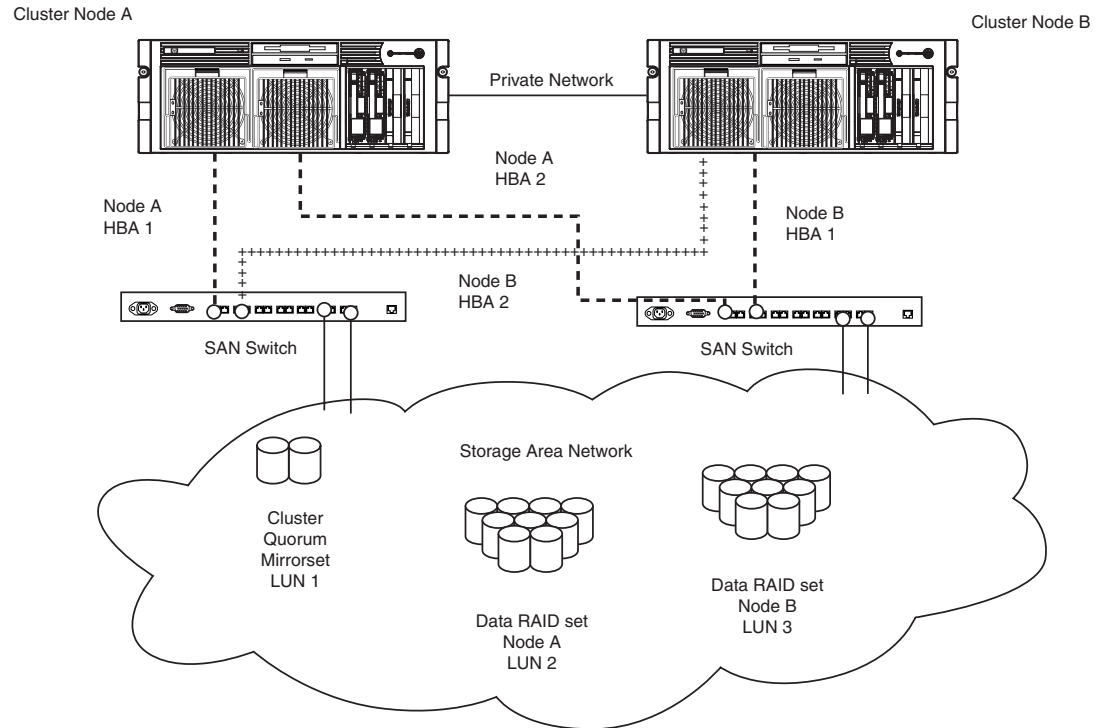
This chapter discusses:

- Cluster Overview
- Cluster Terms and Components
  - Nodes
  - Resource Groups
  - Resources
  - Virtual Servers
  - Failover
  - Quorum Disk
- Cluster Concepts
  - Sequence of Events for Cluster Resources
  - Hierarchy of Cluster Resource Components
- Cluster Planning
  - Storage Planning
  - Network Planning
  - Protocol Planning
- Before Beginning Installation
- HP StorageWorks NAS Software Updates
- Checklists for Cluster Server Installation
  - Shared Disk Requirements
- Cluster Installation
  - Installation Overview
  - Setting Up Networks

- Install Cluster Service Software
  - Configuring the First Node
  - Validating the Cluster Installation
  - Configuring the Second Node
- Verify Installation
- HP StorageWorks NAS Software Updates
- Configuring Physical Disk Resources after the Cluster Install is Complete
- Basic Cluster Administration Procedures
  - Failing Over and Failing Back
  - Restarting One of the Cluster Nodes
  - Shutting Down One of the Cluster Nodes
  - Powering Down both of the Cluster Nodes
  - Powering Up both of the Cluster Nodes
- Cluster Groups and Resources, including File Shares
  - Group Overview
  - Resource Overview
  - File Share Resource Planning Issues
- Persistent Storage Manager
- Non Cluster Aware File Sharing Protocols
- Pathing Software in a Clustered Environment

## Cluster Overview

As introduced in the Quick Start Guide, two server heads (nodes) can be connected to each other and deployed as a no single point of failure (NSPOF) dual redundant cluster. The nodes are connected by a crossover cable and are each connected to network switches or hubs. This connection allows communication between the nodes to track the state of each cluster node. Each node sends out periodic messages to the other node; these messages are called heartbeats. If a node stops sending messages, the cluster service will fail over any resources that the node owns to the other node. For example, if the node that owns the Quorum disk is shut down for any reason, its heartbeat will stop. The other node detects the lack of the heartbeat and takes over ownership of the Quorum disk and the cluster.



**Figure 107: NAS server cluster diagram**

## Cluster Terms and Components

This section provides brief definitions of clustering terms. This information provides basic knowledge of clusters and the terminology used throughout this document.

### Nodes

The most basic parts of a cluster are the server heads. A server node is any individual computer in a cluster or a member of the cluster. If the NAS device is a member of a cluster, then the server heads are referred to as nodes.

### Resources

Hardware and software components that are managed by the cluster service are called cluster resources. Cluster resources have three defining characteristics:

- They can be brought online and taken offline.
- They can be managed in a cluster.
- They can be owned by only one node at a time.

Examples of cluster resources are IP addresses, network names, physical disk resources, and file shares.

### Virtual Servers

A virtual server is a cluster group that consists of a static IP Address resource and a Network Name resource. Several virtual servers can be created. By assigning ownership of the virtual servers to the different server nodes, the processing load on the NAS devices can be distributed between the two nodes.

The creation of a virtual server allows resources dependant on the virtual server to fail over and fail back between the cluster nodes. File Share and physical disks resources are assigned to the virtual server to ensure non disruptive service of file shares to the clients.

### Failover

Failover of cluster groups and resources happens:

- when a node hosting the group becomes inactive. A shutdown of cluster service or a loss of power can cause a failover.
- when all of the resources within the group are dependent on one resource and that resource fails.
- when an administrator forces a failover.

A resource and all of its dependencies must be located in the same group so that if a resource fails over, all of its dependent resources fail over.

When a resource is failed over, the cluster service performs certain procedures. First, all of the resources are taken offline in an order defined by the resource dependencies. Secondly, the cluster service attempts to transfer the group to the next node on the preferred owners list. If the transfer is successful, the resources are brought online in accordance with the resource dependency structure.

The system failover policy defines how the cluster detects and responds to the failure of individual resources in the group. After a failover occurs and the cluster is brought back to its original state, failback can occur automatically based on the policy. After a previously failed node comes online, the cluster service can fail back the groups to the original host. The failback policy must be set before the failover occurs so that failback works as intended.

## Quorum Disk

Each cluster must have a shared disk called the Quorum disk. This physical disk in the common cluster disk array plays a critical role in cluster operations. The Quorum disk offers a means of persistent storage. The disk must provide physical storage that can be accessed by any node in the cluster. If a node has control of the quorum resource upon startup, it can initiate the cluster. In addition, if the node can communicate with the node that owns the quorum resource, it can join or remain in the cluster.

The Quorum disk maintains data integrity by:

- storing the most current version of the cluster database.
- guaranteeing that only one set of active communicating nodes is allowed to operate as a cluster.

## Cluster Concepts

Microsoft cluster concepts are rather straight forward when explained through a diagram. [Figure 108](#) illustrates a typical cluster configuration with the corresponding storage elements. The diagram progresses from the physical disks to the file shares, showing the relationship between both the cluster elements and the physical devices underlying them.

## Sequence of Events for Cluster Resources

The sequence of events in the diagram includes:

1. Physical disks are combined into RAID arrays and LUNs.
2. LUNS are designated as basic disks, formatted and assigned a drive letter via Logical Disk Manager
3. Physical Disks resource are created for each basic disk inside cluster administrator.
4. Directories and folders are created on assigned drives.
5. Cluster components (virtual servers, file shares) are created, organized in groups, and placed within the folders using cluster administrator exclusively.

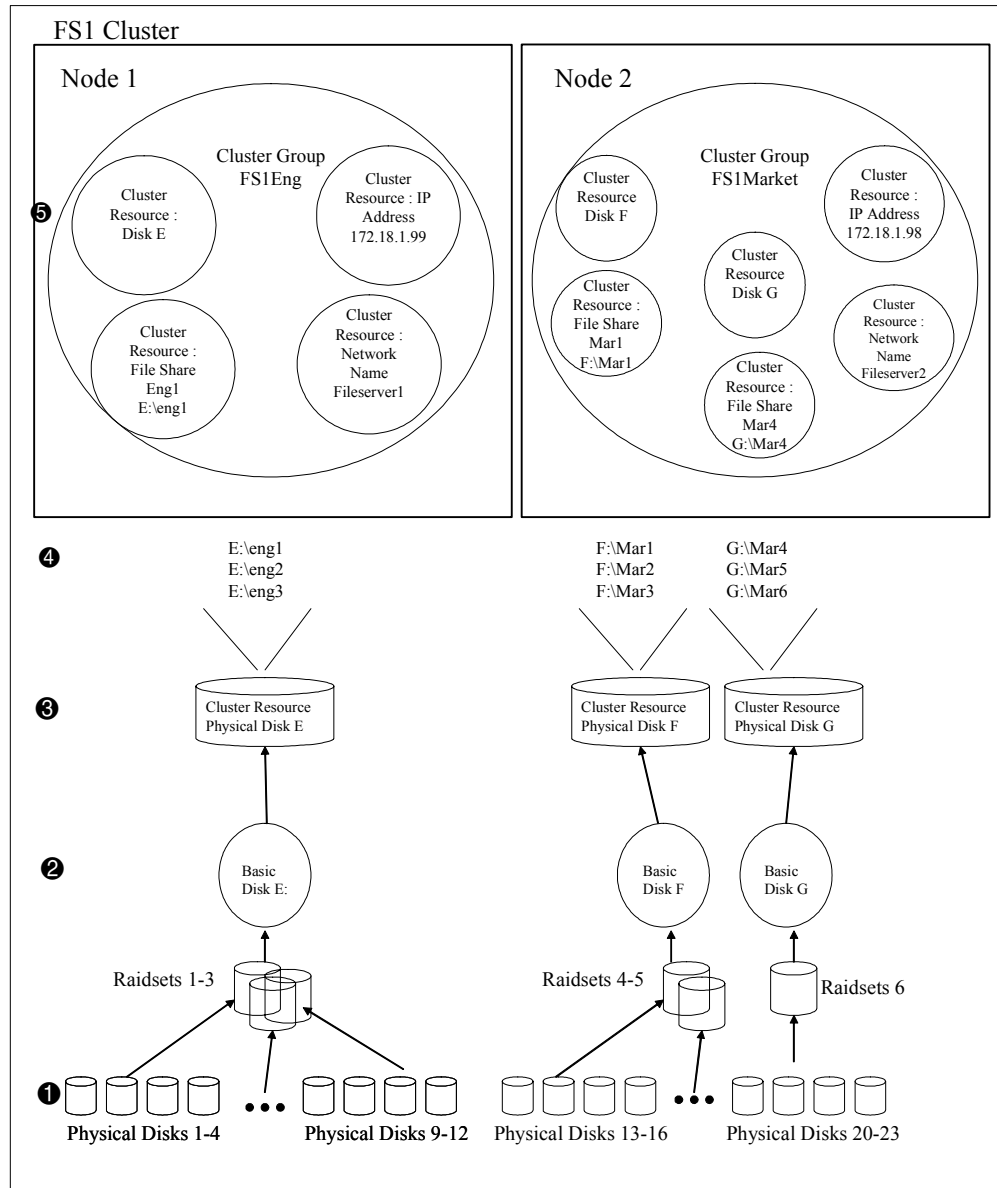


Figure 108: Cluster concepts diagram

## Hierarchy of Cluster Resource Components

The cluster components are referred to as resources and are placed together in groups. Groups are the basic unit of failover between nodes. Resources do not failover individually, rather they failover with the group in which they are contained.

