TotalStorage™ NAS Gateway 300
Model G27

User's Reference

GA27-4321-00

Note

Before using this information and the product it supports, be sure to read the general information in Appendix A, “Notices” on page 115.

First Edition (October 2002)

This edition applies the IBM 5196 TotalStorage NAS Gateway 300, (Model G27, product number 5196-G27) and to all subsequent releases and modifications until otherwise indicated in new editions.

Order publications through your IBM representative or the IBM branch office servicing your locality. Publications are not stocked at the address below.

IBM welcomes your comments. A form for reader’s comments is provided at the back of this publication. If the form has been removed, you can address your comments to:

International Business Machines Corporation
Design & Information Development
Department CGFA
PO Box 12195
Research Triangle Park, NC 27709–9990
U.S.A.

You can also submit comments on the Web at www.ibm.com/storage/support.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Figures

Tables

About this book

This book provides information necessary to configure and administer the IBM 5195 TotalStorage NAS Gateway 300, hereafter referred to as the NAS Gateway 300.

Who should read this book

This book is for NAS Gateway 300 administrators.

The NAS Gateway 300 administrator should have experience in at least the following skills, or have access to personnel with experience in these skills:

- Microsoft® Windows® and Windows Advanced Server
- Networking and network management
- Disk management
- SAN management
- General technologies of the product (such as Microsoft Cluster Service, Services for UNIX®, storage, RAID, and so on)
- Critical business issues (such as backup, disaster recovery, security)

Frequently used terms

This document contains certain notices that relate to a specific topic. The caution and danger notices also appear in the multilingual Safety Information on the Documentation CD-ROM that came with the appliance. Each notice is numbered for easy reference to the corresponding notices in the Safety Information.

The following terms, used within this document or within the Safety Information, have these specific meanings:

Term	Definition in this document
Notes	These notices provide important tips, guidance, or advice.
Attention	These notices indicate possible damage to programs, devices, or data. An attention notice is placed just before the instruction or situation in which damage could occur.
Caution	These notices indicate situations that can be potentially hazardous to you. A caution notice is placed just before descriptions of potentially hazardous procedure steps or situations.
Danger	These notices indicate situations that can be potentially lethal or extremely hazardous to you. A danger notice is placed just before descriptions of potentially lethal or extremely hazardous procedure steps or situations.

Publications

The latest versions of the following product publications are available in softcopy at:

www.ibm.com/storage/support/nas

Descriptions of the NAS Gateway 300 publications

The NAS Gateway 300 library consists of the following publications:

- *Hardware Installation Guide* GA27-4320

This book describes hardware physical specifications, electrical specifications, cabling, environmental specifications, and networking specifications for installing the NAS Gateway 300.

- *User's Reference* GA27-4321

  This book describes such operational and administrative activities as:

  – Using the configuration utilities
  – Administering the NAS Gateway 300
  – Troubleshooting
  – Using the Recovery and Supplementary CD-ROMs

Hardcopy publications shipped with the NAS Gateway 300

The following publications are shipped in hardcopy and are also provided in softcopy (PDF) form at:

www.ibm.com/storage/support/nas

- *NAS Gateway 300 Hardware Installation Guide* GA27-4320
- *Release Notes*

  This document provides any changes that were not available at the time this book was produced.

Note that the *User's Reference* is provided in softcopy only.

Related publications

The following publications contain additional information about the NAS Gateway 300:

- *NAS Gateway 300 Hardware Installation Guide* GA27–4320
- *NAS Gateway 300 Service Guide* GY27–0414
- *NAS Gateway 300, NAS 200, and NAS 100 Planning Guide* GA27–4319
- *UM Services User's Guide* (on the Documentation CD-ROM that came with the appliance)

Additional information on Universal Manageability Services, IBM Director, and Advanced System Management is located on the Documentation CD-ROM that came with the appliance.

Accessibility

The softcopy version of this manual and other related publications are accessibility-enabled for the IBM Home Page Reader.

Web sites

The following Web site has additional and up-to-date information about the NAS Gateway 300:

www.ibm.com/storage/nas/

A highly recommended Web site: for the latest troubleshooting guidance and symptom-fix tip information, go to the IBM support Web site at:

www.ibm.com/storage/support/nas

This site contains additional information, gathered from field experience, not available when this document was developed.

Chapter 1. Introduction

The NAS Gateway 300 connects clients and servers on an IP network to Fibre Channel storage, efficiently bridging the gap between LAN storage needs and SAN storage capacities.

This appliance offers a storage solution for both Windows, UNIX®, and UNIX-like environments, including mixed Windows-UNIX environments that enable Windows and UNIX clients and servers to share the same Fibre Channel storage.

Model G27 replaces Models G01 and G26. Enhancements provided by the new model include:

- More options in configuring Ethernet connections
- More options in configuring Fibre Channel connections
- More options for tape backup
- Faster processor
- Gigabit Ethernet connection
- Faster adapters

The dual-node Model G27 features:

- Two engines (IBM 5187 NAS Model 7RY), each with:
  - Dual 2.4-GHz processors
  - 512 MB of ECC memory standard (plus one upgrade); up to 4.5 GB available
  - Two redundant hot-swap 270 watt power supplies
  - Qlogic 2340 1-port Fibre Channel adapter for storage area network (SAN) connection
  - Four PCI adapter slots for plugging in optional adapters, including three high-performance slots. (Communication between the two engines takes place through an integrated 10/100/1000 Mbps Ethernet port on each engine's planar board.)
- Optional adapters:
  - Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter
  - IBM Gigabit Ethernet SX Server Adapter
  - IBM PCI Ultra160 SCSI adapter (LVD/SE)
  - Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter
  - PRO/1000 XT Server Adapter by Intel
  - Qlogic 2340 1-port or Qlogic 2342 2-port Fibre Channel adapter (to replace single-port Fibre Channel SAN adapter)
  - Qlogic 2340 1-port Fibre Channel adapter for tape backup
  - Remote Supervisor Adapter

In addition, the Model G27 provides clustering and failover protection. This high-availability design helps protect against appliance failure to provide continuous access to data.

Note: Throughout this book, information about the Model G27 node and engine applies to both its nodes and engines.

The preloaded software stack is based on the Windows Powered OS operating system, which is very similar to Microsoft® Windows® 2000 Advanced Server. Preloaded software includes:

Microsoft Windows 2000 for Network Attached Storage
Enables remote administration of the appliance using a Web-based graphical user interface (GUI).

Microsoft Windows Terminal Services
Enables remote administration of the appliance using its Windows desktop.

Microsoft Cluster Service
Provides clustering support and failover protection.

Microsoft Services for UNIX
Provides file access to UNIX and UNIX-based clients and servers through the Network File System (NFS) protocol. Note that the NAS Gateway 300 supports Linux and other platforms that employ NFS.

IBM Director Agent and Universal Manageability Server Extensions
Provides system management support based on industry standards (in conjunction with the IBM Director console application as well as other management software).

IBM Advanced Appliance Configuration Utility agent
Supports management through the IBM Advanced Appliance Configuration Utility console application (supports aggregate Web-based management of all of your IBM appliances).

IBM FAStT Management Suite Java (MSJ)
Provides diagnostics for the Fibre Channel adapters.

Intel® PROSet II
Provides diagnostics for the Intel Ethernet adapters.

Alacritech® SLICuser
Provides diagnostics for the quad-port and accelerated Ethernet adapters.

Columbia Data Products® Persistent Storage Manager (PSM)
Provides 250 persistent images of customer data and enables full online backup of system with Microsoft backup applications.

Tivoli® Storage Manager Client
Provides data backup and archive support (in conjunction with Tivoli Storage Manager Server).

Tivoli SANergy
Provides shared data access to the SAN storage at Fibre Channel speed.

Services for NetWare
Provides interoperability within the Novell environment and a complete set of new interoperability services and tools for integrating the NAS Gateway 300 into existing NetWare environments. Only Netware V5.0 Print and File services are included in the preloaded code and is required for supporting Netware File system protocol. Clustering is not supported by the SFN5.

Storage Manager for SAK
A storage management tool that includes storage reports, directory quotas, and file screening functions.

Roadmap for setting up and configuring the NAS Gateway 300

A suggestion for first-time users . . .

Your understanding of the NAS Gateway 300 and your ability to use it will be greatly enhanced if you first proceed to the NAS Setup Navigator tutorial.

The NAS Setup Navigator maps out the initial configuration tasks and leads you through the tasks in the proper order. The tool detects which NAS appliance it is running on and adjusts the menu and content appropriately. You can follow links to more in-depth information and to the configuration panels used to perform the steps. You can also tailor the instructions to fit your needs by selecting optional topics. The Navigator not only presents information on functions and features–such as clustering–but also allows you to enable the functions and features. To start the NAS Setup Navigator, click on the NAS Setup Navigator icon on the desktop.

After you have become familiar with the NAS Gateway 300, you can refer to this book for more details.

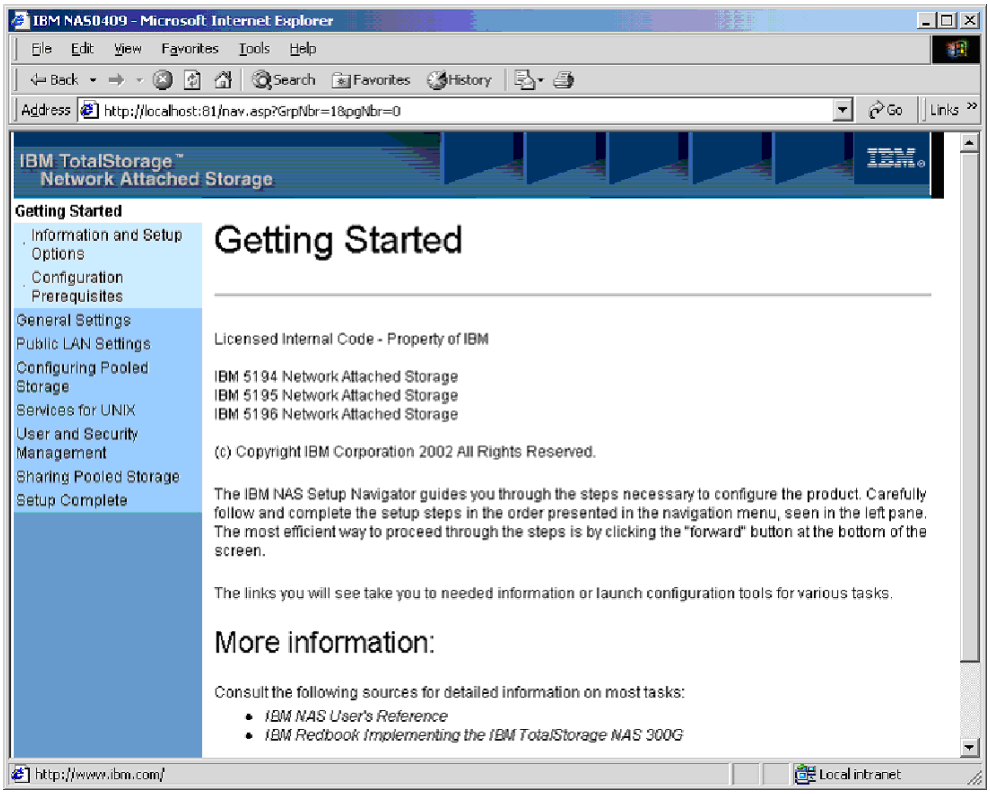

Figure 1. Opening screen of the NAS Setup Navigator

The following roadmap presents the requirements and instructions for setting up and configuring the NAS Gateway 300. Following these directions and referring to the appropriate sections of this book will help you in this task.

Prerequisites

- A domain controller must exist on the network and a login ID must be defined for each node to log on. Each node must join the same domain.

- All Windows shared disks must be defined as basic. Windows 2000 dynamic disks are not supported.
- A Quorum drive must be available to both nodes and have the same drive letter on each node.
- All disks shared between the two cluster nodes must have the same drive letter.
- All shared storage must be defined as NTFS and be on primary partitions.
- Compression cannot be enabled on any disk partition.
- Each node must have one private and one public adapter.