In [Figure 108](#) it is depicted as follows:

- Physical Disk resources are placed in a group and relate to the basic disk created through Logical Disk Manager. It should be noted that when a Physical Disks resource is created through Cluster Administrator a corresponding group should be created for the resource to reside in. Groups are the basic unit of failover on a cluster.
- File Share resources are placed in a group and relate to the actual directory on the drive on which the share is being created.
- An IP Address resource is formed in the group and relates to the IP address by which the group's virtual server is identified on the network.
- A Network Name resource is formed in the group and relates to the name published on the network by which the group is identified.
- A Virtual Server is a group containing an IP Address resource and a Network Name resource. File share and disk resources assigned to this virtual server group can transition from one node to the other during failover conditions.
- The Group is owned by one of the nodes of the cluster, but may transition to the other node during failover conditions.

The diagram illustrates a cluster containing two nodes. Each node has ownership of one group. Contained within each group are singular file shares that are known on the network by the associated Network Name and IP address. In the specific case of Node1, file share Eng1 relates to *E:\Eng1*. This file share is known on the network as *\\Fileserver1\Eng1* with an IP address of 172.18.1.99. *E:\Eng1* relates to the actual Basic Disk E: containing a directory *Eng1*.

For cluster resources to function properly, two very important requirements should be adhered to:

- Dependencies between resources of a group must be established. Dependencies determine the order of startup when a group comes online. In the above case, the following order should be maintained:
  1. File Share—dependent on Physical Disk Resource
  2. NFS File Share—dependent on Physical Disk Resource and Network Name
  3. Network Name—dependent on IP Address

Failure to indicate the dependencies of a resource properly may result in the file share attempting to come online prior to the physical disk resource being available, resulting in a failed file share.

- Groups should have a Network Name resource and an IP Address resource. These resources are used by the network to give each group a virtual name. Without this virtual reference to the group, the only way to address a share that is created as a clustered resource is by node name. Physical node names do not transition during a failover, whereas virtual names do.

For example, if from a client a network share map F: was established and assigned to `\\Node1\Eng1` instead of `\\Fileserver\Eng1`, when Node1 fails and Node2 assumes ownership, the map will become invalid because the reference in the map is to `\\Node1`. If the map were created to the virtual name and Node1 were to fail, the map would still exist when the group associated with Eng1 failed over to Node2.

The previous diagram is an example and is not intended to imply limitations of a single group or node. Groups can contain multiple physical disks resources and file shares and nodes can have multiple groups, as shown by the group owned by Node2.

## Cluster Planning

Clustering the NAS b3000 v2 or e7000 v2 greatly enhances the availability of file service by enabling file shares to fail over to a second NAS device, if problems arise. Clients see only a brief interruption of service as the file share resource transitions from one server node to the other.

Requirements for taking advantage of clustering include:

- Storage planning
- Network planning
- Protocol planning

## Storage Planning

For clustering, a storage unit (LUN) must be designated for the cluster and configured as a mirrorset. This LUN is used for the Quorum disk. The Quorum disk is the shared storage used by the cluster nodes to coordinate the internal cluster state.

One or more RAID arrays are dedicated to each cluster node for data storage. Each cluster node will assume ownership of at least one physical disk resource. That owner node will serve all shares within that physical disks resource, until a failover condition occurs. When a failover occurs, the physical disk resource and all associated shares will transition over to the remaining node and will remain there until the other node is returned to service. Some types of shares are not cluster aware and will not be available during a failover condition. See the “Protocol Planning” section for additional information.

To prepare a basic disk for use in a cluster, a cluster group for each basic disk should be created to allow each resource to failover separately. Once the group is created, a physical disk resource is created in each of the groups. Cluster groups may contain more than one physical disk depending on the site-specific requirements. This physical disk resource is required for basic disk to successfully work in a cluster environment protecting it from simultaneous access from each node.

---

**Note:** The LUN underlying the basic disk should be presented to only one node of the cluster using selective storage presentation SAN switch zoning or having only one node online at all times until such times as the physical resource for the basic disk is established.

---



In preparing for the cluster installation:

- All software components listed in the SAN connection tool must be installed and the fiber cables attached to the HBA(s) before the cluster installation is started.
- All shared disks, including the quorum disk, must be accessible from both nodes. When testing connectivity between server and LUN, only one server should be given access to the LUN at a time or the non-testing server should be powered off.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

## Network Planning

Clusters require more sophisticated networking arrangements than a stand alone NAS device. For example, because a cluster must be deployed into a domain environment, workgroups are not supported. A Windows NT domain or Active Directory domain must be in place to contain the cluster names, virtual server names, and user and group information. A cluster cannot be deployed into a non domain environment.

All cluster deployments have at least seven network addresses and network names:

- The cluster name (Unique NETBIOS Name) and IP address
- Node A's name and IP address
- Node B's name and IP address
- At least one virtual server name and IP address for Node A
- At least one virtual server name and IP address for Node B
- Cluster Interconnect static IP addresses for Node A and Node B.

Virtual names and addresses are the only identification used by clients on the network. Because the names and addresses are virtual, their ownership can transition from one node to the other during a failover, preserving access to the shares on the virtual disks.

In addition, a cluster will use at least two network connections on each node:

- The cluster interconnect or “heartbeat” crossover cable connects to the first network port on each cluster node.
- The client network subnet connects to a second network port on each cluster node. The cluster node names and virtual server names will have IP addresses residing on these subnets.

---

**Note:** If the share is to remain available during a failover, each cluster node must be connected to the same network subnet. It is impossible for a cluster node to serve the data to a network to which it is not connected.

---

## Protocol Planning

The NAS b3000 v2 and e7000 v2 both support many file sharing protocols, including sharing protocols for Windows, UNIX, Linux, Novell, Macintosh, Web, and FTP clients. However, not all of these protocols can take advantage of clustering. If a protocol does not support clustering, the share will not be available to the clients until the owner cluster node is brought back online.

HP recommends placing cluster aware and non cluster aware protocols on different file shares.

Use the information in [Table 7](#) to determine whether it is advantageous to use clustering.

**Table 7: Sharing Protocol Cluster Support**

| Protocol  | Client Variant   | Cluster Aware (supports failover) | Supported |
|-----------|--|-----------------------------------|-----------|
| CIFS      | Windows NT<br>Windows 2000<br>Windows 95<br>Windows 98<br>Windows ME | Yes                               | Yes       |
| NFS       | UNIX<br>Linux  | Yes                               | Yes       |
| HTTP      | Web  | No                                | Yes       |
| FTP       | Many   | Yes                               | Yes       |
| NCP       | Novell   | No                                | Yes       |
| AppleTalk | Apple  | No                                | No        |

---

**Note:** AppleTalk is not supported on clustered disk resources. AppleTalk requires local memory for volume indexing. On failover events, the memory map is lost and data corruption can occur.

---

---

## Before Beginning Installation

This section provides the steps necessary to cluster HP StorageWorks NAS servers. Confirm that the following specifications have been met before proceeding:

- The SAN connection tool must be completed and all the necessary software components for connecting to the desired storage must be installed before the configuration of cluster services. The SAN connection tool can be found under the HP Utilities tab in the WebUI.
- It is required that at least one LUN has been presented for the configuration of the Quorum disk. Additional LUNS may also be presented for use as shared disk resources.

More detailed information about setting up clusters is available at <http://www.microsoft.com/windows2000/techinfo/planning/server/clustersteps.asp>

Each node of the cluster will be shut down during the Cluster Setup.

## HP StorageWorks NAS Software Updates

Please be sure to read the section further in this chapter that discusses the process to update HP StorageWorks NAS software to clustered versions. This should only be done after the cluster install is complete and must be run on **both cluster nodes**.

## Checklist for Cluster Server Installation

This checklist assists in preparing for installation. Step-by-step instructions begin after the checklist.

## Shared Disk Requirements

---

**Note:** Do not let both nodes access the shared storage devices at the same time until Cluster service is installed on at least one node and that node is online. This can be accomplished through selective storage presentation, SAN switch zoning or having only one node online at all times.

---

- All software components listed in the SAN connection tool must be installed and the fiber cables attached to the HBA(s) before the cluster installation is started.
- All shared disks, including the quorum disk, must be accessible from both nodes.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

# Cluster Installation

## Installation Overview

During the installation process, both nodes will be shut down and both nodes will be rebooted. These steps are necessary to guarantee that the data on disks that are attached to the shared storage bus is not lost or corrupted. This can happen when multiple nodes try to simultaneously write to the same disk that is not yet protected by the cluster software.

Use [Table 8](#) to determine which nodes and storage devices should be presented during each step.

**Table 8: Power Sequencing for Cluster Installation**

| Step                         | Node 1 | Node 2 | Storage       | Comments  |
|------------------------------|--------|--------|---------------|---|
| Setting Up Networks          | On     | On     | Not Presented | Verify that all storage devices on the shared bus are not presented. Power on both nodes. |
| Setting up Shared Disks      | On     | Off    | Presented     | Shutdown both nodes. Present the shared storage, then power on the first node.            |
| Verifying Disk Configuration | Off    | On     | Presented     | Shut down first node, power on second node.   |
| Configuring the First Node   | On     | Off    | Presented     | Shutdown all nodes; power on the first node.  |
| Configuring the Second Node  | On     | On     | Presented     | Power on the second node after the first node was successfully configured.                |
| Post-installation            | On     | On     | Presented     | At this point all nodes should be on.   |

To configure the Cluster service on the HP StorageWorks NAS server, the account must have administrative permissions on each node. Both nodes must be member servers within the same domain. It is not acceptable to have a mix of domain controllers and member servers in a cluster.