Cluster setup requirements

See "Cluster setup requirements" on page 6.

Configuration and administration tools

The NAS Gateway 300 is a network-attached storage appliance that has several different methods of configuration depending on your environment.

First, determine how you will manage the device. You can manage the NAS Gateway 300 in "headless" mode or with a keyboard, display, and mouse directly attached to each node. See "Using a keyboard, monitor, and mouse for setup and configuration" on page 13 for information on managing this device using a keyboard, display, and mouse. For "headless" management of the NAS Gateway 300, you can use one of the following tools:

- Terminal Services, for remote configuration and management from another device on the network
- Universal Manageability Services (UMS) for management through a Web browser
- Windows 2000 for NAS, a Web-based GUI for those not familiar with the Windows desktop
- IBM Advanced Appliance Configuration Utility (IAACU) for setup and configuring multiple devices or other appliances on a single network

After you determine how you will manage the NAS Gateway 300, you can begin setup and configuration of the device.

For more information on configuration and administration tools, see Chapter 3, "Configuration and administration tools" on page 13.

Step 1 - Initial network setup

Configure both nodes to enable access over the network. The general steps to do this are given below. More details are given in Chapter 2, "Getting started" on page 9.

1. Use Dynamic Host Configuration Protocol (DHCP) or static addressing to set up one public network connection in each node.
   a. If you are operating with a keyboard, display, and mouse, set up a public network connection to access the device.
   b. If you are operating in a headless environment, use one of the following methods:
      - If DHCP is installed and the IP address requested can be determined, you can use DHCP for initial setup, but you should change this address to static later in the configuration.

- If you have multiple appliances or cannot determine the DHCP address, you can install the IAACU utility to identify appliances and define IP addresses. The tool will also allow you to set static addresses.

2. Complete the steps in “Setting the date and time” on page 10 and “Setting up the network” on page 11.

Step 2 - Define storage and setup partitions

The NAS Gateway 300 attaches to your SAN-attached storage device, through the Fibre Channel, and provides your Ethernet LAN-attached clients access to that storage. You must define storage arrays and logical drives on the SAN-attached storage device and set up Windows partitions on the logical drives as defined in Chapter 4, “Setting up storage” on page 27.

For more information on defining storage and setting up partitions, see Chapter 4, “Setting up storage” on page 27.

Step 3 - Complete network setup and cluster installation

1. Power on either node. (This becomes the first node.)
2. Set up the first node:
   a. Networking setup

      See “Networking setup” on page 33. Note the cautionary statement at the beginning of that section.

   b. Domain setup

      See “Joining a node to a domain” on page 36.
3. Shut down the first node (see “Shutting down and powering on the NAS Gateway 300” on page 87 for more information on shutting down the NAS Gateway 300).
4. Power on the other node (the joining node).
5. Set up the joining node:
   a. Networking setup

      See “Networking setup” on page 33.

   b. Shared storage setup

      For the joining node, the only part of this step that you must complete is assigning drive letters on the shared storage; make sure that the drive letters are the same as those on the first node.

      Also, if you have trouble with the Fibre Channel connection, you can use the steps in “Fibre Channel adapter” on page 104 to diagnose the problem.

   c. Domain setup

      See “Joining a node to a domain” on page 36.

   d. Shut down the joining node (see “Shutting down and powering on the NAS Gateway 300” on page 87 for more information on shutting down the NAS Gateway 300).
6. Power on the first node and complete “Cluster setup” on page 39.
7. Power on the joining node and complete “Cluster setup” on page 39.

For more information on network setup and cluster installation, see Chapter 5, “Completing networking, clustering, and storage access setup” on page 33.

Step 4 - Cluster administration

At this point you can add users, file shares, and complete other configuration tasks to improve operations of the NAS Gateway 300 in a cluster environment.

1. Add users (see “Creating users” on page 42).
2. Add file shares (see “Creating clustered file shares (CIFS and NFS)” on page 50). Note that you must configure Server for NFS before NFS file sharing can be used.

For more information on cluster administration, see “Configuring clusters” on page 41.

Step 5 - Additional functions

Additional functions are available for backup, persistent images, and adding more storage areas. It is recommended that after you complete the setup and configuration procedures, you use the Persistent Storage Manager Disaster Recovery option (“Disaster Recovery” on page 66) or other method to back up the system configuration in the event of a failure.

Also, **it is imperative to use the system shutdown procedure** described in “Shutting down and powering on the NAS Gateway 300” on page 87 to ensure system integrity.

For more information, see Chapter 6, “Managing and protecting the network and storage” on page 55.

Cluster setup requirements

Before you configure the NAS Gateway 300 nodes for clustering, ensure that the following requirements are met:

Network requirements

- A unique NetBIOS cluster name.
- You will need ***at least*** seven static IP addresses: five for the node and cluster setup, and two for each file share served by the cluster. A formula for the number of static IP addresses is: 5 + (2 x *number_of_file shares*). The IP addresses required for node and cluster setup are:
  - At least three unique, static IP addresses for the public network: one for each node (for client access through the PCI NIC adapter) and one for the cluster itself (the administration IP address).

    Table 1 on page 8 shows a summary of the networking information necessary for the public connection.
  - Two static IP addresses for the cluster interconnect on a private network or crossover, through the onboard Ethernet adapter. The default IP addresses for the private network adapters are 10.1.1.1 for the first node in the cluster, and 10.1.1.2 for the node that joins the cluster. (The top node in the NAS Gateway 300 is considered the first node, and the bottom node is considered the joining node.)

Notes:

1. If you are not the system administrator, contact that person for the applicable IP addresses.
2. Each node in a cluster must join the same domain and be able to access a Primary Domain Controller (PDC) and DNS server, but it is not required that the nodes log into the domain.

3. Each node in the cluster must have at least two network adapters: at least one for the public network and the other for the private interconnect.

Shared disk requirements

- All shared disk arrays and devices, including the quorum disk, must be physically attached to a shared storage bus.
- All shared disks must be configured as **basic** (not dynamic) disks.
- All shared disks **must have the same drive letter** on each node.
- All partitions on these disks must be formatted with NTFS.
- All partitions on these disks must also be Primary Partitions.
- Compression must not be enabled.

Shutting down and powering on the NAS Gateway 300

The clustering function requires special considerations when you need to shut down and power on the NAS Gateway 300. See “Shutting down and powering on the NAS Gateway 300” on page 87 for details.

Table 1. Networking information worksheet for the public connection

Cluster component	Information needed
Cluster	**Cluster name:** **IP address:** **Subnet mask:**
First node	**Computer name (example: IBM5196–23H1234):** **IP address:** **Subnet mask:** **Gateway:** **Preferred DNS:** **WINS server (optional):**
Joining node	**Computer name:** **IP address:** **Subnet mask:** **Gateway:** **Preferred DNS:** **WINS server (optional):**
Domain to join	**Domain name:**

Chapter 2. Getting started

This chapter gives details to set up the initial communication to the NAS Gateway 300 to enable setup and configuration. These instructions refer specifically to a base configuration as shipped and do not cover the setup of additional storage units, which can be purchased separately.

Note: You must follow these procedures for both nodes.

Methods for setting up the NAS Gateway 300

The following sections detail how to set up the NAS Gateway 300. You must first ensure that the network recognizes the new appliance. Which method you should use depends on several conditions:

- In "headless" mode (without a keyboard, monitor, and mouse directly attached to the unit), use one of the following methods:
  - **IBM Advanced Appliance Configuration Utility**

    If you have multiple appliances or cannot determine the DHCP address, install the IAACU to identify appliances and define IP addresses. The tool also allows you to set static addresses.

    If you are using this method, proceed with "Installing the IBM Advanced Appliance Configuration Utility" on page 17.
  - **Windows Terminal Services**

    If DHCP is installed and the IP address requested can be determined, use this method for initial setup, but you should change the address to static later in the configuration. This condition is most appropriate when using Windows Terminal Services for operation of the NAS Gateway 300.

    If you are using this method, proceed with "Initial setup and configuration" on page 10.
- The use of a keyboard, display, and mouse is most appropriate when there is a single or few appliances in the network and you use static setup and definition.

  If you are using this method, proceed with "Initial setup and configuration" on page 10.

Accessing Universal Manageability Services

1. You will be prompted to authenticate with the administrative user name ("Administrator") and password (initially "password," but you can change it later; note that the password is case-sensitive, but the user name is not).

   If this is the first time you have accessed the UM Services browser (on any appliance) from this workstation, you will also be prompted to install the Swing and XML Java libraries in your Web browser. You can download these libraries from the NAS Gateway 300 through the network link.
2. The UM Services browser starts. In the left pane, Microsoft Windows 2000 for Network Attached Storage is automatically selected on the Appliance tab. In the right pane, Windows 2000 for Network Attached Storage starts.
3. Again, you are prompted to authenticate with the administrative user name and password.
4. Click **Administer this server appliance** to bring up the Microsoft Windows 2000 for Network Attached Storage GUI.

You are now ready to begin administering the appliance. Details for this task are described in “Initial setup and configuration”.

Initial setup and configuration

This section provides details on the initial setup and configuration of the NAS Gateway 300.

Note that if you are administering the NAS Gateway 300 without a keyboard, monitor, and mouse (“headless” mode), you can use one of two methods:

- Terminal Services, which provides full administrative function. (See “Terminal Services and the IBM NAS Administration console” on page 15.)
- Windows 2000 for Network Attached Storage, which provides a subset of the full administrative function in Terminal Services. (See “Windows 2000 for Network Attached Storage” on page 23.)

In general, you administer the appliance by adjusting information contained in the following task groups:

Note: In this example, you access the task groups through the Windows 2000 for Network Attached Storage Web-based GUI.

- “Setting the date and time”
- “Setting up the network” on page 11

Although you can modify multiple appliance and network attributes in each task group, the information given here is the minimum you need to know to administer the appliance and network.

You can find more information on administration elsewhere in this book and in the online help.

You can access these task groups in one of three ways:

1. Click the **Home** tab and then select the task group link.
2. Click the top tab associated with that task group.
3. Click the **Back** button on the browser until you arrive Home and then select the task group link.

Setting the date and time

To change the date and time, click **Date and Time**. (Remember that you can also access all of these task groups by clicking the titled tabs at the top of the page.) The Set Date and Time page appears, allowing you to adjust information as necessary.

Setting up the network

Note: All appliances have an initial default user name of “Administrator” and password of “password.”

As part of the Network task group, you can change the administrator password and (optionally) you can configure the properties of each network interface that resides on the appliance.

To change the administrator password, click **Change Administrator Password**. The Change Administrator Password page appears, allowing you to change the password. Note the warning on the page that any information that you enter can be viewed by others on the network. To prevent others from seeing your information, set up a secure administration Web site as described in the online help.

To change IP addresses, click **Interfaces**. The Network Adapters on Server Appliance page appears. Use this page primarily to change IP addresses from dynamic (DHCP, which is the system default) to static.

Note: During the initial setup, you should configure the nonplanar Ethernet adapter only. The NAS Gateway 300 engine uses the Ethernet adapter that is integrated on the planar board as the interconnect private network for clustering.

If you want to use an Ethernet adapter other than the default Ethernet adapter (in slot 2) as the network interface to be attached to the subnet, then you can change the order of precedence later with the Windows Networking Properties option. The order of precedence for the initial configuration is: PCI slot 2, then PCI slot 3.

Note that you might need to enable some of the NAS Gateway 300 NIC connections, because the NICs in slots 1, 3, and 4 are not enabled. During initial setup, the IAACU first looks for a 10/100 adapter in slot 2, which is enabled by default. If there is no adapter in slot 2, the IAACU looks for a Gigabit adapter card in slot 3 and it should be enabled. If the Gigabit adapter card is not enabled, right-click the adapter icon to enable it. After the initial setup, you can then enable all other NIC interfaces installed.