---

**Note:** Do not let both nodes access the shared storage devices at the same time until Cluster service is installed on at least one node and that node is online. This can be accomplished through selective storage presentation, SAN switch zoning or having only one node online at all times.

---

## Setting Up Networks

Each cluster node requires at least two network adapters—one to connect to a public network, and one to connect to a private network consisting of cluster nodes only.

The private network adapter establishes node-to-node communication, cluster status signals, and cluster management. Each node's public network adapter connects the cluster to the public network where clients reside.

Verify that all network connections are correct, with private network adapters connected to other private network adapters only, and public network adapters connected to the public network.

## Configure the Private Network Adapter

The following procedures are Best Practices provided by Microsoft and should be configured on the private network adapter.

- On the General tab of the private network adapter, deselect all items except TCP/IP
- In the DNS tab under advanced settings for the private network adapter, deselect **Register this connection's addresses in DNS**
- Set the Link Speed and Duplex to 100Mps/Full Duplex under the advanced tab for the Ethernet card used for the private network adapter
- If a network adapter is in a disconnected state during installation, the Cluster service does not detect the adapter because there are no protocols bound to the adapter. If Media Sense is disabled the network adapter still shows the “disconnected” status, but the cluster installation process can detect the adapter as available for cluster communication. To make this change, refer to the online guide.

If the private network connection is made using a crossover cable, then the procedures outlined in Knowledge Base (KB) article Q242430 (<http://support.microsoft.com/support/kb/articles/Q242/4/30.ASP>) should be followed and the node rebooted prior to installing Cluster service. If this procedure is not completed, and the second node is powered off while installing Cluster service on the first node, the private network adapter may not be detected.

This will prevent configuring the adapter during Cluster service installation. However, after Cluster service is installed on both nodes and both nodes are powered on, the adapter can be added as a cluster resource and configured properly for the private network in Cluster Administrator.

## Configure the Public Network Adapter

While the public network adapter's IP address can be automatically obtained if a DHCP server is available, this is not recommended for cluster nodes. We strongly recommend setting static IP addresses for all network adapters in the cluster, both private and public. If IP addresses are obtained via DHCP, access to cluster nodes could become unavailable if the DHCP server goes down. If DHCP must be used for the public network adapter, use long lease periods to assure that the dynamically assigned lease address remains valid even if the DHCP service is temporarily lost. In all cases, set static IP addresses for the private network connector. Keep in mind that Cluster service will recognize only one network interface per subnet. For assistance with TCP/IP addressing in Windows 2000, please see Windows 2000 Online Help (<http://www.microsoft.com/windows2000/techinfo/proddoc/default.asp>).

## Rename the Local Area Network Icons

HP recommends changing the names of the network connections for clarity. The naming will help to identify a network and correctly assign its role.

## Verifying Connectivity and Name Resolution

To verify name resolution, ping each node from a client using the node's machine name instead of its IP number.

## Verifying Domain Membership

Both nodes in the cluster must be members of the same domain and able to access a domain controller and a DNS Server.

## Setting Up a Cluster User Account

The Cluster service requires a domain user account under which the Cluster service can run. This user account must be created before installing Cluster service, because setup requires a user name and password. This user account should not belong to a user on the domain. This user account will need to be granted administrator privileges.

## About the Quorum Disk

To proceed, power off both nodes. Power up node one.

The quorum disk is used to store cluster configuration database checkpoints and log files that help manage the cluster. The quorum disk must be a shared disk resource. We make the following quorum disk recommendations:

- Create a small partition [A minimum of 50 megabytes (MB) to be used as a quorum disk. We generally recommend a quorum disk to be 500 MB.]
- Dedicate a separate disk resource for a quorum disk. As the failure of the quorum disk would cause the entire cluster to fail, it is strongly recommended that the disk resource be a RAID 1 configuration.

During the Cluster service installation, a drive letter must be provided for the quorum disk. We recommend the drive letter Q for the quorum disk.

## Configuring Shared Disks

Use Disk Management to configure the quorum and shared disk resources. Verify that all shared disks are formatted as **NTFS** and are designated as **Basic**.

## Verifying Disk Access and Functionality

Write a file to each shared disk resource to verify functionality.

At this time, shut down the first node, power on the second node and repeat the Verifying Disk Access and Functionality step above. When it has been verified that both nodes can read and write from the disks, turn off the second node and power on the first, and then continue with this guide.

# Install Cluster Service Software

## Configuring the First Node

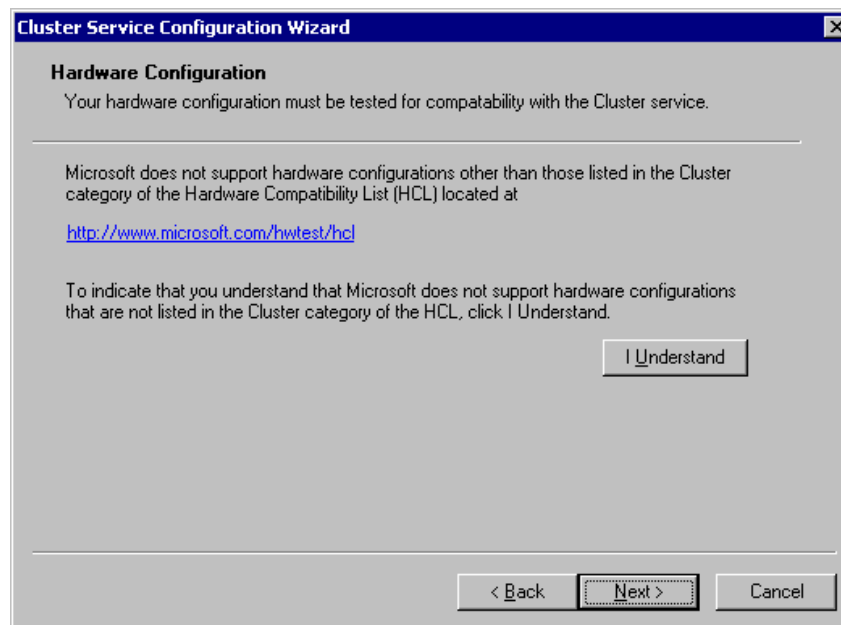
---

**Note:** During installation of Cluster service on the first node, the second node must either be turned off, or stopped prior to the HP StorageWorks NAS server booting. All shared storage devices should be powered up.

---

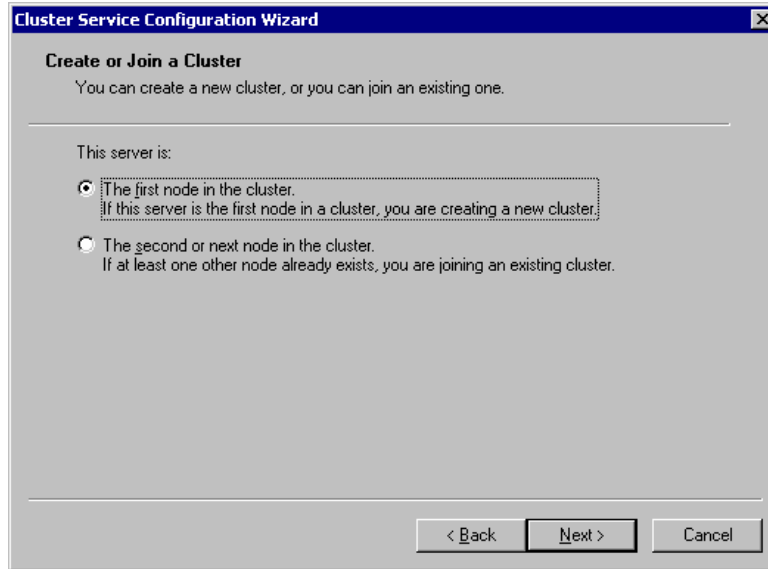
In the first phase of installation, all initial cluster configuration information must be supplied so that the cluster can be created. This is accomplished using the **Cluster Service Configuration Wizard** via the iLO board or Terminal Services.

1. Click **Start**, click **Settings**, and click **Control Panel**.
2. Double-click **Add/Remove Programs**.
3. Double-click **Add/Remove Windows Components**.
4. Select **Cluster Service**. Click **Next**.
5. When prompted for files on the Windows Powered OS CD or the Windows 2000 SP3 cd, browse to the directory *C:\Compaq\Windows Components\i386*.
6. The window shown in [Figure 109](#) below appears. Click **I Understand** to accept the condition that Cluster service is supported on hardware from the Hardware Compatibility List only.



**Figure 109:** Hardware Configuration certification screen

7. Because this is the first node in the cluster, the cluster must be created. Select **The first node in the cluster**, as shown in [Figure 110](#) and then click **Next**.



**Figure 110: Create new cluster**

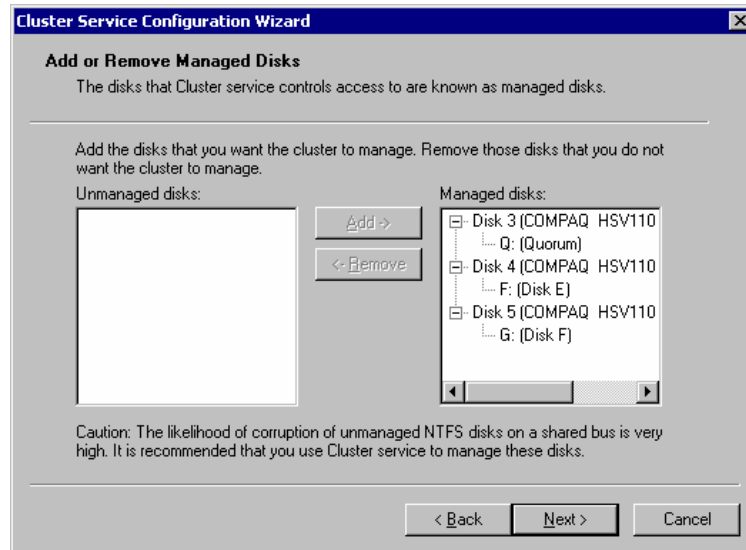
8. Enter a name for the cluster (up to 15 characters), and click **Next**.
9. Type the **user name** and **password** of the cluster service account that was created during the pre-installation. Type the **domain name**, and click **Next**.  
At this point the **Cluster Service Configuration** Wizard validates the user account and password.
10. Click **Next**.