You must modify the adapter by completing the **IP** task (to modify IP configurations) and then choosing one or more of the following tasks, as appropriate:

- **DNS** (to modify DNS configurations)
- **WINS** (to modify WINS configurations)
- **HOSTS** (to modify host configurations)

Chapter 3. Configuration and administration tools

Attention

Changing the preloaded software configuration of this product, including applying or installing unauthorized service packs or updates to preinstalled software, or installing additional software products that are not included in either the preloaded image or on the Supplementary CD-ROM, might not be supported and could cause unpredictable results. For updated compatibility information, refer to the IBM Web site:

www.ibm.com/storage/nas

To correct problems with a preloaded software component, back up your user and system data. Then, use the Recovery CD-ROM to restore the preloaded software image.

The NAS Gateway 300 appliance comes with the following configuration programs that you can use to configure and administer the appliance:

- Terminal Services Client (page 15)

  This tool enables you to remotely administer the appliance.
- IBM Advanced Appliance Configuration Utility (IAACU, page 16)

  You can use the IAACU to set up and configure the network configuration on the appliance.
- Universal Manageability Services (page 20)

  This tool allows you to remotely manage your appliance using a Web browser.
- Windows 2000 for Network Attached Storage (page 23)

  This is a Web-based GUI for administrators who are not familiar with Windows.

This chapter describes these tools in general and then in detail.

Using a keyboard, monitor, and mouse for setup and configuration

It is recommended that you directly attach a keyboard, monitor, and mouse to the NAS Gateway 300 when:

- Initially setting up and configuring the device
- Changing or adding to RAID arrays defined on the fibre-attached storage
- Troubleshooting the device

Summary of configuration and administration tools

There are several ways to set up and administer the NAS Gateway 300.

Terminal Services Client

The Terminal Services Client, when installed on a workstation that is attached to the same network as the NAS Gateway 300 desktop. If you are familiar with administrative tasks using a Windows desktop, you can use Terminal Services.

See “Terminal Services and the IBM NAS Administration console” on page 15 for more information.

IBM Advanced Appliance Configuration Utility (IAACU)

The IBM Advanced Appliance Configuration Utility (IAACU) aids in setting up and reconfiguring the network configuration on your appliances. The IAACU agent works with the IAACU console to automatically detect the presence of appliances on the network.

After the appliance is detected by the IAACU console, you can use the IAACU to:

- Set up and manage the network configuration for the appliance, including assigning the IP address, default gateway, network mask, and DNS server to be used by the appliance. (See the note in "Setting up the network" on page 11, regarding the Ethernet adapter that is integrated on the planar board.)
- Start Universal Manageability Services on the appliance, enabling you to perform advanced systems-management tasks.

See "IBM Advanced Appliance Configuration Utility" on page 16 for more information.

Universal Manageability Services

Universal Manageability Services (UM Services) provides point-to-point remote management of client systems using a Web browser. Use UM Services to:

- Learn detailed inventory information about your computers, including operating system, memory, network cards, and hardware.
- Track your computers with features such as power management, event log, and system monitor capabilities.
- Integrate with Tivoli Enterprise, Tivoli NetView®, Computer Associates Unicenter, Microsoft SMS, and Intel® LANDesk Management Suite.

In addition, you can link to Windows 2000 for Network Attached Storage and Terminal Services from UM Services.

See "Universal Manageability Services" on page 20 for more information.

Windows 2000 for Network Attached Storage

The NAS Gateway 300 provides a Web-based GUI, Microsoft Windows 2000 for Network Attached Storage (Windows 2000 for NAS). Using Windows 2000 for NAS, you navigate through administrative task categories by clicking the appropriate tabs and then selecting a task from that category.

See "Windows 2000 for Network Attached Storage" on page 23 for more information.

Terminal Services and the IBM NAS Administration console

If you are familiar with Windows operating systems, you can use Terminal Services. In some cases, you must use Terminal Services to complete administrative tasks.

You can access Terminal Services in two ways:

1. Through the UM Services browser, as described in “Starting UM Services” on page 21.
2. By using the Terminal Services Client software.

Installing Terminal Services

To use the Terminal Services Client, complete the following steps to install it on the remote workstation and connect to the NAS Gateway 300 appliance:

1. Insert the Supplementary CD-ROM into the workstation CD-ROM drive.
2. Select **Start** → **Run**.
3. In the Open field, type (with quotation marks)

   "*x*:\Terminal Services Client\Disk 1\setup.exe"

   where *x* is the drive letter assigned to the CD-ROM drive.
4. Click **OK** to begin the Terminal Services Client Setup program.
5. Accept the defaults in each window that opens or refer to the Microsoft Windows documentation for more instructions.
6. When the Terminal Services Client Setup program completes, ensure that the workstation has network-connectivity to the NAS appliance so that you can administer the appliance.

Connecting to the desktop through Terminal Services

To connect to Terminal Services from your workstation, do the following:

1. Click **Start** → **Programs** → **Terminal Services** → **Terminal Services Client**.
2. In the Server field, select the computer name of the appropriate NAS Gateway 300. If that NAS Gateway 300 is not listed, type the IP address or the computer name of the NAS Gateway 300. The computer name is predefined as IBM5196-*xxxxxxx*, where *xxxxxxx* is the serial number located in the lower right corner of the bezel on the front of the appliance. If you have changed the computer name from the predefined value, use that name instead.

   Note: Although you can do so, it is recommended that you not change the default computer name to avoid the chance of propagating misidentification through the system. And, if you are using IBM Director to manage your appliance, and you change the default name, the default name continues to appear in IBM Director.
3. For **Size**, select a screen size in which the NAS Gateway 300 desktop will appear. It is recommended that you choose a size other than full screen.
4. Click **Connect** to start the Terminal Services Client session. A user login window opens.
5. Log in. Type *Administrator* in the Username field, type *password* in the Password field, and then click **OK** to log in. After you log in, you can begin using Terminal Services Client to configure and manage the NAS Gateway 300, as if a keyboard, mouse, and monitor were directly attached to it. The NAS Gateway 300 desktop contains a shortcut, titled **IBM NAS Admin**, to a special console, the IBM NAS Administration console.

IBM NAS Administration console

The IBM NAS Administration console includes all the standard functions provided by the standard Computer Management console available on any Windows 2000 desktop, plus the following functions specific to the NAS Gateway 300:

- Cluster Administration (see “Configuring clusters” on page 41)
- These advanced functions (see Chapter 6, “Managing and protecting the network and storage” on page 55):
  - FAStT MSJ
  - NAS Backup Assistant
  - Persistent Storage Manager
  - Tivoli SANergy

Determining who is using the network-attached storage

Occasionally, you might want to know who is using the network-attached storage. To determine this information:

1. Start a Windows Terminal Services session from the administrator’s console to the NAS Gateway 300.
2. Click the **IBM NAS Admin** icon on the desktop.
3. In the left pane, click **File Systems** → **Shared Folders** → **Sessions**.
4. The users currently using the storage are displayed. To close those sessions, use a right-click. Before you close a session, notify the user that you are going to close the session by clicking **Start** → **Programs** → **Accessories** → **Command Prompt**, and then issuing the **net send** *hostname messagetext* command.

IBM Advanced Appliance Configuration Utility

Note: Although you can do so, it is recommended that you not change the default computer name of your NAS appliance to avoid the chance of propagating misidentification through the system. Also, The IBM Advanced Appliance Configuration Utility (IACCU) depends on the original name to function.

The IBM Advanced Appliance Configuration Utility helps you to set up and reconfigure the network configuration on the NAS Gateway 300 appliance, as well as other IBM appliances.

The IAACU agent, preinstalled on the NAS Gateway 300 appliance, works with the IAACU console, a Java-based application that is installed on a network-attached system. You can use the IAACU as a systems-management console to automatically detect the presence of NAS Gateway 300 appliances on the network. After the NAS Gateway 300 appliance is detected by the IAACU console, use the IAACU to set up and manage the appliance’s network configuration, including assigning the IP address, default gateway, network mask, and DNS server to be used by the appliance. You can also use the IAACU to start Universal Manageability Services (UM Services) on the appliance, enabling you to perform more advanced systems-management tasks.

For networks that are not currently running DHCP servers, the IAACU is useful for automatically configuring network settings for newly added appliances, such as the NAS Gateway 300.

However, networks with DHCP servers will also benefit from using the IAACU because it enables you to reserve and assign the appliance IP address in an orderly, automated fashion. Even when you use DHCP and do not reserve an IP

address for the appliance, you can still use the IAACU to discover appliances and to start UM Services Web-based systems management.

Notes:

1. The IAACU configures and reports the TCP/IP settings of the first adapter (excluding the integrated Ethernet controller that is used for the interconnection of the two engines) on each appliance. The “first” adapter is defined by its position: if there is an adapter in slot 2, it is the first adapter; if there is an adapter in slot 3, it is the first adapter.

   Be sure to connect the first adapter to the same physical network as your systems-management console. You can do this by manually configuring the network adapter to be on the same subnetwork as the systems-management console.
2. The IAACU must be running to configure newly installed appliances automatically.
3. The system running the IAACU console automatically maintains a copy of its database (ServerConfiguration.dat) in the Advanced Appliance Configuration Station installation directory (Program files\IBM\iaaconfig). To remove previous configuration data, close the IAACU, delete this file, and then restart the utility. This deletes all previously configured Families. However, the IAACU will automatically discover connected appliances and their network settings.

Installing the IBM Advanced Appliance Configuration Utility

These instructions assume that you have installed and powered on the appliance according to the installation guide procedures. You are now ready to install the IAACU console application from the Supplementary CD-ROM.

Install the IACCU console application from the Supplementary CD-ROM onto a Windows NT 4.0 or Windows 2000 workstation that is attached to the same IP subnetwork to which the appliance is attached.

Note: The IAACU creates a private database that is specific to the IP subnetwork to which it is attached. Therefore, do not install it on more than one systems-management console residing on the same IP subnetwork.

For information on how to install the IAACU console, see “Installing the IBM Advanced Appliance Configuration Utility”.

After you install the IACCU console application, the following steps will take you to the point where you can administer the appliance.

1. Start the IACCU console application by clicking its icon.
2. On the left pane of the Advanced Appliance Configuration console, select the appliance to administer. Initially, the appliance name is **IBM5196**-*serial number*; the serial number is located in the lower right corner of the bezel on the front of the appliance.
3. Click **Start Web Management** to start the UM Services browser. This will open a separate Web browser window.
4. Proceed to “Accessing Universal Manageability Services” on page 9.

For more information on the IAACU, see “IAACU console” on page 18.

Initial network adapter selection and connection to the IAACU

Unlike the limited number of network adapter placement options in the previous release, in this release there are an increased number of network adapter types and locations from which you can connect. Assuming you have a keyboard and monitor attached, perform the following steps to take into account the new adapter placement options:

1. Decide which adapter will be used to connect to the IAACU, and connect the appropriate cable type.
2. Open the Network and Dial-up Connections panel. (From the desktop, right-click **My Network Places**, and select **Properties**.)
3. Determine the connection name of the adapter that you have selected to use. Move the mouse cursor over the adapter name, and a description of the adapter type will appear. If this is inconclusive, right-click the adapter, and select **Properties**. Under the General tab, click **Configure**. The line that contains the location information will provide the adapter's slot location. For example, *Location 1* means the adapter is in PCI slot number 1. Close the adapter properties panel.
4. On the Network and Dial-up Connections menu bar, select **Advanced** and then **Advanced Settings**. From the Connections menu, select the adapter's connection name. Then using the down arrow, move the selection down to the next-to-last position in the list. (The last entry in the list should be the *remote access connections*, shown as the telephone icon.) Save your changes by clicking **OK**.
5. The IAACU will now detect the appliance using the adapter that you have just enabled.

IAACU agent

The IAACU agent is preinstalled on the NAS Gateway 300 appliance.

After you connect the NAS Gateway 300 to your network, the IAACU agent automatically reports the appliance serial number and type, the MAC address of its onboard Ethernet controller, and whether DHCP is in use by the appliance. Furthermore, it reports the host name, primary IP address, subnet mask, primary DNS server address, and primary gateway address if these are configured on the system.

Note: The IAACU agent periodically broadcasts the appliance IP settings. To prevent the service from broadcasting this data periodically, stop the iaaconfig service.

IAACU console

The IAACU console is a Java application that you install on one system in your network for use as a systems-management console. For information on how to install the IAACU console, see "Installing the IBM Advanced Appliance Configuration Utility" on page 17.