## Configuring Cluster Disks

1. The **Add or Remove Managed Disks** dialog box shown in [Figure 111](#) specifies which shared disks will be used by Cluster service. Add or remove disks as necessary and then click **Next**.

**Note:** If a managed disk does not appear in the managed disk list, then a reboot is required for it to appear.



**Figure 111: Add or remove managed disks**


In production clustering scenarios more than one private network for cluster communication must be used to avoid having a single point of failure. Cluster service can use private networks for cluster status signals and cluster management. This provides more security than using a public network for these roles. It is possible to use a public network for cluster management, or use a mixed network for both private and public communications. In any case, make sure at least two networks are used for cluster communication, as using a single network for node-to-node communication represents a potential single point of failure. HP recommends that multiple networks be used, with at least one network configured as a private link between nodes and other connections through a public network. If there is more than one private network, make sure that each uses a different subnet, as Cluster service recognizes only one network interface per subnet.

This document is built on the assumption that only two networks are in use. It shows how to configure these networks as one mixed and one private network.

The order in which the Cluster Service Configuration Wizard presents these networks may vary. In this example, the public network is presented first.

2. Click **Next** in the **Configuring Cluster Networks** dialog box.
3. Make sure that the network name and IP address correspond to the network interface for the *public* network.
4. Check the box **Enable this network for cluster use**.

5. Select the option **All communications (mixed network)** as shown in [Figure 112](#).
6. Click **Next**.




The dialog box is titled "Cluster Service Configuration Wizard" and has a "Network Connections" section. It contains the following fields and options:

- Network name: Public Cluster Connection
- Device: Compaq NC3134 Fast Ethernet NIC #10
- IP address: 172.18.1.243
- Enable this network for cluster use
- This network performs the following role in the cluster:
  - Client access only (public network)
  - Internal cluster communications only (private network)
  - All communications (mixed network)

Buttons at the bottom: < Back, Next >, Cancel

**Figure 112: Public network connection**

7. The next dialog box shown in [Figure 113](#) configures the private network. Make sure that the network name and IP address correspond to the network interface used for the *private* network.
8. Check the box **Enable this network for cluster use**.
9. Select the option **Internal cluster communications only**.



The dialog box is titled "Cluster Service Configuration Wizard" and has a "Network Connections" section. It contains the following fields and options:

- Network name: Cluster
- Device: Compaq NC3134 Fast Ethernet NIC #9
- IP address: 10.10.10.53
- Enable this network for cluster use
- This network performs the following role in the cluster:
  - Client access only (public network)
  - Internal cluster communications only (private network)
  - All communications (mixed network)

Buttons at the bottom: < Back, Next >, Cancel

**Figure 113: Private network connection**

10. Click **Next**.

11. In this example, both networks are configured in such a way that both can be used for internal cluster communication. The next dialog window offers an option to modify the order in which the networks are used. Because **Private Cluster Connection** represents a direct connection between nodes, it is left at the top of the list. In normal operation this connection will be used for cluster communication. In case of the **Private Cluster Connection** failure, cluster service will automatically switch to the next network on the list—in this case **Public Cluster Connection**. Make sure the first connection in the list is the Private Cluster Connection and click **Next**.



**Caution:** Always set the order of the connections so that the Private Cluster Connection is first in the list.

12. Enter the unique cluster IP address and Subnet mask, and click **Next**.

**Figure 114: Cluster IP Address**

The **Cluster Service Configuration** Wizard shown in [Figure 114](#) automatically associates the cluster IP address with one of the public or mixed networks. It uses the subnet mask to select the correct network.

13. Click **Finish** to complete the cluster configuration on the first node.

The **Cluster Service Setup** Wizard completes the setup process for the first node by copying the files needed to complete the installation of Cluster service. After the files are copied, the Cluster service registry entries are created, the log files on the quorum resource are created, and the Cluster service is started on the first node.

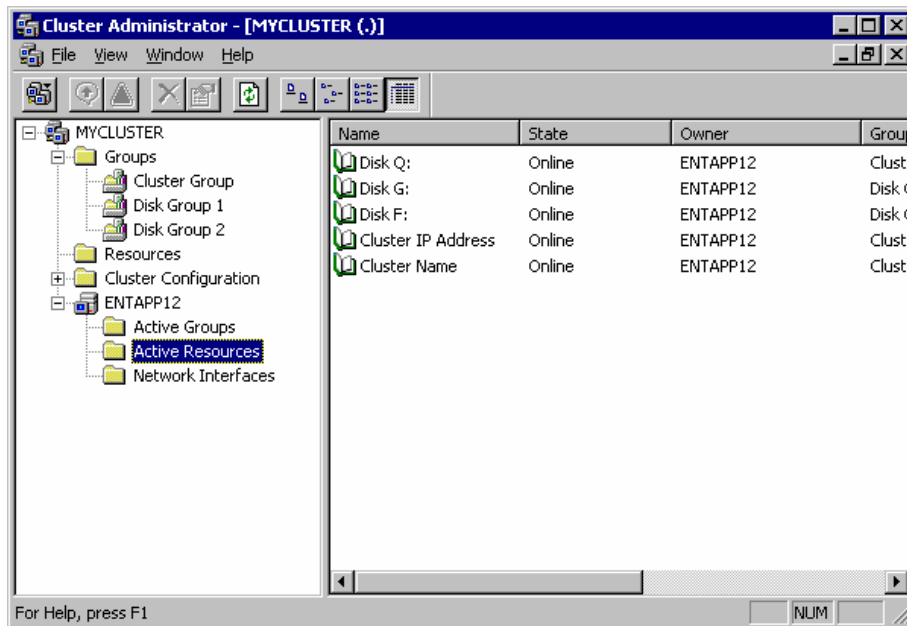
A dialog box appears showing that Cluster service has started successfully.

14. Click **OK**.
15. Close the **Add/Remove Programs** window.

## Validating the Cluster Installation

Use the Cluster Administrator snap-in to validate the Cluster service installation on the first node.

1. Click **Start>Programs>Administrative Tools**, and click **Cluster Administrator** via Terminal Services or the iLO port. Alternatively, a cluster administrator link is provided in the Cluster tab of the WebUI.



**Figure 115: Cluster administrator**

If the displayed snap-in window is similar to that shown in [Figure 115](#), the Cluster service was successfully installed on the first node. It is now possible to install the Cluster service on the second node.

## Configuring the Second Node

---

**Note:** For this section, leave the first node on and power up the second node.

---

Installing Cluster service on the second node requires less time than on the first node. Setup configures the Cluster service network settings on the second node based on the configuration of the first node.

Installation of Cluster service on the second node begins exactly as for the first node. During installation of the second node, the first node must be running.

Follow the same procedures used for installing Cluster service on the first node, with the following differences:

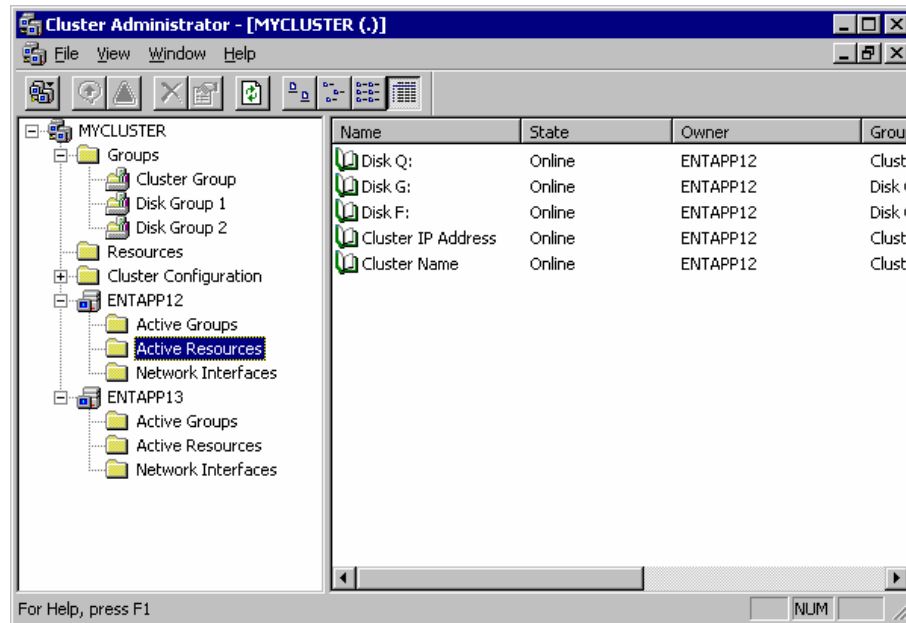
1. In the **Create or Join a Cluster** dialog box, select **The second or next node in the cluster**, and click **Next**.
2. Enter the cluster name that was previously created (in this example, **MyCluster**), and click **Next**.

3. Leave **Connect to cluster** as unchecked. The Cluster Service Configuration Wizard will automatically supply the name of the user account selected during the installation of the first node. Always use the same account used when setting up the first cluster node.
4. Enter the password for the account (if there is one) and click **Next**.
5. At the next dialog box, click **Finish** to complete configuration.
6. The Cluster service will start. Click **OK**.
7. Close **Add/Remove Programs**.

## Verify Installation

There are several ways to verify a successful installation of Cluster service. Here is a simple one:

1. Click **Start>Programs>Administrative Tools**, and click **Cluster Administrator**.



**Figure 116: Cluster resources**

The presence of two nodes (Entapp12 and Entapp13 in [Figure 116](#)) shows that a cluster exists and is in operation.

2. Right-click the cluster group and select the option **Move**. The group and all its resources will be moved to the other node. After a short period of time the cluster resources will be brought online on the second node. Watch the screen to see this shift. Close the **Cluster Administrator** snap-in.

If this test fails then the cluster configuration was not successful. Further information can be found in the event logs and it may be necessary to reinstall cluster services on one of the nodes.

Installation of Cluster service on both nodes is complete. The server cluster is fully operational.

## HP StorageWorks NAS Software Updates

After cluster installation is complete it is necessary to upgrade Services for UNIX and NAS DataCopy (if installed) to make them cluster aware. Select the link at the bottom of the Cluster Setup Guide located under the HP Utilities tab in the WEB User Interface to perform the software updates.

Complete this procedure on both cluster nodes.

## Configuring Physical Disk Resources after the Cluster Install is Complete

After the cluster installation is complete it may be necessary to add additional physical disk resources into the cluster. The procedure below outlines the steps to complete this process. This can be accomplished at the server console or remotely through terminal services or the ILO board.

- Present the LUN(s) to Node A, do not present the LUN(s) to Node B at this time.
- On Node A, open Disk Manager and write the disk signatures and create the partitions. Remember that disks can only be basic in a cluster configuration.
  - To open Disk Manager select **Start>Run>diskmgmt.msc**
  - A volume label must be assigned when the partition is formatted.
- Ensure that all cluster resources reside on Node B.
- Reboot Node A; this is necessary so that the cluster service will recognize the newly created partitions.
- After Node A has rebooted open Cluster Administrator and create the Physical Disk Resources.
  - Select **Start>Run>Cluadmin** to open Cluster Administrator
  - Move the cluster group the new physical disk resource needs to reside in to Node A, or create a new group on Node A.
  - To create a physical disk resource in Cluster Administrator select **File>New>Resource**.
  - If the new disks do not appear in the disk selection list, Node A must be rebooted.
- After all new LUN(s) have been added to the cluster, present the LUN(s) to Node B using selective storage presentation or switch zoning.
- On Node B, open Device Manager and scan for new devices.
- Move the new Physical Disk Resources to Node B
  - To move the Physical Disk Resource, right-click the cluster group the disk resides in and select **Move Group**.

---

## Basic Cluster Administration Procedures

- Failing over and failing back
- Restarting one cluster node
- Shutting down one cluster node
- Powering down both cluster nodes
- Powering up both cluster nodes

### Failing Over and Failing Back

As previously mentioned, when a node goes offline, all of the resources dependent on that node are automatically failed over to the other node. Processing continues, but in a reduced manner because all operations must be processed on the remaining node.

Because operating environments differ, the administrator must indicate whether the system will automatically fail the resources (organized by resource groups) back to their original node or will leave the resources failed over, waiting for the resources to be moved back manually. See “Managing Cluster Resource Groups” for information on allowing or preventing failback and moving these resources from one node to another.

---

**Note:** If the NAS server is not set to automatically fail back the resources to their designated owner, the resources must be moved back manually each time a failover occurs. See “Managing Cluster Resource Groups” for information on overriding this default setting.

---

### Restarting One Cluster Node



**Caution:** Restarting a cluster node should be done only after confirming that the other node in the cluster is functioning normally. Adequate warning should be given to users connected to resources of the node being restarted.

Attached connections can be viewed through the NAS Management Console on the NAS Desktop using Terminal Services. From the NAS Management Console, select File Sharing, Shared Folders, and Sessions.

---

The physical process of restarting one of the nodes of a cluster is the same as restarting a NAS device in single node environment. However, additional caution is needed.

Restarting a cluster node causes all file shares served by that node to fail over to the other node in the cluster. Until the failover process completes, any currently executing read and write operations will fail. The other node will be placed under a heavier load by the extra work until the restarted node comes up.

## Shutting Down One Cluster Node



**Caution:** Shutting down a cluster node must be done only after confirming that the other node in the cluster is functioning normally. Adequate warning should be given to users connected to resources of the node being restarted.

---

Shutting down a cluster node causes file shares served by that node to fail over to the other node. This will cause any currently executing client read and write operations to fail until the cluster failover process completes. The other node will be placed under a heavier load by the extra work until the second node is powered up and rejoins the cluster.

## Powering Down Both Cluster Nodes

The power down process for the NAS cluster is similar to the process for a single node, but with the cluster, extra care must be taken with the storage subsystem and the sequence of the shutdown.

The power down process is divided into two main steps:

1. Shutting down the cluster nodes
2. Removing power from the cluster nodes

The sequence of these steps is critical. The devices must be shut down before the storage subsystem. Improperly shutting down the nodes and the storage subsystem causes corruption and loss of data.



**Caution:** Before powering down the cluster nodes, follow the proper shutdown procedure as previously illustrated. See "Shutting Down One Cluster Node." Only one cluster node should be shut down at a time.

---



**Caution:** The cluster nodes should never be powered on when the storage subsystem is not available.

---

## Powering Up Both Cluster Nodes

The power up process for the NAS cluster is more complex than it is for a single node because extra care must be taken with the storage subsystem.

The sequence of the power up steps is critical. Improper power up procedures can cause corruption and loss of data.



**Caution:** Do not power up the cluster nodes without first powering up the storage subsystem, and verifying it is operating normally.

---



Nodes should be powered up separately allowing one node to form the cluster prior to powering up the second node. To power up the cluster nodes:

1. After the storage subsystem is confirmed to be operating normally, power up a single node by pressing the power button on the front of the device. Wait for the node to come completely up before powering up the second node.

If both nodes are powered up at the same time, the first node that completes the sequence will gain ownership of the cluster quorum and will control the cluster database. Designate a particular node as the usual cluster quorum owner by always powering up that node first and letting it completely restart before powering up the second cluster node.

2. Power up the second cluster node by pressing the power button on the front of the device.

As each node starts, the monitor displays the logon dialog. Background processes will start the cluster service and form the cluster.

## Cluster Groups and Resources, including File Shares

Management tasks for a cluster include creating and managing cluster resources and cluster groups. The Microsoft Cluster Administrator Tool provides complete online help for all cluster administration activities. As mentioned previously, cluster resources are created and then assigned to logical, organizational groups. Ownership of these groups should be assigned in a balanced arrangement between the server nodes, distributing the processing load between the two nodes.

Cluster resources include administrative types of resources as well as file shares. The following paragraphs include overview and planning issues for cluster groups, cluster resources, and clustered file shares.

Creating and managing these resources and groups must be managed through Cluster Administrator, available from the Cluster tab of the WebUI. Complete online help for creating the various cluster objects is available in the Cluster Administrator tool.

### Cluster Group Overview

A default cluster group is automatically created when the cluster is first created. This default cluster group contains an Internet Protocol (IP) Address resource, a Network Name resource, and the Quorum disk resource. When the new cluster is created, the (IP) address and the cluster name that were specified during setup are set up as the IP address and network name of this default cluster group.



**Caution:** Do not delete or rename the Cluster Group or IP Address. Doing so will result in losing the cluster and will require reinstallation of the cluster.

---

When creating groups, the administrator's first priority is to gain an understanding of how to manage the groups and their resources. Administrators may choose to create a resource group and a virtual server (IP Address resource and Network Name resource) for each node that will contain all resources owned by that node, or the administrator may choose to create a resource group and virtual server for each physical disk resource. Additionally, the administrator should try to balance the load of the groups and their resources on the cluster between the two nodes.

## Node Based Cluster Groups

Creating only one resource group and one virtual server for each node facilitates group and resource administration. This setup allows administrators to include all file share resources under one group. Clients access all of the resources owned by one node through a virtual server name.

In node based cluster groups, each group has its own network name and IP address. The administrator decides on which node to place each physical disk resource. This configuration provides a very coarse level of granularity. All resources within a group must remain on the same node. Only two IP addresses and network names are required. This configuration creates less overhead for resource and network administration. A possible disadvantage of this approach is that the resource groups can potentially grow large when many file shares are created.

## Load Balancing

The creation of separate cluster groups for each virtual server provides more flexibility in balancing the processing load on the cluster between the two nodes. Each cluster group can be assigned to a cluster node with the preferred owner parameter. For example, if there are two cluster groups, the cluster could be set up to have the first cluster group owned by node A and the second cluster group owned by node B. This allows the network load to be handled by both devices simultaneously. If only one cluster group exists, it can only be owned by one node and the other node would not serve any network traffic.

## Cluster Resource Overview

Hardware and software components that are managed by the cluster service are called cluster resources.

Resources represent individual system components. These resources are then organized into groups and managed as a group.

Some resources are created automatically by the system and other resources must be set up manually.

Resource Types:

- IP Address resource
- Cluster name resource
- Cluster Quorum disk resource
- Physical Disk resource
- Virtual server name resources
- CIFS file share resources
- NFS file share resources

## File Share Resource Planning Issues

CIFS and NFS are cluster aware protocols that support the Active/Active cluster model, allowing resources to be spread out and processed on both nodes at the same time. For example, some NFS file share resources can be assigned to a group owned by a virtual server for NodeA and additional NFS file share resources can be assigned to a group owned by a virtual server for NodeB.

Configuring the file shares as cluster resources provides for high availability of file shares. Because the resources are placed into groups, ownership of the files can easily move from one node to the other, as circumstances require. If the cluster node owning the group of file shares should be shut down or fail, the other node in the cluster will begin sharing the directories until the original owner node is brought back on line. At that time, ownership of the group and its resources can be brought back to the original owner node.

### Resource Planning

1. Create at least one virtual server for each node in the cluster.  
A virtual server is a resource group consisting of an IP Address resource and a Network Name resource. Ownership of these virtual servers should be assigned to the different server nodes. In addition to providing load balancing capabilities, the virtual server allows for the transition of group resources in failover situations.
2. Create a virtual server group for each node in the cluster.  
Cluster resource groups are used to balance the processing load on the servers. Distribute ownership of the groups between the virtual servers.
3. For NFS environments, configure the NFS server.  
NFS specific procedures include entering audit and file lock information as well as setting up client groups and user name mappings. These procedures are not unique to a clustered deployment and are detailed in the “UNIX File System Management” chapter. Changes to NFS setup information are automatically replicated to all nodes in a cluster.
4. Create the file share resources.  
In a clustered environment, file shares are created as a type of cluster resource. Creating cluster resources and file shares is documented later in this chapter.
5. Assign ownership of the file share resources to the resource groups.
  - a. Divide ownership of the file share resource between the resource groups, which are in turn distributed between the virtual servers, for effective load balancing.
  - b. Make sure that the physical disk resource for this file share is also included in this group.
  - c. Make sure that the resources are dependent on the virtual servers and physical disk resources from which the file share was created.

### Permissions and Access Rights on Share Resources

File Share and NFS Share permissions must be managed via the Microsoft Cluster Administrator tool versus the individual shares on the file system themselves via Windows Explorer. Administering them through the Cluster Administrator tool allows the permissions to migrate from one node to other. In addition, permissions established using Explorer will be lost once the share is failed or taken offline. To access the permissions, open up Cluster Administrator, right click on the file share resource and select properties, click the parameters tab, and then click on permissions.

## NFS Cluster Specific Issues

In addition to the user name mapping best practices outlined in the “UNIX File System Management” chapter, there are additional recommendations.