Note: The IAACU creates a private database that is specific to the IP subnetwork to which it is attached. Therefore, do not install it on more than one systems-management console residing on the same IP subnetwork.

The IAACU console enables you to:

- Automatically discover NAS Gateway 300 appliances, as well as other IBM appliances that run the IAACU agent and are attached to the same physical subnet as the IAACU console.
- Use a GUI-based application to configure the appliance network settings.

  Use the IAACU to assign network parameters such as IP addresses, DNS and gateway server addresses, subnet masks, and host names.
- Start UM Services Web-based systems-management console.

  Launch UM Services on your appliances and perform advanced systems-management tasks on a selected appliance with a single mouse click.

The IAACU console is divided into two panes:

- **Tree View Pane**

  The Tree View Pane, located on the left side of the IAACU console window, presents a list of all discovered NAS Gateway 300 appliances. The Tree View Pane also includes groups for appliances that were not configured using the IAACU or that have IP addresses that conflict with other devices on your network. When you click any item in the Tree View, information about that item (and any items that are nested below that item in the tree view) appears in the Information Pane.
- **Information Pane**

  The Information Pane, located on the right side of the IAACU console, displays information about the item that is currently selected in the Tree View Pane. The information that appears in the Information Pane varies depending on the item that is selected. For example, if you select the All Appliances item from the Tree View Pane, the Information Pane displays configuration information (IP settings, host name, serial number, and so on) about each of the NAS Gateway 300 appliances that have been discovered by the IAACU console.

The IAACU console also features the following menus:

File Use the File menu to import or export the IAACU console configuration data, to scan the network, or to exit the program.

Appliance
Use the Appliance menu to remove a previously discovered appliance from a group.

Help Use the Help menu to display product information.

Discovering NAS Gateway 300 appliances

Any NAS Gateway 300 appliance, or other IBM appliance, that is running and is connected to the same subnet as the system running the IAACU console is automatically discovered when you start the IAACU console. Discovered appliances appear in the IAACU console tree view (in the left pane of the IAACU console window). Every discovered appliance is listed in the tree view under All Appliances.

Universal Manageability Services

Universal Manageability Services (UM Services) is a Windows application that functions as both a stand-alone management tool for the system it is installed on and a client to IBM Director.

As a Director Client, it receives and sends information to the Director Server as controlled from the IBM Director Console.

As a stand-alone tool, it provides a Web-browser based interface and a Microsoft Management Console (MMC) interface, where you can view the system status, perform certain management tasks and configure alerts.

The UM Services GUI enhances the local or remote administration, monitoring, and maintenance of IBM systems. UM Services is a lightweight client that resides on each managed computer system. With UM Services, you can use a Web browser and UM Services Web console support to inventory, monitor, and troubleshoot IBM systems on which UM Services is installed.

This “point-to-point” systems-management approach, in which you use a Web browser to connect directly to a remote-client system, enables you to effectively maintain IBM systems without requiring the installation of additional systems-management software on your administrator console.

In addition to point-to-point systems-management support, UM Services also includes support for UM Services Upward Integration Modules. These modules enable systems-management professionals who use any supported systems-management platform (including Tivoli Enterprise, CA Unicenter TNG Framework, and Microsoft Systems Management Server [SMS]) to integrate portions of UM Services into their systems-management console. Because it was designed to use industry-standard information-gathering technologies and messaging protocols, including Common Information Model (CIM), Desktop Management Interface (DMI), and Simple Network Management Protocol (SNMP), UM Services adds value to any of these supported workgroup or enterprise systems-management platforms.

You can use UM Services to perform the following tasks:

- View detailed information about your computers, including operating system, memory, network cards, and hardware.
- Track your computers with features such as power management, event log, and system monitor capabilities.
- Upwardly integrate with Tivoli Enterprise, Tivoli Netview, Computer Associates Unicenter, Microsoft SMS, and Intel LANDesk Management Suite.

Complete documentation on how to use UM Services is included on the Documentation CD-ROM that came with the appliance.

System requirements

The UM Services client is preinstalled on the NAS Gateway 300 appliance. However, you must have a Web browser installed on your systems-management console. It is recommended that you set Microsoft Internet Explorer 5.x (or later) as the default browser.

Notes:

1. You must install the optional Java Virtual Machine (VM) support to access a client system running UM Services.
2. If you reinstall Internet Explorer after installing UM Services, you must reapply the Microsoft VM update. The UM Services client requires Microsoft VM Build 3165 or later. Download the latest Microsoft VM from www.microsoft.com/java
3. If you install UM Services before you install MMC 1.1 (or a later version), you will not have an icon for MMC in the IBM Universal Manageability Services section of the Start menu.

Starting UM Services

You can use IAACU or Terminal Services Client to configure the network setting remotely, or you can attach a keyboard and mouse to your appliance and configure the Network settings using the Windows Control Panel. After you have configured the network settings for your appliance, you are ready to use UM Services.

To start UM Services:

1. Start a Web browser and then, in the **Address** or **Location** field of the browser, type:

   http://*ip_address***:1411**

   where *ip_address* is the IP address of the NAS Gateway 300, and then press **Enter**.

   Or, type:

   http://*computer_name***:1411**

   where *computer_name* is the computer name of the NAS Gateway 300. The computer name is predefined as: IBM5196-*xxxxxxx*, where *xxxxxxx* is the serial number located in the lower right corner of the bezel on the front of the appliance.

   If you have changed the computer name from the predefined value, use that name instead. A user log in window opens.

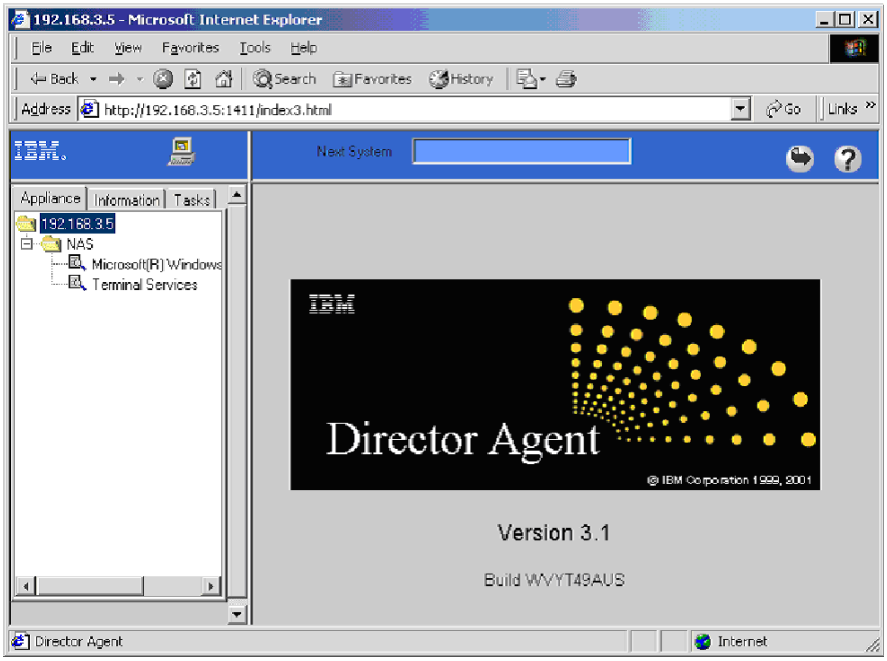

Figure 2. UM Services default page

2. Type *Administrator* in the User Name field, and type *password* in the Password field. You can leave the Domain field blank. Make sure the **Save this password in your password list** check box is ***not*** selected, and then click **OK**.

   Note: To ensure system security, change the Administrator password from "password" to something else. After you do, or if you create another user in the Administrator group in the future, use your new username/password combination instead of the default username/password combination.

   The first time you connect, you might be prompted to install XML and Swing components. Follow the on-screen instructions to install these components and then close and restart Internet Explorer before you proceed.

   You are now connected to the NAS Gateway 300 through UM Services. In addition to the standard UM Services functionality, the appliance includes functionality for administering the appliance, available from the Appliances tab in the left pane of the UM Services browser. The default view (in the right pane of the UM Services browser) when you connect to the appliance is Windows 2000 for NAS. The other selectable view in the Appliance tab is Windows 2000 Terminal Services, which displays a Terminal Services Web Connection page.
3. To start Windows 2000 for NAS, click **Administer this server appliance** in the right pane of the UM Services browser. To connect to the NAS Gateway 300 and manage it as though you were running Terminal Services Client from the desktop, select **Terminal Services** in the Appliance tab of the UM Services browser, and then follow the instructions for connecting to the NAS Gateway 300 using Terminal Services described in "Terminal Services and the IBM NAS Administration console" on page 15.

Launching UM Services from the configuration utility

You can use the IAACU to launch UM Services on the NAS Gateway 300 appliances.

Note: The selected appliance must be running UM Services as a UM Services client. Also, the systems-management console (the system that is running the IAACU console) must use a Web browser that is supported for use with UM Services. If you have not used UM Services from this system, you must install several plug-ins before proceeding.

To use the IAACU console to start UM Services on an appliance:

1. Click the appliance in the IAACU console Tree View Pane.

   When you select the appliance from the tree view, information about the selected appliance appears in the Information Pane.
2. Click **Start Web-Based Management**.

   Your default Web browser starts, loading the UM Services browser automatically.
3. Log in to the UM Services browser. Refer to Step 2 on page 22 for login instructions.

For more information on using UM Services to manage your appliances, see the *Universal Manageability Services User's Guide*, included on the Documentation CD-ROM that came with the appliance.

Windows 2000 for Network Attached Storage

While you can perform most administrative tasks using Windows 2000 for NAS, you must use Terminal Services Client for some advanced tasks. See "Terminal Services and the IBM NAS Administration console" on page 15 for more information.

Task categories available to you through Windows 2000 for NAS include:

- Status
- Network
- Disks
- Users
- Shares
- Maintenance
- Controller

To start Windows 2000 for NAS, use one of these methods:

- UM Services, described in Step 3 on page 22
- Web browser, by entering **http://***ip_address***:8099** or **http://***computer_name***:8099** and then logging on to the NAS Gateway 300
- NAS Gateway 300 desktop while using Terminal Services Client and starting a browser

You can access online help for Windows 2000 for NAS in two ways:

1. Click the **Help** button at the top of any Web page. This displays a table of contents that you can navigate to find help for any Windows 2000 for NAS task.
2. Click the question mark (**?**) button at the top of any Web page. This displays context-sensitive help for the task you are currently performing.

Determining the tool to use

Table 2 suggests which tool to use for specific functions, but does not list all options or combinations. The administrator's training level or preferences might determine an alternate approach from that suggested in the table.

Table 2. Summary of configuration and administration tools for the NAS Gateway 300

Administration tool	Main functions
Windows Domain Controller (not NAS appliance)	Users and user groups can be defined and authenticated by the Windows Domain Controller, although this is not required.
IBM Advanced Appliance Configuration Utility (IAACU)	Access a headless NAS Gateway 300 node, particularly for the initial setup of the network connectivity. (Alternatively, you can attach a keyboard, mouse, display to each node of the NAS Gateway 300.) IAACU enables you to: • Set time, date, and initial network connectivity parameters • Access to Windows 2000 for NAS GUI, Terminal Services (NAS Desktop), and Universal Manageability Services
Windows 2000 for NAS GUI	Provides ease-of-use administration, but not all the capabilities of Terminal Services and IBM NAS Administration. The GUI enables you to: • Configure networking connectivity, private (for clustering) and public LAN connections • Create and format logical drives • Join domains • Set up access permissions and disk quotas for CIFS, NFS, HTTP, FTP, and Novell NetWare shares • Use Persistent Storage Manager
IBM NAS desktop and IBM NAS Admin program, through a Terminal Services session or a directly-connected keyboard and monitor	Provides in-depth administration of all aspects of NAS Gateway 300. Provides all of the Windows 2000 for NAS GUI functions above, plus the ability to: • Use NAS Backup Assistant, or NT Backup and Restore wizard • Learn detailed inventory information about hardware, OS, and so on, using Universal Manageability Services • Cluster administration: – Set up cluster – Define failover for each volume – Manually fail over cluster resources – Set up cluster resource balancing by assigning preferred node • Diagnose system problems: – Check Ethernet adapters using PROSet II and 10/100 Quad-Port Ethernet adapter using SLICuser – Check Fibre Channel card using FAStT MSJ
Disaster Recovery	Restores a previously saved PSM image of the system partition to a failed machine. This restores all configuration information on the failed node. You create the recovery boot diskette from the PSM tools in the Windows for 2000 NAS GUI.
Recovery CD-ROM Set	Reinstalls the software to the original state as shipped on the machine; however, does not restore configuration information (configuration changes you applied to the original shipped configuration are lost). You must first boot with the Recovery Enablement Diskette, and then reboot with the Recovery CD-ROM. To create the Recovery Enablement Diskette, run enablement_disk_*x.y*.exe (where *x.y* are the version number of the disk), located on the Supplementary CD-ROM. You will be prompted to insert a blank disk into drive **a:**.
Integrated System Management Processor (ISMP) configuration program	Configures the ISMP that is integrated on the engine planar board.