For convenience, all suggestions are listed below:

- **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.
- **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.
- **In a clustered deployment, create user name mappings using domain user accounts.**  
Because the security identifiers of local accounts are recognized only by the local server, other nodes in the cluster will not be able to resolve those accounts during a failover. Do not create mappings using local user and group accounts.
- **In a clustered deployment, administer user name mapping on a computer that belongs to a trusted domain.**  
If NFS administration tasks are performed on a computer that belongs to a domain that is not trusted by the domain of the cluster, the changes are not properly replicated among the nodes in the cluster.
- **In a clustered deployment, if PCNFS password and group files are being used to provide user and group information, these files must be located on each node of the system.**  
Example: If the password and group files are located at *c:\maps* on node 1, then they must also be at *c:\maps* on node 2. The contents of the password and group files must be the same on both nodes as well.  
  
These password and group files on each server node must be updated periodically to maintain consistency and prevent users or groups from being inadvertently squashed.

## Persistent Storage Manager

Persistent Storage Manager objects, as any other file system resource, are dependent on the underlying physical disk resources and reside on the node that owns the resource. PSM management features therefore are node specific and can only operate against those PSM objects that exist on the node. As a consequence of this dependency, the following facts should be observed:

- A snapshot can only be created on the node that owns the disk resource.

- A scheduled persistent image will only succeed if the cluster node owns the disk resource(s) named in the scheduled image. If the disk resource has been moved or has failed over to the other cluster node, the scheduled snapshot will fail.
- Restoring from a persistent image will cause cluster physical disks and shares to be taken offline. Clients connected to these shares should be notified that access to the shares will not be available while the restore is taking place. After the restore completes, cluster physical disks are automatically brought online, however, the administrator must manually bring file shares back online.
- Cache File management can only be performed via the Web UI on the node that currently owns the underlying physical disk resource associated with the underlying file system.
- Only file share resources can be created for snapshot directories as NFS shares are not supported.

PSM snapshots behave in the exact same manner as any other file system object and can be shared out as a clustered file share resource via the Microsoft Cluster Administration interface. The standard pages in the Web UI for PSM function in a cluster in like manner as they do in stand alone configurations.

## Non Cluster Aware File Sharing Protocols

Services for Macintosh (SFM), File and Print Services for NetWare, HTTP file sharing protocols are not cluster aware and will experience service interruption if installed on a clustered resource during failover events of the resource. Service interruptions will be similar to those experienced during a server outage. Data that has not been saved to disk prior to the outage will experience data loss. In the case of SFM, it is not supported because SFM maintains state information in memory. Specifically, the Macintosh volume index is located in paged pool memory. Using SFM in clustered mode is not supported and may result in data loss similar in nature to a downed server should the resource it is based on fail over to the opposing node.

## Pathing Software in a Clustered Deployment

Cluster configurations should be deployed with dual data paths for high availability. Dual data paths from each node enable a path failure to occur that does not force the failover of the node. Clusters can be configured with single path, but if a failure in the path does occur, the entire node's resources will be failed to the non-affected node.

The NAS b3000 v2 and the NAS e7000 v2 both support AutoPath and Secure Path. Secure Path is designed for MSA1000, MA, and EVA storage subsystems. AutoPath is designed for use with VA and XP storage subsystems. Both software products are included with the NAS device, however, while Secure Path licensing is included licensing for AutoPath must be obtained separately from HP.

An online tool in the HP Utilities tab of the Web UI provides guidance for installing and configuring each of these products. It is strongly recommended that the SAN Connection Tool be utilized when configuring each of these products.



# Remote Access Methods and Monitoring

## 10

The HP StorageWorks NAS server 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 server
- Telnet Server
  - Enabling Telnet Server
  - Configuring Telnet Server
- Remote Shell Daemon
- HP Systems Management
  - HP Systems Management Console
  - HP Systems Management Agent Web Interface

These options let administrators use interfaces with which they are already familiar.

## Web Based User Interface

The NAS server 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 machine name or IP address>:3201/`

Extensive procedural online help is included in the WebUI.

## Terminal Services

The NAS server 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 server desktop
- The NAS Management Console
- A command line interface
- Backup software
- Antivirus programs
- Pathing software
- 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. The integrated Lights-Out port is 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.



## Features

The integrated Lights-Out port provides the following features:

---

**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.

---

- Hardware based graphical remote console access
- Remote restart
- Server failure alerting
- Integration with HP Systems Management
- 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 server 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 to use the **F8** feature.

---

- 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 Server

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.

---

**Note:** The iLO port can also be accessed from the HP Utilities tab of the WebUI by clicking the remote management link.

---

2. At the Integrated Lights-Out Account Login screen, supply the username and password for the iLO and click **Login**.
3. Click the Remote Console tab. The Remote Console Information screen is displayed.
4. Click on the Remote Console choice in the menu on the left side of the screen.
5. Press the **Ctrl-Alt-Del** button to login to the console.
6. Supply an administrator username and password. The NAS server desktop is displayed.

---

**Note:** The remote desktop feature of the iLO port requires a license key. The key is included with the product inside the Country Kit. See the iLO Advanced License Pack for activation instructions.

---

## 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 server, but must be activated before use.

---

**Note:** Telnet Server is non-clusterable.

---



**Caution:** For security reasons, the Telnet Server service must be restarted each time the server is restarted.

---

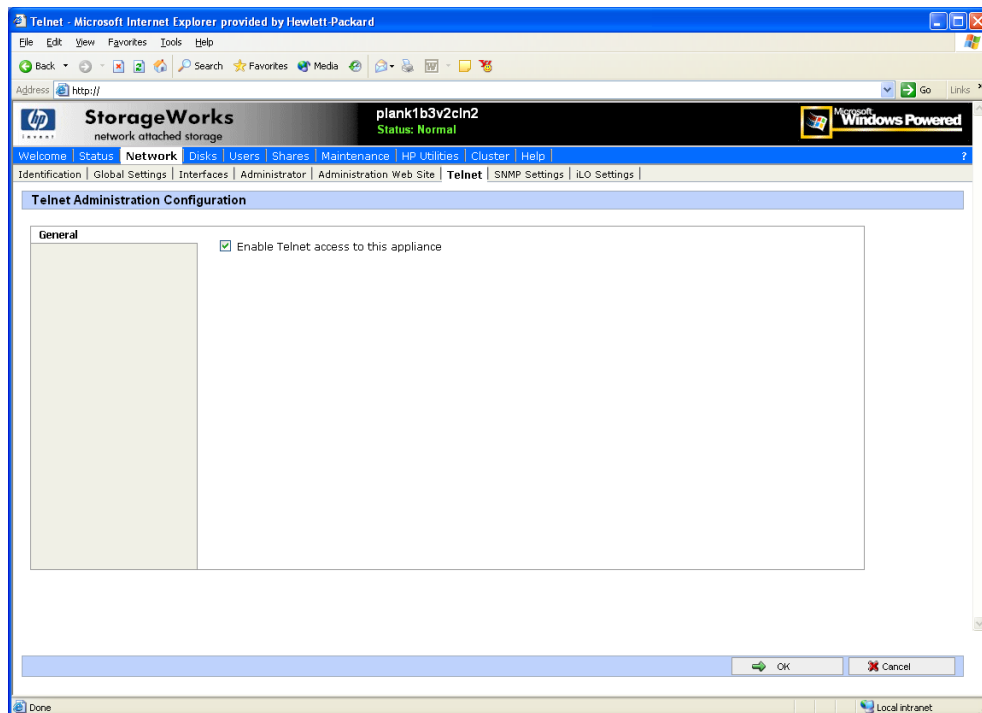
## Enabling Telnet Server

Telnet Server can be enabled in two ways. The first is to use Terminal Services to access a command line interface and enter the following command:

```
net start tlntsvr
```

The second is to open the WebUI:

1. Click **Network**.
2. Click **Telnet**.
3. Check the **Enable Telnet access to this appliance** box.
4. Click **OK**.



**Figure 117: Enabling Telnet access**

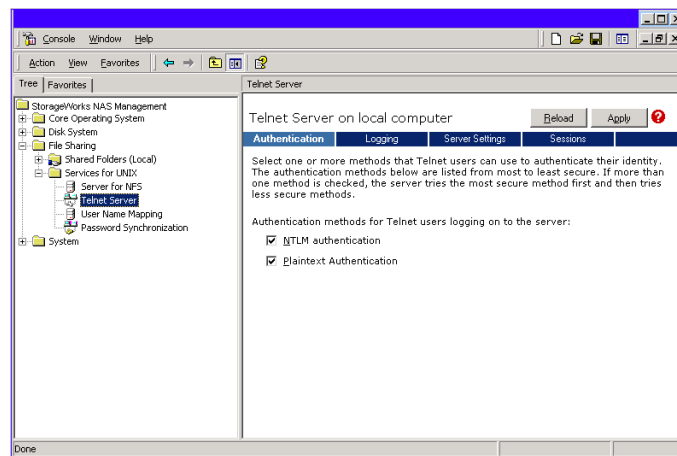
## Configuring Telnet Server

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 118: 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 server 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 server. The administrator will need to start this service by entering the following command:

```
net start rshsvc
```

---

**Note:** For security reasons, each time the server 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.

## HP Systems Management Version 7


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

HP Systems Management is a comprehensive management tool that monitors and controls the operation of HP servers and clients. HP Systems Management Version 7.0 or later is needed to successfully manage the NAS server. HP Systems Management consists of two components:

- Windows-based console application available on the Insight Manager 7 CD-ROM loaded on a separate client for NAS devices
- 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.