Table 2. Summary of configuration and administration tools for the NAS Gateway 300 (continued)

Administration tool	Main functions
Remote Supervisor Adapter (RSA) configuration program	Configures the optional RSA.

Telnet Server support

Attention: When you Telnet to another machine, your user name and password are sent over the network in plain, unencrypted, text.

The NAS Gateway 300 includes Telnet server capability. The Telnet server provides limited administrative capability. This can be useful in cases where you need to remotely administer the NAS Gateway 300, but do not have access to a Windows-based workstation (from which you could remotely administer the appliance through a supported Web browser or Terminal Services Client).

To access the NAS Gateway 300 from any Telnet client, specify the IP address or host name of the NAS Gateway 300, then log in using an ID and password (defined on the NAS Gateway 300) with administrative authority. From the command line, you can issue DOS-like commands (such as **dir** and **cd**), and some UNIX-like commands (such as **grep** and **vi**). You can launch some applications, but only character-mode applications are supported.

By default, the Telnet server is disabled. To enable the Telnet server, from the Windows 2000 for NAS user interface, go to the Network task group, then select **Telnet**. On the Telnet Administration Configuration page, select the **Enable Telnet access to this appliance** check box. If you do not require Telnet access to the NAS Gateway 300, then it is recommended that you leave the Telnet server disabled.

SNMP support

Support for the Simple Network Management Protocol (SNMP) is enabled. To manage the NAS Gateway 300 from an SNMP-capable management application, you must install the management information base (MIB) files for various components of the NAS Gateway 300 on the management application workstation, so that the application can recognize those SNMP elements (values, alerts, and so on) supported by the components.

Chapter 4. Setting up storage

This chapter gives details for setting up and configuring the fibre-attached storage for the NAS Gateway 300.

Note: You need to configure the storage on one node only. For the other node (the joining node), the only part of shared storage setup that you will need to complete is assigning drive letters on the shared storage, making sure that the drive letters are the same as those on the first node.

Configuring arrays and logical drives on the fibre-attached storage

You will need to configure the RAID arrays and logical drives (LUNs) on the fibre-attached storage, or contact your disk administrator to do the configuration. The specific procedures for configuring arrays and LUNs are defined by the fibre-attached storage, so its documentation must be consulted for those procedures.

You will need to get the World Wide Name (WWN) of the Fibre Channel host adapter to set up an association between the NAS Gateway 300 node and each LUN that you create on the fibre-attached storage. To get the WWN:

1. Click the **IBM NAS Admin** icon.
2. Select **NAS Management**.
3. Select **Storage**.
4. Select **NAS Utilities**.
5. Select **IBM Fibre WWN**.

   The IBM Fibre WWN panel will display information for each Fibre Channel adapter installed in your NAS Gateway 300, including PCI slot number and World Wide Name. The slot number given is not the physical PCI slot location of the adapter within the system but rather a reference slot to the PCI bridge of the PCI system. If you have only one Fibre Channel adapter, the information for that adapter will appear immediately in the fields on the right side of the panel.

   If you have multiple Fibre Channel adapters, select the radio button next to each adapter listing until you find the one whose displayed PCI slot number matches the PCI slot number of the actual adapter. Make a note of the World Wide Name that is displayed so that you can provide this when configuring your fibre-attached storage to be accessed through the NAS Gateway 300.

There is an additional requirement imposed on the configuration of the fibre-attached storage: one LUN must be defined for the Quorum drive. The Quorum drive is used by Microsoft Cluster Service to manage clustered resources, including the fibre-attached storage. The requirements for the Quorum drive LUN are the following:

- The array in which you create the Quorum drive LUN should be RAID 5. This is recommended for performance and redundancy.
- The Quorum drive LUN should be at least 500 MB in size, but no larger than 1 GB.
- The Quorum drive LUN should be completely dedicated for use by Microsoft Cluster Service; no other data should be stored on this LUN. However, it is acceptable for the array in which this LUN is created to have other LUNs.
- See “Recovering from a corrupted Quorum drive” on page 52 in the event of a power loss to both nodes or a hardware failure that corrupts Quorum data.

You can configure other arrays and LUNs for user data as required. However, do not create any arrays that are RAID 0, as this is not supported. It is recommended that all arrays be RAID 5 arrays.

Expanding the LUN

LUN expansion is enabled by the DiskPart command line utility. Using DiskPart and array/LUN management software, you can dynamically expand an existing logical drive into unallocated space that exists in a LUN.

Note that you cannot use DiskPart to dynamically expand an existing LUN in an array. You can do this only with array/LUN management software such as Storage Manager Application. DiskPart cannot change the size of the drive that the external storage has configured; it can only change how much of the drive that Windows can use.

Attention: It is highly recommended that you always perform a backup of your data before using the DiskPart utility.

To perform LUN expansion, use the following two DiskPart commands:

select This command focuses on (selects) the volume that you want to expand. The format of the command and its options are

```
select volume[=n/l]
```

You can specify the volume by either index, drive letter, or mount point path. On a basic disk, if you select a volume, the corresponding partition is put in focus. If you do not specify a volume, the command displays the current in-focus volume.

extend

This command extends the current in-focus volume into contiguous unallocated space. The unallocated space must begin where the in-focus partition ends. The format of the command and its options are

```
extend [size=n]
```

where *size* is the size of the extension in MB.

Note that if the partition had been formatted with the NTFS file system, the file system is automatically extended to occupy the larger partition, and data loss does not occur. However, if the partition had been formatted with a file system format other than NTFS, the command is unsuccessful and does not change the partition.

DiskPart blocks the extension of only the current system or boot partition.

Several other commands are useful when you expand the LUN:

assign

Use this command to assign a letter or mount point to the current selected (in-focus) partition. If you do not specify a drive letter, the next available drive letter is assigned. If the letter or mount point is already in use, an error is generated.

You can use this command to change the drive letter that is associated with a removable drive. The drive letter assignment is blocked on the system, boot, or paging volumes. You cannot use this command to assign a drive letter to an OEM partition or any globally unique identifier (GUID) partition table (GPT) partition, other than the Msdata partition.

The format of the command and its options are:

```
assign [letter=l] or [mount=path]
```

convert

You can use several commands to convert disks. The format and options for each of the commands are:

```
convert mbr
convert gpt
convert dynamic
convert basic
```

convert mbr sets the partitioning style of the current disk to master boot record (MBR). The disk can be a basic disk or a dynamic disk but the disk must not contain any valid data partitions or volumes.

convert gpt sets the partitioning style of the current disk to GPT. The disk may be a basic or a dynamic disk but it must not contain any valid data partitions or volumes. This command is valid only on Itanium™-based computers; it can be unsuccessful on x-86-based computers.

convert dynamic changes a basic disk into a dynamic disk. The disk can contain valid data partitions.

convert basic changes an empty dynamic disk to basic.

list

You can use several commands to display summaries of disk configuration. The format for each of the commands is:

```
list disk
list partition
list volume
```

list disk displays summary information about each disk in the computer. The disk with the asterisk (*) has the current focus. Only fixed disks (for example, IDE or SCSI) or removable disks (for example, 1394 or USB) are listed. The removable drives are not displayed.

list partition displays information about each partition on the in-focus disk.

list volume displays information about each volume in the computer.

Using DiskPart with clustering

To expand a volume, you first need to add free space at the end of the volume that you want to expand. This free space is now directly behind the existing volume that is to be extended.

Verify which computer is the owner of the volume that you want to expand:

1. Open the IBM NAS Admin utility and select **Cluster Tools** → **Cluster Administration**.
2. Select the group where the disk is located and see what computer is the owner of the disks.
3. Shut down the other node. You can move all cluster resources to the other server prior to this shutdown by right-clicking the group in the Cluster Administrator and selecting **Move Group**.

Important

Stop all I/O to the disk while performing this procedure by setting offline all the resources in the cluster group that contains the disks. One way to do this is by bringing offline the cluster group that contains the disk in the cluster administration utility, and then bringing online only the physical disk. This should close all open handles to the disk.

4. Open a command prompt window and issue the **diskpart** command. Or, in the IBM NAS Admin utility, select **Storage** → **Diskpart** → **Diskpart**.
5. In the DiskPart utility, you can list the volumes in the computer by issuing the **List Volume** command. Check the number for your volume by comparing the label.
6. Select the volume you want to extend by entering:

   ```
   select volume X
   ```

   where *X* is the number of the volume you want to extend.
7. Issue the command **extend** to extend the selected volume.
8. Rescan the disks with the **rescan** command and list all volumes with the capacity changed.
9. Exit the utility by issuing the **exit** command.
10. Bring online all resources in the cluster group containing the disk by right-clicking the group and selecting **Bring Online**.
11. Power on the node that is down and move the group that owns the disk to ensure proper operation. When the node is up, perform these steps:
    a. Open the IBM NAS Admin utility and select **Cluster Tools** → **Cluster Administration**.
    b. Right-click the group in the Cluster Administrator and select **Move Group**.
12. When the group has been moved, check that the volume size has been increased by opening the IBM NAS Admin tool and selecting **Storage** → **Disk Management (local)**.

Formatting the logical drives

Note the following restrictions when formatting logical drives:

1. Disk 0 (the internal hard disk drive) is an 18-GB drive, preformatted into two partitions: a 6-GB partition (label System, drive letter C:) and a 12-GB partition (label MAINTENANCE, drive letter D:). Do not reformat or repartition these partitions. Doing so could wipe out important data and seriously impair the functioning of your system.
2. Do ***not*** upgrade any disks to dynamic. Only basic disks are supported for clustering. In addition, all partitions used for clustering must be primary partitions.
3. Do ***not*** use drive letter F: as a volume drive letter. This drive letter is reserved for Persistent Storage Manager-based backup using NAS Backup Assistant.

Follow this procedure to format logical drives.

1. Open IBM NAS Admin and select **Disk Management (Local)** in the Storage folder.
2. At the Write Signature and Upgrade Disk Wizard, click **Cancel**.

3. Right-click **Disk 1** and select **Write Signature**.
4. Write Signature to all disks that will be accessed by the NOS (all disks in view).
5. On each disk:
   a. Right-click and select **Create Partition** and click **Next**.
   b. Select **Primary Partition** and click **Next**.
   c. Select the entire disk size and click **Next**.
   d. Specify *NTFS* as the file system. If this is the Quorum disk, specify *Quorum disk* as the Volume Label; otherwise, specify whatever name you want to assign to the partition.
   e. Click **Finish**. Do not enable disk compression, and select **Finish**.
6. Format all other drives but do not enable compression. Use all space available for each drive for each logical drive assigned by the operating system. Assign a drive letter of *G* for the first drive (the Quorum drive), *H* for the second drive (the first user volume), and so on.
7. Shut down the first node and make sure the drives are available on the joining node. Change the drive letters to match those on the first node. Rescan the disks if the LUNs do not show up.

At this point, you have completed shared storage setup. You can now continue with Chapter 5, “Completing networking, clustering, and storage access setup” on page 33.

Chapter 5. Completing networking, clustering, and storage access setup

The NAS Gateway 300 uses Microsoft Cluster Server (MSCS) software to provide clustering technology for your storage. Clustering ensures availability of the storage, regardless of individual component failure.

After installing the clustering function, you can use Cluster Administration to set up the failover function. Then, if a node or a node's component were to fail, the NAS Gateway 300 detects the failure and begins a failover process in less than 10 seconds, and completes the process within 60 seconds. Failover/Failback includes Active/Active support for the CIFS and NFS protocols.

Active/Active support is available for HTTP and FTP. See the online cluster administration guide for this procedure.

Novell NetWare and Apple Macintosh shares are available on both nodes, but not through clustering services. If either node fails, the shares become unavailable until the node is brought back up.

This chapter gives the details for installing and initially configuring MSCS on the NAS Gateway 300. Administrative concepts and procedures are provided in the online help and at the following Web sites:

- www.microsoft.com/windows2000/library/technologies/cluster/default.asp
- www.microsoft.com/ntserver/support/faqs/clustering_faq.asp
- http://support.microsoft.com/default.aspx?scid=kb;EN-US;q248025

Networking setup

Attention: Before you install cluster services on the first node, make sure that the joining node is shut down (see "Shutting down and powering on the NAS Gateway 300" on page 87 for more information on shutting down the NAS Gateway 300). This is required to prevent corruption of data on the shared storage devices. Corruption can occur if both nodes simultaneously write to the same shared disk that is not yet protected by the clustering software.

Note: After you complete this procedure on the first node, you must complete it on the joining node with the first node shut down.

Configuring the interconnect (private) network adapter

To configure the interconnect (private) network adapter, perform the following steps on both nodes. The **Private** connection is the "heartbeat" interconnect for the cluster.

1. Right-click **My Network Places** and then select **Properties**.
2. Select the network connection that uses the integrated Ethernet controller.
3. Right-click the adapter icon and click **Properties**.
4. Click **Configure**, select the **Advanced** tab, and verify that the following characteristics are set:

Link speed and Duplex

100 Mbps / Full Duplex

1000 Mbps / Full Duplex

5. Click **OK**.
6. If prompted to restart the node, select **No**.
7. In the Properties panel for the integrated Ethernet controller connection, select **Internet Protocol (TCP/IP)** from the components section, and click **Properties**.
8. The default IP addresses should be:
   - 10.1.1.1 for the first node
   - 10.1.1.2 for the joining node

   If they are not, it is recommended that you set them to those values.
9. Ensure a Subnet Mask of 255.255.255.0.
10. Click **Advanced**, and select the **WINS** tab.
11. Select the **Disable NetBIOS over TCP/IP** radio button.
12. Click **OK**.
13. Select **Yes** at the prompt to continue using an empty Primary WINS address.
14. Click **OK** on the Internet Protocol (TCP/IP) Properties panel.
15. Click **OK** on the Local Area Connection Properties (Private) panel.
16. Rename the connection to **Private**.

Configuring the public local area connection

Note: 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. It is strongly recommended that you set static IP addresses for all network adapters in the cluster. If IP addresses are obtained through DHCP, access to cluster nodes could become unavailable if the DHCP server goes down.

To configure each public local area connection, perform the following steps on each node:

1. Right-click **My Network Places**, then click **Properties**.
2. Select a **Local Area Connection**.

   When you perform this step, the connection that uses the integrated Ethernet controller is the private connection. The other active connection is the public connection. Use that other active connection for this step and the next step.
3. To rename the connection, click **Rename**, and then type (for example) **Public 1**, and press **Enter**. Ensure that local area connection names are unique.

   When you perform these renaming steps for the joining node, ensure that the local area connection name for each physically connected network is identical on each server. See Table 3 on page 35 for a further example.
4. Use the networking information in Table 1 on page 8 to enter the networking addresses:
   a. Right-click **My Network Places**.
   b. Click **Properties**.
   c. Right-click the **Public** icon, and then click **Properties**.
   d. Select **Internet Protocol (TCP/IP)**.

e. Click **Properties**, select **Use the following IP address:**, and enter the addresses for the IP, subnet mask, default gateway, and preferred DNS server.

5. If needed, configure the DNS, WINS, HOSTS, or whichever method you will be using for names resolution. To view this information, click **Advanced** on the Properties window.

   Note: NetBIOS should be disabled.

6. Click **OK** on each panel to return to the Properties window.

Do not place paired adapters on the same IP network unless you are going to use adapter teaming or adapter load balancing.

Verifying network connectivity and names resolution

Verify network connectivity and names resolution after you have installed clustering on the joining node.

To verify that the private and public networks are communicating properly:

1. Click **Start → Run**, type **cmd** in the text box, and click **OK** to bring up an MS-DOS prompt.
2. Type **ping** *ipaddress* where *ipaddress* is the IP address for the corresponding network adapter in the other node, and press **Enter**.

For example, assume that the IP addresses are set as follows:

Table 3. Example of local area connection names and network adapter IP addresses

Node	Local area connection name	Network adapter IP address
1	Private	10.1.1.1
1	Public 1	192.168.1.12
1	Public 2	192.168.2.12
2	Private	10.1.1.2
2	Public 1	192.168.1.13
2	Public 2	192.168.2.13

In this example, you would type **ping 192.168.1.12** and **ping 192.168.2.12** for the first node, and you would type **ping 192.168.1.13** and **ping 192.168.2.13** for the joining node. You can do this from any machine that is physically connected to the network of each node.

To view the addresses, use the **ipconfig** command on each node:

1. Click **Start → Run**, type **cmd** in the text box, and click **OK** to bring up an MS-DOS prompt.
2. Type **ipconfig /all** and press **Enter**. IP information should appear for all network adapters in the machine.

Checking or changing the network binding order

The clustering function requires the following binding order:

- Private
- Public 1
- Public 2
- :

The top-most connection is first in the binding order. Typically, this is the most frequently used network adapter.

To check the binding order and change it:

1. From the desktop, right-click **My Network Places** and then select **Properties**.
2. Select **Advanced Settings** from the Advanced menu.
3. Reorder the position of the adapters by selecting them, then pressing the up or down arrow keys, then clicking **OK**.

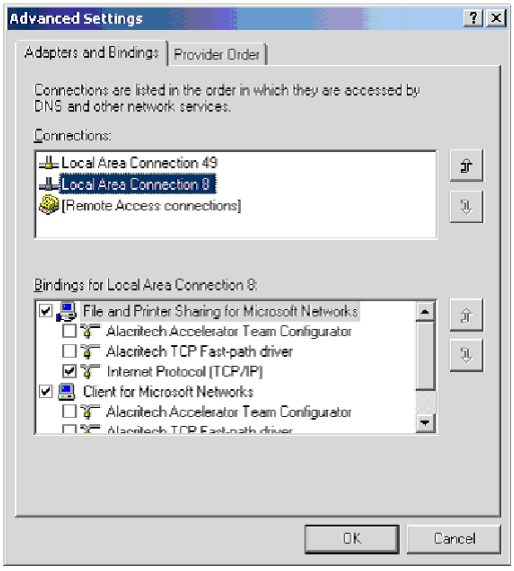

Figure 3. Advanced Settings for binding order

If prompted to restart, click **No**. If you change the binding order, you do not have to reboot until after you join the node to the domain.

Joining a node to a domain

For the Windows Cluster service to form a cluster on a given node, the service must authenticate with a Windows domain. If a Windows domain controller is available on a public network to which both nodes will be physically connected, follow the instructions below. Otherwise, follow the instructions in “Creating an Active Directory Domain” on page 37 to create a new domain that will encompass just the cluster itself. All nodes in the cluster must be members of the same domain and be able to access a Primary Domain Controller (PDC) and a DNS server.

All nodes in the cluster must be members of the same domain and be able to access a PDC and a DNS server.

1. Right-click **My Computer**, and click **Properties**.
2. Click **Network Identification**. The System Properties dialog box displays the full computer name and workgroup or domain.
3. Click **Properties** and perform these steps to join a domain:
   a. Select the **Domain** radio button.
   b. Type the name of your domain and click **OK**.
   c. When prompted, enter the Administrator user ID and password and click **OK**.
4. Close the **System Properties** window.
5. Restart the node, and proceed with "Cluster setup" on page 39.

   After the computer restarts, it is recommended that you do not log on to the domain. If you do, you will see the Windows 2000 Configure Your Server window. Click the **I will configure this server later** radio button, and then click the **Next** button. On the next window, clear the **Show this screen at startup** check box and click **Finish**.

Creating an Active Directory Domain

The Windows 2000 Cluster service runs in the context of a Windows-based domain security policy, typically created specifically for the Cluster service to use. For the Cluster service to form a cluster on a given node, the service must first authenticate itself using the credentials of this policy. A domain controller must be available for the domain that issued the policy for authentication to occur. If the Cluster service does not have access to a domain controller, it cannot form a cluster.

Note: For Active Directory to function properly, DNS servers must provide support for Service Location (SRV) resource records described in RFC 2052, *A DNS RR for specifying the location of services (DNS SRV)*. SRV resource records map the name of a service to the name of a server offering that service. Active Directory clients and domain controllers use SRV records to determine the IP addresses of domain controllers. Although not a technical requirement of Active Directory, it is highly recommended that DNS servers provide support for DNS dynamic updates described in RFC 2136, *Observations on the use of Components of the Class A Address Space within the Internet*.

The Windows 2000 DNS service provides support for both SRV records and dynamic updates. If a non-Windows 2000 DNS server is being used, verify that it at least supports the SRV resource record. If not, it must be upgraded to a version that does support the use of the SRV resource record. A DNS server that supports SRV records but does not support dynamic update must be updated with the contents of the Netlogon.dns file created by the Active Directory Installation wizard while promoting a Windows 2000 Server to a domain controller.

By default, the Active Directory Installation wizard attempts to locate an authoritative DNS server for the domain being configured from its list of configured DNS servers that will accept a dynamic update of an SRV resource record. If found, all the appropriate records for the domain controller are automatically registered with the DNS server after the domain controller is restarted.

If a DNS server that can accept dynamic updates is not found, either because the DNS server does not support dynamic updates or because dynamic updates are

not enabled for the domain, the following steps are taken to ensure that the installation process is completed with the necessary registration of the SRV resource records:

1. The DNS service is installed on the domain controller and is automatically configured with a zone based on the Active Directory domain.

   For example, if the Active Directory domain that you chose for your first domain in the forest was example.microsoft.com, a zone rooted at the DNS domain name of example.microsoft.com is added and configured to use the DNS service on the new domain controller.
2. A text file containing the appropriate DNS resource records for the domain controller is created.

   The file called Netlogon.dns is created in the %systemroot%\System32\config folder and contains all the records needed to register the resource records of the domain controller. Netlogon.dns is used by the Windows 2000 NetLogon service and to support Active Directory for non-Windows 2000 DNS servers.

If you are using a DNS server that supports the SRV resource record but does not support dynamic updates (such as a UNIX-based DNS server or a Windows NT® Server 4.0 DNS server), you can import the records in Netlogon.dns into the appropriate primary zone file to manually configure the primary zone on that server to support Active Directory.

If you are configuring the first node, complete these steps to create the Active Directory Domain Controller:

1. Start the Active Directory Installation Wizard from the IBM NAS Admin console by selecting **Local Domain Controller Setup** in the Cluster Tools folder.
2. Read the first page and click **Next**.
3. On the Domain Controller Type page, select **Domain controller for a new domain**; then click **Next**.
4. On the Create Tree or Child Domain page, select **Create a new domain tree**; then click **Next**.
5. On the Create or Join Forest page, select **Create a new forest of domain trees**; then click **Next**.
6. On the New Domain Name page, type the full DNS name for the new domain. Write down this value now; it will be needed later on. Click **Next**.
7. On the NetBIOS Domain Name page, click **Next**.
8. On the Database and Log Locations page, click **Next** to accept the default values.
9. On the Shared System Volume page, click **Next** to accept the default value.
10. On the Permissions page, select **Permissions compatible only with Windows 2000 servers**; then click **Next**.
11. On the Directory Services Restore Mode Administrator password, type your password. Write down this value now; it will be needed later on. Click **Next**.
12. On the Summary page, review the values; then, click **Next** and wait for the Configuring Active Directory process to complete.
13. On the Active Directory Installation Wizard page, click **Finish**. If prompted, defer the system reboot.
14. Restart the node.