# PSM Error Codes



If problems are experienced when 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 9: PSM Error Codes**

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

**Table 9: PSM Error Codes**

| Error Code | Description   |
|------------|---|
| 0x00000057 | An invalid parameter was passed to a function.<br>Action: Programmatically, verify the parameters being passed to PSM are correct.  |
| 0x00000079 | I/O timed out while reading from the cache file.<br>Action: Verify the hard drive is operational.   |
| 0x0000007A | Buffer size supplied is insufficient to hold requested information.<br>Action: Save the system eventlog and contact technical support.  |
| 0x000000A1 | An invalid path was given for the cache file.<br>Action: Save the system eventlog and contact technical support.  |
| 0x000000EA | Buffer size supplied is insufficient to hold requested information.<br>Action: Save the system eventlog and contact technical support.  |
| 0x000003E6 | An exception occurred.<br>Action: Save the system eventlog and contact technical support.   |
| 0x00000456 | PSM was stopped because the media of a device being PSM'ed was changed.<br>Action: You can take a new snapshot now  |
| 0x0000045D | An error occurred on the device.<br>Action: Save the system eventlog and contact technical support.   |
| 0x000005AA | There is insufficient memory available.<br>Action: Close unnecessary applications or add more memory.   |
| 0x000006F8 | Buffer size supplied is insufficient to hold requested information.<br>Action: Save the system eventlog and contact technical support.  |
| 0x000006F8 | Invalid buffer address passed for I/O.<br>Action: Save the system eventlog and contact technical support.   |
| 0x80000005 | Specified buffer size is too low.<br>Action: Save the system eventlog and contact technical support.  |
| 0x8000001C | PSM was stopped because the media of a device being PSM'ed was changed.<br>Action: Take a new snapshot.   |
| 0xA0000004 | The cache file is <x>% full. The oldest snapshot(s) will automatically be deleted at <y>%.<br>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).)<br>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.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC0000002 | Function is not yet implemented.<br>Action: Save the system eventlog and contact technical support.   |



**Table 9: PSM Error Codes**

| Error Code | Description  |
|------------|--|
| 0xC0000005 | An Access Exception occurred.<br>Action: Save the system eventlog and contact your vendor's technical support.   |
| 0xC0000008 | User performing PSM function without opening PSM.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC000000D | An invalid parameter was passed to a function.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC000000E | Device name is not recognized by PSM.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC0000010 | An invalid IOCTL was sent to the driver.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC0000013 | Access to a virtual volume has been attempted after it has been destroyed.<br>Action: Do not access virtual volumes after they have been destroyed.  |
| 0xC000001C | An invalid IOCTL was sent to the driver.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC0000022 | An access exception occurred.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC0000022 | You do not have sufficient rights to the cache file directory.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC0000023 | Specified buffer size is too small.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC0000034 | Cache file name is invalid.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC000003A | An invalid path was given for the cache file.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC000003B | An invalid path was given for the cache file.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC000003E | Bad sector was detected in the cache file.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC0000043 | A file cannot be opened because the share access flags are incompatible.<br>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.<br>Action: Save the system eventlog and contact technical support.   |
| 0xC00000B5 | I/O timed out while reading from the cache file.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC00000BA | The cache location must be a file rather than a directory.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC00000E8 | Invalid buffer address passed for I/O.<br>Action: Save the system eventlog and contact technical support.  |

**Table 9: PSM Error Codes**

| <b>Error Code</b> | <b>Description</b>   |
|-------------------|--|
| 0xC000010A        | PSM was told to shut down.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC0000184        | Something has gone wrong with PSM.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC0000185        | An error occurred on the device.<br>Action: Save the system eventlog and contact technical support.  |
| 0xC0000206        | Buffer size supplied is insufficient to hold requested information.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE0001001        | PSM could not start due to the server being constantly busy for minutes.<br>Action: Take a snapshot when the system demands are lower.   |
| 0xE0001002        | PSM detected a deadlock.<br>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.<br>Action: Do not delete volumes with active snapshots.  |
| 0xE0001004        | PSM was specified for a volume that is currently not being PSM'ed.<br>Action: Save the system eventlog and contact technical support.  |
| 0xE0001005        | Cache file overflow caused all existing snapshots to be deleted.<br>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.<br>Action: Programmatically, a program must register with PSM prior to sending it commands.  |
| 0xE0001007        | Invalid license code.<br>Action: Contact vendor for a valid license.   |
| 0xE0001008        | Another application already has PSMed locked exclusively.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE0001009        | PSM needs to be locked exclusive for this function to work.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE000100A        | Wrong version of the driver has been loaded on this system.<br>Action: Verify the PSM version, save the system eventlog and contact technical support.   |
| 0xE000100B        | A reboot is required before PSM can operate.<br>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.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE000100D        | An incompatible DLL from another version of PSM is already loaded.<br>Action: Verify the PSM version, save the system eventlog and contact technical support.  |
| 0xE000100E        | Out of memory.<br>Action: Close unnecessary applications or add more memory.   |

**Table 9: PSM Error Codes**

| <b>Error Code</b> | <b>Description</b>   |
|-------------------|--|
| 0xE000100F        | Invalid parameter.<br>Action: Save the system eventlog and contact technical support.  |
| 0xE0001010        | Invalid handle.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE0001011        | Not implemented yet.<br>Action: Save the system eventlog and contact technical support.  |
| 0xE0001012        | Object type is not expected object.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE0001013        | User buffer is not large enough.<br>Action: Save the system eventlog and contact technical support.  |
| 0xE0001014        | Out of available structures.<br>Action: In Windows 2000 for NAS (Disks/Persistent Storage Manager), delete some snapshots.   |
| 0xE0001015        | PSM is shutting down.<br>Action: This is not an error but is a status message.   |
| 0xE0001016        | The device, volume or object does not exist.<br>Action: Verify that the device, volume, or object exists.  |
| 0xE0001017        | Unsuccessful.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE0001018        | The device does not have any media loaded.<br>Action: If the snapshot has been deleted, it cannot be accessed  |
| 0xE0001019        | Object already exists.<br>Action: Save the system eventlog and contact technical support.  |
| 0xE000101A        | Specified path is a directory and not a file.<br>Action: Provide a full path and filename  |
| 0xE000101B        | Invalid path was specified.<br>Action: Ensure the CacheFile name is correct  |
| 0xE000101C        | The static volume was not mounted.<br>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.<br>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.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE000101F        | The volume the cache file resides on is out of space.<br>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.<br>Action: The cache file for each volume resides on the volume itself. Do not dismount the volume.   |

**Table 9: PSM Error Codes**

| Error Code | Description   |
|------------|---|
| 0xE0001021 | The server was shutdown.<br>Action: Do not shut down the machine while snapshots are in progress.   |
| 0xE0001022 | Unable to create cache file.<br>Action: Save the system eventlog and contact support.   |
| 0xE0001023 | PSM recovery could not find a snapshot entry.<br>Explanation: A snapshot was lost during the recovery process. It is unknown which snapshot it was.<br>Action: Save the system eventlog and contact technical support.  |
| 0xE0001024 | PSM recovery could not open the index file.<br>Explanation: All snapshots are corrupt.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE0001025 | PSM recovery encountered error <x> inserting key (<y>:<z>) into dictionary.<br>Explanation: <x> is the error that occurred and can be found in this list of errors.<br>Action: Look up the error in this list and take the specified action.  |
| 0xE0001026 | PSM recovery encountered corrupt index sector %2.<br>Explanation: An index entry was found to be corrupt during the last boot.<br>Action: Save the system eventlog and contact technical support.   |
| 0xE0001027 | A snapshot could not be created due to error 0x<x>.<br>Explanation: <x> is the error that occurred.<br>Action: Look up the error in this list and take the specified action.  |
| 0xE0001028 | The cache file is <x>% full. Snapshots have been deleted.<br>Explanation: The oldest snapshots have been deleted.<br>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.<br>Explanation: PSM cannot create any more snapshots because the configured maximum number of snapshots that PSM can keep concurrently has been reached.<br>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.<br>Action: Contact your vendor's technical support for a non-evaluation version.   |
| 0xE000102B | There is not enough free cache space to perform the operation.<br>Action: Delete some snapshots to free up some cache space or 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.<br>Action: Increase the maximum snapshot number. This is a status message.   |

**Table 9: PSM Error Codes**

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



**A**

access control list. See ACL  
 access rights, managing 199  
 ACL  
   defined 118  
   translating 145  
 adaptive load balancing 32  
 ADG (Advanced Data Guarding) 57  
 AFP  
   installing services for AppleTalk 126  
   installing services for Macintosh 126  
   protocol settings 128  
   shares, setting up 127  
   sharing  
     volume mount points 87  
 alerts, e-mail, setting up 46  
 AppleTalk. See AFP.  
 array controller  
   purpose 52  
 arrays  
   defined 51  
 audience 12  
 Authentication software, installing 133  
 authorized reseller, HP 16

**B**

backup  
   mappings 150  
   snapshots 66  
 basic disk 59

**C**

cache  
   changing size 75  
   clearing from system 70  
   deleting images 75  
   full 87  
   size 74  
   usage 74

caching  
   automatic for documents 123  
   automatic for programs 123  
   manual for documents 123  
 CIFS  
   add a new user or group 123  
   modify access 123  
   protocol settings 128  
   remove access 123  
   share support 119  
   sharing 122  
 CIFS administration 90  
 Cisco Fast EtherChannel 32  
 client groups  
   adding NFS 142  
   deleting NFS 142  
   editing NFS 143  
   managing NFS 141  
 cluster  
   components, hierarchy 179  
   concepts 177  
   concepts, diagram 178  
   diagram 175  
   disks, configuring 189  
   groups, node-based 198  
   installation 184  
   load balancing 198  
   managing access rights 199  
   managing file share permissions 199  
   NFS issues 200  
   node  
     configuring first 187  
     configuring second 192  
   nodes  
     powering down 196  
     powering up 197  
     restarting 195  
   overview 173  
   pathing software 201  
   planning 180

- protocols, non cluster aware 201
- resource overview 198
- resources, defined 176
- setting up user account 186
- terms and components 176
- validating installation 192
- verify installation 193
- configuring shared disks 186
- connectivity, verifying 185
- conventions
  - document 13
  - equipment symbols 13
  - text symbols 13

## D

- data blocks 52
- data guarding explained 56
- data recovery, PSM 65
- data replication software
  - installing 39
- data striping 52, 54
- date, system, changing 42
- disk access, verifying 186
- disk quotas
  - creating new entries 101
  - deleting entries 102
  - disabling 100
  - enabling 100
  - modifying entries 102
- document
  - conventions 13
  - prerequisites 12
  - related documentation 13
- domain controller
  - configuring 90
- domain environment 21
- domain membership, verifying 185
- drive defragmentation 65
- drive mirroring explained 55
- drive quotas
  - defined 98
  - managing 98
- dynamic disk 59

## E

- e-mail alerts, setting up 46
- encoding types 137
- environments
  - domain compared to workgroup 90
  - overview 21
- equipment symbols 13
- error codes 211
- Ethernet NIC teams

- adding 30
- checking status 36
- configuring 31
- configuring properties 33
- configuring TCP/IP 34
- renaming the connection 33
- setting up 28
- showing connection icon 34
- troubleshooting 37
- events, SFU, logging 132
- explicit mapping 148
- explicit mappings 144

## F

- fail on fault setting 31
- failover
  - automatic 195
  - defined 176
- fault tolerance
  - for NIC teams 31
  - methods supported 54
- features
  - hardware 18
  - redundancy 19
  - software 18
- FEC 32
- File and Print Services for NetWare. See FPNW.
- file level permissions 112
- file share permissions, managing 199
- files, ownership 117
- folders
  - auditing access 115
  - compress tab 109
  - creating new 108
  - creating new share 110
  - deleting 109
  - general tab 108
  - managing 106
  - managing shares for 111
  - modifying properties 109
  - navigating to 107
- FPNW
  - accessing 161
  - described 159
  - installing 160
- FTP
  - protocol settings 128
  - sharing 124