If you are configuring the joining node, perform these tasks to join this node to the existing Active Directory Domain Controller (previously created on the first node):

1. Start the Active Directory Installation Wizard from the IBM NAS Admin console by selecting **Local Domain Controller Setup** in the Cluster Tools folder.
2. Read the first page and click **Next**.
3. On the Domain Controller Type page, select **Additional domain controller for an existing domain**; then click **Next**.
4. On the Network Credentials page, type the user name (Administrator), password, and domain name (enter the Active Directory Domain name created on the other node); then click **Next**.
5. On the Additional Domain Controller page, type the full DNS name of the existing domain; then click **Next**.
6. On the Database and Log Locations page, click **Next** to accept the default values.
7. On the Shared System Volume page, click **Next** to accept the default value.
8. On the Directory Services Restore Mode Administrator password page, type your password. Write down this value now; it will be needed later on. Click **Next**.
9. On the Summary page, review the values; then click **Next**. Wait for the Configuring Active Directory process to complete.
10. On the Active Directory Installation Wizard page, click **Finish**. If prompted, reboot the system now.

Cluster setup

At this step, you have completed the cluster installation steps on each node and are ready to set up the cluster.

Perform the following steps:

1. Power on the first node. The joining node should be shut down (see “Shutting down and powering on the NAS Gateway 300” on page 87 for more information on shutting down the NAS Gateway 300).
2. To begin setting up the cluster on the node, open **IBM NAS Admin**, then the **Cluster Tools** folder, and click the **Cluster Setup** icon.
3. At the prompt, verify that you have completed the steps that precede this cluster setup step. If you have, click **Continue**.
4. If this is the first node, click **First Node**. If this is the joining node, go to Step 12 on page 40 and continue from there.
5. The Cluster Information panel appears. Enter the data for the following fields (some of this data comes from Table 1 on page 8):
   - Administrator ID and password

     Note: The ID and password are any valid user ID and password with administrator privileges on the domain.
   - Domain name
   - Cluster name
   - Cluster IP address
   - Subnet mask
   - Quorum drive (select from the pulldown menu)

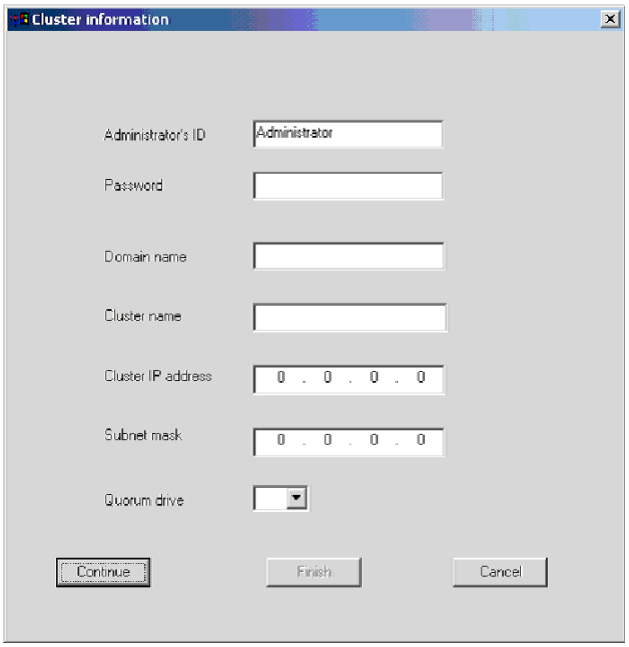

Figure 4. Cluster Information panel

6. After you enter the data, click **Continue**.
7. Verify the information. If it is correct, click **Yes** to start the configuration. Configuration takes a few minutes.
8. If you are prompted to select a user account, enter the user name and password for the domain account that you want the cluster service to use.
9. If you are prompted to select a disk on which to store cluster checkpoint and log files, do the following:
   a. Select the disk on which the Quorum is located (for instance, *G*, if this is what you specified earlier) and click **Next**.
   b. Click **Finish** at the Cluster Information panel.
10. Cluster configuration completes for the first node.
11. Power on the joining node. (You will join this node to the cluster.)
12. In the Cluster Setup wizard, click **Joining Node**.
13. In the First Node Information panel, enter the name of the ***first*** node.
14. At the prompt, specify the domain.
15. If prompted to confirm the Administrator name and password, enter that information and click **Finish**.

You will see a message that configuration takes a few minutes. When configuration completes, the Cluster Administration function starts.

Go to “Verifying network connectivity and names resolution” on page 35 and complete the procedure to verify network connectivity and names resolution.

You have now completed cluster setup.

Configuring clusters

This section contains procedures to assist you in configuring basic cluster functions. It is assumed that the cluster installation procedures in “Cluster setup” on page 39 have completed without errors, and both cluster nodes are running.

It is recommended that you review the *Cluster Administration Guide*, located in the IBM NAS Admin in the Cluster Tools folder, before continuing with the following steps.

Configuring cluster state and properties

You must complete the following steps on the first node to reset the size of the logfile and set the priority and purpose of the private network.

1. Select **Cluster Administration**, located in IBM NAS Admin, in the Cluster Tools folder.

   If prompted for a cluster name, enter the name of the cluster, and then click **Open**.
2. The cluster name appears in the left panel. Click the cluster name to see the status of the cluster nodes in the right pane. The state of both nodes should be “Up”.
3. Right-click the cluster name and select **Properties**.
   a. Select **Quorum Disk**, and change the *Reset quorum log at:* field from 64 KB to 4096 KB.
   b. Select **Network Priority** to view all networks acknowledged by the cluster server, and then select the private network connection and move it to the top for cluster communication priority by clicking **Move Up**.

      This provides internal communication to the private network before attempts are made to communicate over any public networks that are installed. Do not change the communication options for the public network adapters as they should support both network and cluster traffic.
4. Open the properties for the private network and select **Internal cluster communication only (private network)** to ensure that no client traffic will be placed on the private network.
5. Click **Apply**, **OK**, and then **OK**.

Setting up cluster resource balancing

When you configure cluster resources, you should manually balance them on the disk groups to distribute the cluster resource functions between the two nodes. This allows for a more efficient response time for the clients and users accessing these resources.

To set up cluster resource balancing:

1. Select a disk group and bring up its Properties panel by right-clicking it.
2. Click the **General** tab.
3. Click the **Modify** button to the right of the *Preferred owners:* field.
4. In the Available nodes pane, select a node and click the → button to move the node to the Preferred Owners pane.
5. Complete Steps 1 through 4 for each disk group.

Each disk group has a preferred owner so that, when both nodes are running, all resources contained within each disk group have a node defined as the owner of

those resources. Even though a disk group has a preferred owner, its resources can run on the other node in the cluster following a failover. If you restart a cluster node, resources that are preferentially owned by the restarted node switch to the standby system when the cluster service detects that the node is operational, and provided that the defined failover policy allows this to occur. If you have not defined the node as the preferred owner for the resources, then they do not switch to the standby system.

Note: You must reboot before you can see changes made to the cluster resource balancing.

Setting up failover

The failover of resources under a disk group on a node enables users to continue accessing the resources if the node goes down. Individual resources contained in a group cannot be moved to the other node; rather, the group it is contained in is moved. If a disk group contains a large number of resources and any one of those resources fails, then the whole group will perform a failover operation according to the group's failover policy.

The setup of the failover policies is critical to data availability.

To set up the failover function:

1. Open the Properties panel for the disk group.
2. Select the **Failover** tab to set the Threshold for Disk Group Failure.

   For example, if a network name fails, clustering services attempts to perform a failover operation for the group 10 times within six hours, but if the resource fails an eleventh time, the resource remains in a failed state and administrator action is required to correct the failure.
3. Select the **Failback** tab to allow, or prevent, failback of the disk group to the preferred owner, if defined.

In allowing failback of groups, there is a slight delay in the resources moving from one node to the other. The group can also be instructed to allow failback when the preferred node becomes available or to perform a failover operation during specific off-peak usage hours.

Each resource under each disk group has individual resource properties. The properties range from restart properties, polling intervals to check if a resource is operational, to a timeout to return to an online state. The default settings for these properties are selected from average conditions and moderate daily use.

Creating users

The creation of users is performed through normal procedures. Users do not need to be created exclusively for use on the cluster resources. You must define properties of the resources for users to access the resources within the domain policies. All user-accessible cluster resources have the same properties as standard Microsoft Windows resources, and should be set up following the same policies.

Note: If your storage will be accessed by UNIX or UNIX-based clients and servers, continue with "Defining UNIX users and groups" on page 44. The NAS Gateway 300 is on a Windows domain and inherits those Windows users, eliminating the need to define local Windows users and groups. Also, shares are created in the clustering setup.

- If Windows clients and servers will access your storage, follow the steps in "Defining Windows users and groups".
- If UNIX and UNIX-based clients and servers will access your storage, follow the steps in "Defining UNIX users and groups" on page 44.
- If both Windows and UNIX clients and servers will access your storage, follow the steps in "Defining Windows users and groups" and then follow the steps in "Defining UNIX users and groups" on page 44.

Defining Windows users and groups

This section describes how to set up Windows users and groups that will access the NAS Gateway 300 storage.

You can define new local users and groups on the NAS Gateway 300 and also allow existing users and groups to access the NAS Gateway 300 storage. You can also add the NAS Gateway 300 to an existing Windows domain that is controlled by a PDC and define new users and groups on the PDC who can access the NAS Gateway 300.

If you are defining local Windows users and groups, follow the steps in "Defining local Windows users and groups". If you are giving access to the NAS Gateway 300 storage to users and groups in an existing Windows domain, follow the steps in "Giving storage access to Windows domain users and groups" on page 44.

Defining local Windows users and groups: If you are defining local Windows users and groups, you can use the Windows 2000 for NAS user interface. In the Users task group, you create and manage local users and groups on the NAS Gateway 300. To go to the users page, click **Users**. From this page you can create, edit, and delete local users and groups on the NAS Gateway 300 by clicking either **Local Users** or **Local Groups**.

To create new local users:

1. Click **Local Users**.
2. Click **New...**.
3. Type user name, password, and description (optional).
4. Click **OK**. The new user name should appear in the list of user names.
5. Repeat Steps 1 through 4 for each new local user that you want to add.
6. When you finish adding new users, click **Back** to return to the Users and Groups page.

To create new local groups:

1. Click **Local Groups**.
2. Click **New...**.
3. Type group name and description (optional).
4. Click **Members**.
5. For each user that you want to add to the group, select the user name from the list of users, and then click **Add**.
6. Click **OK**. The new group name should appear in the list of group names.
7. Repeat Steps 1 through 6 for each new local group that you want to add. If your storage is also going to be accessed by UNIX or UNIX-based clients and servers, continue with "Defining UNIX users and groups" on page 44. Otherwise, continue with "Creating shares" on page 49.

Giving storage access to Windows domain users and groups: You must first join the NAS Gateway 300 to the Windows domain. You can use the Windows 2000 for NAS user interface to do this. Start the Windows 2000 for NAS user interface, and then do the following:

1. Click **Network**.
2. Click **Identification**.
3. Select the radio button labeled **Domain**, and specify the name of the domain being joined.
4. Specify a user name and password that can be used to log on to the domain.
5. Click **OK**.
6. Shut down and restart the NAS Gateway 300.

Users and groups already defined in the domain can now be given access to any file shares that you create on the NAS Gateway 300. If you need to add new users and groups to the domain, consult the online documentation on the PDC for information on performing this procedure, or if you are not the administrator of the domain (PDC), contact the domain administrator to have the users and groups defined.

If your storage is also going to be accessed by UNIX or UNIX-based clients and servers, continue with "Defining UNIX users and groups". Otherwise, continue with "Creating shares" on page 49.

Defining UNIX users and groups

This section describes how to set up UNIX users and groups to access the NAS Gateway 300 storage using the Network File System (NFS) protocol.