## G

- getting help 15
- granule size
  - rules 69



- update utility 69
- group names
  - examples 91
  - managing 91
- groups
  - adding from a domain 98
  - adding local users 97
  - adding to permissions list 113
  - local, adding 96
  - local, deleting 96
  - local, managing 95
  - local, modifying properties 97
  - properties, general tab 97
  - properties, members tab 97
  - removing local users 98

**H**

- hard drives
  - best practices 59
  - online spares 54
  - physical 51
  - RAID 19
- help, obtaining 15
- HP
  - authorized reseller 16
  - storage website 16
  - technical support 15
- HP Network Teaming Utility
  - installing 29
  - opening 30
- HTTP protocol settings 128

**I**

- iLO. See Integrated Lights-Out Port
- image directory 73
- image groups 79
- inactive period 72
- inactive time-out 73
- Insight Manager
  - defined 19
  - described 210
- Integrated Lights-Out port
  - accessing NAS B2000 207
  - activating 41
  - configuration 206
  - described 204
  - features 205
  - license key 41

**L**

- LAN icons, renaming 185
- LDM (Logical Disk Manager) 59, 61
- license key, iLO port 41

- load balancing 32, 198
- localhost 132
- locks, NFS 139
- logging, SFU events 132
- logical drives. See LUNs
- logical storage elements 59
- logs
  - accessing 44
  - audit 44
  - options 44
- LUNs
  - and storage controller subsystems 53
  - expanding 53
  - largest size 53
  - maximum number 53

**M**

- Macintosh, installing services for 126
- management, storage 49
- managing system storage 38
- mapping
  - best practices 145
  - data stored 146
  - explicit 148
  - simple 147
- mappings
  - backup and restore 150
  - creating 146
  - explicit 144
  - NFS 144
  - simple 144
  - squashed 145
- Microsoft Services for UNIX. See SFU
- mount points 87
- mount points, creating 59

**N**

- NAS B2000
  - defined 18
  - desktop 24
  - hardware features 18
  - restarting 43
  - shutting down 43
  - software features 18
  - supported fault tolerance methods 54
  - using iLO to access 207
  - utilities 19
- NAS Data Copy
  - described 39
  - installing 40
- NAS Management Console 25
- NCP
  - adding new user 125

- creating new share 164, 168
  - properties, modifying 166
  - protocol settings 128
  - removing access 125
  - shares, deleting 165
  - shares, modifying properties 171
  - sharing 125
- NetWare
- adding local users 162
  - enabling user accounts 163
  - installing services for 160
  - supervisor account 163
  - volume mount points 87
- Network File System. See NFS.
- network planning 181
- network settings, changing 46
- networks
- setting up 184
- NFS
- async/sync settings 139
  - client groups 141
    - adding 142
    - deleting 142
    - editing 143
  - cluster specific issues 200
  - compatibility issues 119
  - deleting shares 135
  - file share, creating 134
  - file sharing tests 152
  - group mappings 144
  - locks 139
  - modifying share properties 135
  - network protocols 130
  - protocol properties settings 138
  - protocol settings 128
  - sharing 124
  - user mapping server 131
  - user mappings 144
- NIC teams. See Ethernet NIC teams
- node
- defined 176
- NTFS partition size limit 53
- O**
- online spares 54
- P**
- partitions 59
- passwords 154
- default, WebUI 22
  - modifying local user's 94
  - synchronization
    - advanced settings 156
    - best practices 155
    - customizing 157
    - implementing 155
    - installing 156
    - requirements 155
- pathing software
- in a cluster 201
- performance, snapshots 68
- permissions
- file level 112
  - list
    - adding users and groups 113
    - removing users and groups 113
  - modifying 113
  - resetting 115
- persistent image. See snapshots.
- Persistent Storage Manager
- in a cluster 200
- Persistent Storage Manager. See PSM
- physical disk resources, configuring 194
- physical storage best practices 59
- planning
- cluster 180
  - network 181
  - protocol 182
  - storage 180
- prerequisites 12
- properties
- editing PSM schedule 77
  - editing snapshot 83
- protocols
- AFP settings 128
  - CIFS settings 128
  - FTP settings 128
  - HTTP settings 128
  - NCP settings 128
  - NFS properties settings 138
  - NFS settings 128
  - non cluster aware 201
  - parameter settings 127
  - planning 182
  - planning for compatibility 119
  - supported 21, 127
- PSM
- accessing 71
  - and creating mount points
  - creating new schedule 76
  - creating snapshots 81
  - data recovery 65
  - deleting schedule 78
  - deleting snapshots 82
  - editing schedule properties 77
  - editing snapshot properties 83
  - elements overview 61

- error codes [211](#)
- global settings [72](#)
- image directory [73](#)
- image groups [79](#)
- inactive period [72](#)
- inactive time-out [73](#)
- managing snapshots [80](#)
- overview [63](#)
- restore defaults [73](#)
- restore snapshot [85](#)
- schedules [75](#)
- storage limitations [71](#)
- undo snapshot changes [84](#)
- volume configuration settings [74](#)
- volume display [70](#)
- volume settings [73](#)
- public network adapter, configuring [185](#)

## Q

- quorum disk
  - defined [177](#)
  - recommendations [186](#)

## R

- rack stability, warning [15](#)
- RAID
  - ADG advantages [58](#)
  - ADG disadvantages [58](#)
  - ADG explained [57](#)
  - level on server [19](#)
  - RAID 0 [52](#)
  - RAID 0 advantages [54](#)
  - RAID 0 disadvantages [54](#)
  - RAID 0 explained [54](#)
  - RAID 1 advantages [55](#)
  - RAID 1 disadvantages [56](#)
  - RAID 1 explained [55](#)
  - RAID 1+0 explained [55](#)
  - RAID 5 advantages [57](#)
  - RAID 5 disadvantages [57](#)
  - RAID 5 explained [56](#)
  - summary of methods [58](#)
- rapid startup utility
  - defined [19](#)
- redundancy [19](#)
- related documentation [13](#)
- remote access
  - iLO port [204](#)
  - Insight Manager [210](#)
  - methods listed [203](#)
  - remote shell daemon [210](#)
  - Telnet Server [208](#)
  - Terminal Services [204](#)

- WebUI [204](#)
- remote shell [210](#)
- Remote Shell Service [153](#)
- resources, cluster [176](#)
- restarting the server [43](#)
- restore
  - PSM defaults [73](#)
  - snapshot [85](#)

## S

- scheduled shutdown [43](#)
- schedules
  - creating new PSM [76](#)
  - deleting PSM [78](#)
  - editing PSM properties [77](#)
  - PSM [75](#)
- security
  - auditing [115](#)
  - file level permissions [112](#)
  - ownership of files [117](#)
- services for AppleTalk, installing [126](#)
- services for Macintosh, installing [126](#)
- setup
  - completing [28](#)
  - e-mail alerts [46](#)
  - Ethernet NIC teams [28](#)
- SFU
  - commands [154](#)
  - event logging [132](#)
- SFU, described [131](#)
- shared disk requirements [183](#)
- shared disks, configuring [186](#)
- shares
  - administrative [119](#)
  - AFP
  - CIFS tab [122](#)
  - creating new [110](#), [120](#)
  - creating new NCP [164](#), [168](#)
  - deleting [121](#)
  - deleting NCP [165](#)
  - FTP [124](#)
  - managing [118](#)
  - managing for a volume or folder [111](#)
  - modifying NCP properties [166](#), [171](#)
  - modifying NFS properties [135](#)
  - modifying properties [122](#)
  - NCP [125](#), [164](#)
  - NFS [124](#)
  - NFS tests [152](#)
  - NFS, creating [134](#)
  - NFS, deleting [135](#)
  - path [111](#)
  - setting up AppleTalk [127](#)
  - standard [119](#)

- web (HTTP) 124
- shutting down the server 43
- simple mapping 147
- simple mappings 144
- smart switch 31
- snapshots
  - always keep 65
  - attributes 64
  - automated deletion 65
  - backup 66
  - creating 64, 81
  - deleting 82
  - drive defragmentation 65
  - editing properties 83
  - facts 61
  - managing 80
  - maximum number 72
  - performance impact 68
  - read/write 65
  - reading 64
  - read-only 64
  - recovering 68
  - restoring 85
  - undo changes 84
  - use with Veritas Backup Exec 67
- software
  - data replication 39
  - installing Authentication 133
  - updating 46
- squashed mappings 145
- squashing 131
- storage controller subsystems and LUNs 53
- storage management elements 49
- subfolder, navigating to 107
- symbols in text 13
- symbols on equipment 13
- synchronization 154
- system date, changing 42
- system storage
  - managing 38
- system time, changing 42

## T

- TCP/IP, configuring on NIC team 34
- technical support, HP 15
- Telnet Server
  - auditing log 209
  - authentication tab 209
  - configuring 209
  - enabling 208
  - sessions information 209
  - settings 209
  - using 153
- Terminal Services

- defined 45
- described 204
- exiting 45
- opening 45
- using 153
- text symbols 13
- time, system, changing 42
- troubleshooting
  - PSM error codes 211
  - PSM known issues
    - PSM known issues 87

## U

- UNIX
  - authenticating user access 131
  - converting ACL 145
  - group ID 131
  - remote shell 210
  - See also NFS
  - Telnet Server 153
  - user ID 131
  - volume mount points 87
- user access, authenticating 131
- user account, setting up 186
- user interfaces 22
- users
  - adding to permission list 113
  - local
    - adding 93
    - deleting 93
    - managing 92
    - modifying properties 94
  - names, managing 91
  - NetWare
    - adding 162
    - enabling 163

## V

- validating cluster installation 192
- verifying
  - cluster installation 193
  - connectivity 185
  - disk access 186
  - domain membership 185
  - name resolution 185
- Veritas Backup Exec, use with snapshots 67
- virtual server, defined 176
- virtual storage 49
- volume settings, PSM 73
- volumes
  - available for snapshots 73
  - creating new share 110

---

- creating Novell [159](#)
- managing shares for [111](#)
- mount points [87](#)
- navigating to [107](#)
- NCP [164](#)
- planning [60](#)
- PSM configuration settings [74](#)
- re-extending [70](#)

## W

warning

- rack stability [15](#)
- symbols on equipment [13](#)
- web sharing [124](#)
- websites
  - HP storage [16](#)
- WebUI
  - accessing [22](#)
  - default password [22](#)
  - defined [19](#)
  - launching [204](#)
- workgroup environment [21](#)