Support for NFS is provided in the NAS Gateway 300 by a preloaded and preconfigured software component, Microsoft Services for UNIX. The levels of NFS supported by Services for UNIX, and in turn the NAS Gateway 300, are NFS Versions 2 and 3. Any client or server that is using an NFS software stack supporting NFS Version 2 or NFS Version 3, regardless of the operating system, should be able to connect to the NAS Gateway 300 and access its storage through NFS.

You administer NFS file shares and other attributes with standard Windows administration tools, including those provided as part of the IBM NAS desktop and the Microsoft Windows 2000 for NAS user interface. Additional configuration of the User Name Mapping component of Services for UNIX, which maps the UNIX user name space to the Windows user name space, is required to support NFS security.

Consult the online documentation for Services for UNIX for more information on configuring User Name Mapping. To view the online documentation for Services for UNIX on the NAS Gateway 300:

1. From the NAS Gateway 300 desktop, click the **IBM NAS Admin** icon.
2. On the left pane of the IBM NAS Admin console, expand **File Systems**.
3. Expand **Services for UNIX**.
4. Select any of the items that appear under Services for UNIX.
5. Click anywhere on the right pane of the IBM NAS Admin console, and then press the F1 key to bring up the online documentation for Services for UNIX in a separate window.

You can define a local UNIX name space on the NAS Gateway 300 by configuring the Server for PCNFS component of Services for UNIX. Alternately, you can point Services for UNIX to an existing Network Information Service (NIS) domain that defines the UNIX name space. In both cases, you must configure the User Name Mapping component to map the UNIX name space that you select to the Windows name space, because file shares and individual file and directory permissions on the NAS Gateway 300 are defined in the context of the Windows name space.

To define a local UNIX name space, continue with "Using a local UNIX name space". To use a UNIX name space defined on a NIS domain, continue with "Using the UNIX name space on an NIS domain" on page 47.

Using a local UNIX name space: This procedure should be performed only once. You might have to add more groups and users in the **Server for PCNFS** page if you add more users and groups to your UNIX environment and NAS Gateway 300 or Windows domain at a later time.

1. Open the IBM NAS Administration console by double-clicking the **IBM NAS Admin** icon on the NAS desktop.
2. In the left pane, select **File Systems**; then select **Services for UNIX**.
3. In the left pane, click **Server for NFS**.
4. In the right pane, in the Computer name: field, type `localhost`.
5. In the left pane, click **Server for PCNFS**.
6. In the right pane, click **Groups**.
7. On the Groups page, you must add the groups from your UNIX host to which all of your UNIX users belong. You need to know both the group name and the group ID (GID) number. This information can be found in the /etc/group file on most UNIX systems, or can be copied to the c:\winnt\system32\drivers\etc directory.

   As an example, on an AIX system, in the following line from an /etc/group file, the fields are separated by a colon (:). The first field ("staff") is the group name; the third column ("1") is the GID:

   ```
   staff:!:1:pemodem,ipsec,netinst,protcs
   ```

   To add a group, type the group name and GID number in the Group name and Group number (GID) fields, and then click **New**.
8. When you finish adding groups, click **Apply**.
9. Click **Users**.
10. On the Users page, you can add all of the UNIX users who will be accessing and storing files on the NAS Gateway 300 through an NFS share. For each user you will need to know the Windows user name, the UNIX user name, the primary group, and the user ID (UID) number. This information can be found in the /etc/passwd and /etc/group files on most UNIX systems or these files can be copied to the c:\winnt\system32\drivers\etc directory.

    As an example, on an AIX system, in the following line from an /etc/passwd file, the fields are separated by a colon (:). The first field ("user1") is the user name; the third field ("3135") is the UID, and the fourth field ("1") is the GID of the user's primary group. This will correspond to a line in the /etc/group file, where you can find the primary group name corresponding to the GID.

    ```
    user1:!:3135:1:User 1:/home/user1:/bin/ksh
    ```

    To add a user, click **New**, type the required information, and then click **OK**.

Services for UNIX supports a limited syntax in the *passwd* file. In particular, it seems to work best when the second field of each line—the password field—is filled in with a random 13-character string. This need not have anything to do with the user's password, so a string such as *0123456789012* is acceptable. Some UNIX systems use shadow passwords and fill in this field with a meaningless token value such as *!* or *x*, and you will need to change this.

11. When you finish adding users, click **Apply**.
12. In the left pane, click **User Name Mapping**.
13. In the right pane, select Personal Computer Network File System (PCNFS).
14. In the **Password file path and name** field, type
    `c:\winnt\system32\drivers\etc\passwd`
15. In the **Group file path and name** field, type
    `c:\winnt\system32\drivers\etc\group`
16. Next, delete all special users and groups, leaving just the actual users and groups that will be used in accessing NFS resources. An example of a special user is *root*, usually, and UID numbers from 0 to 99 are generally reserved for system accounts and should not be mapped.
17. Click **Apply**.
18. Click **Maps**.

    On the Maps page, you can configure simple maps or advanced maps. Configure simple maps if the Windows user name and UNIX user name is the same for each UNIX user to be mapped, and the Windows group name and UNIX group name is the same for each UNIX group to be mapped. Otherwise, you should configure advanced maps.
19. To configure simple maps, select the **Simple maps** check box and continue with Step 20.

    To configure advanced maps, clear the **Simple maps** check box and continue with Step 21.
20. Under Simple maps, select the Windows domain name from the drop-down list, and then continue with Step 22 on page 47. (If your Windows users are defined locally on the NAS Gateway 300, select the entry containing the computer name of the NAS Gateway 300, preceded by two backslash characters ("\\"). Otherwise, select the name of the Windows domain where the users are defined from the list.)
21. Under Advanced maps, perform the following steps.
    a. Define user mappings:
        1) Click **Show user maps**.
        2) Select the Windows domain name from the drop-down list. (If your Windows users are defined locally on the NAS Gateway 300, select the entry containing the computer name of the NAS Gateway 300, preceded by two backslash characters ("\\"). Otherwise, select the name of the Windows domain where the users are defined from the list.)
        3) Click **Show Windows Users** to display all of the Windows user names in the Windows domain that you selected.
        4) Click **Show UNIX Users** to display all of the UNIX user names in the NIS domain that you selected.
        5) Type a Windows user name, or select one from the list of Windows user names.
        6) Type a UNIX user name to be mapped to the Windows user name you specified, or select one from the list of UNIX user names.
        7) Click **Add** to add the mapping between the UNIX user name and Windows user name to the list of maps.

8) If multiple Windows user names are mapped to one UNIX user name, select one Windows user name to be the primary user name. Select the mapping corresponding to the primary user name from the list of maps, and then click **Set Primary**.

b. Define group mappings:
   1) Click **Show group maps**.
   2) Select the Windows domain name from the drop-down list. (If your Windows users are defined locally on the NAS Gateway 300, select the entry containing the computer name of the NAS Gateway 300, preceded by two backslash characters ("\\"). Otherwise, select the name of the Windows domain where the users are defined from the list.)
   3) Click **Show Windows Groups** to display all of the Windows group names in the Windows domain you selected.
   4) Click **Show UNIX Groups** to display all of the UNIX group names in the NIS domain you selected.
   5) Type a Windows group name, or select one from the list of Windows group names.
   6) Type a UNIX group name to be mapped to the Windows group name that you specified, or select one from the list of UNIX group names.
   7) Click **Add** to add the mapping between the UNIX group name and Windows group name to the list of maps.
   8) If multiple Windows group names are mapped to one UNIX group name, you must select one Windows group name to be the primary group name. Select the mapping corresponding to the primary group name from the list of maps, and then click **Set Primary**.

22. Click **Apply**.

User Name Mapping rereads its enumeration source on a schedule. By default, this occurs once a day. You can reset the refresh period. To force User Name Mapping to reread the enumeration source, you can click **Synchronize Now** on the Configuration panel.

Note: If maps do not seem to synchronize, you might need to stop and restart User Name Mapping. You can do this through the GUI, or by the commands:

```
net stop mapsvc

net start mapsvc
```

You can now continue with "Creating shares" on page 49.

Using the UNIX name space on an NIS domain: The following procedure applies whether your NIS server is UNIX-based or Windows-based (implemented as a Windows domain controller running a Microsoft Server for NIS).

1. To open the IBM NAS Administration console, double-click the **IBM NAS Admin** icon on the NAS desktop.
2. In the left pane, expand File Systems; then expand Services for UNIX.
3. In the left pane, click **Server for NFS**.
4. In the right pane, in the Computer name: field, type `localhost`
5. In the left pane, click **User Name Mapping**.
6. In the right pane, select Network Information Services (NIS); then click **Maps**.

   On the Maps page, you can configure simple maps or advanced maps. Configure simple maps if the Windows user name and UNIX user name is the

same for each UNIX user to be mapped, and the Windows group name and UNIX group name is the same for each UNIX group to be mapped. Otherwise, you should configure advanced maps.

7. To configure simple maps, select the **Simple maps** check box and continue with Step 8.

   To configure advanced maps, clear the **Simple maps** check box and continue with Step 9.

8. Under Simple maps, perform the following steps:
   a. Select the Windows domain name from the drop-down list. (If your Windows users are defined locally on the NAS Gateway 300, select the entry containing the computer name of the NAS Gateway 300, preceded by two backslash characters ("\\"). Otherwise, select the name of the Windows domain where the users are defined from the list.)
   b. In the NIS domain box, type the NIS domain name. You can also type the name of a specific NIS server in the NIS server box.
   c. Continue with Step 10 on page 49.

9. Under Advanced maps, perform the following steps:
   a. Define user mappings as follows:
      1) Click **Show user maps**.
      2) Select the Windows domain name from the drop-down list. (If your Windows users are defined locally on the NAS Gateway 300, select the entry containing the computer name of the NAS Gateway 300, preceded by two backslash characters ("\\"). Otherwise, select the name of the Windows domain where the users are defined from the list.)
      3) In the NIS domain field, type the NIS domain name. You can also type the name of a specific NIS server in the NIS server field.
      4) Click **Show Windows Users** to display all of the Windows user names in the Windows domain you selected.
      5) Click **Show UNIX Users** to display all of the UNIX user names in the NIS domain you selected.
      6) Select a Windows user name from the list of Windows user names.
      7) Select a UNIX user name to be mapped to the Windows user name that you specified.
      8) Click **Add** to add the mapping between the UNIX user name and Windows user name to the list of maps.
      9) If multiple Windows user names are mapped to one UNIX user name, you must select one Windows user name to be the primary user name. Select the mapping corresponding to the primary user name from the list of maps, and then click **Set Primary**.
   b. Define group mappings as follows:
      1) Click **Show group maps**.
      2) Select the Windows domain name from the drop-down list. (If your Windows users are defined locally on the NAS Gateway 300, select the entry containing the computer name of the NAS Gateway 300, preceded by two backslash characters ("\\"). Otherwise, select the name of the Windows domain where the users are defined from the list.)
      3) In the NIS domain field, type the NIS domain name. You can also type the name of a specific NIS server in the NIS server field.
      4) Click **Show Windows Groups** to display all of the Windows group names in the Windows domain that you selected.
      5) Click **Show UNIX Groups** to display all of the UNIX group names in the NIS domain that you selected.
      6) Select a Windows group name from the list of Windows group names.

7) Select a UNIX group name to be mapped to the Windows group name that you specified.
8) Click **Add** to add the mapping between the UNIX group name and Windows group name to the list of maps.
9) If multiple Windows group names are mapped to one UNIX group name, you must select one Windows group name to be the primary group name. Select the mapping corresponding to the primary group name from the list of maps, and then click **Set Primary**.

10. Click **Apply**.

You can now continue with “Creating shares”.

Creating shares

To create new file shares on the NAS Gateway 300, do the following:

1. Start the Windows 2000 for NAS user interface.
2. Click the **Shares** tab.
3. Click the **Shares** task.
4. Click **New...**.
5. Specify the share name (the name that clients and servers will use to access the share).
6. Specify the share path and select the **Create folder if it does not already exist** check box.
7. By default, the Microsoft Windows (CIFS) and UNIX (NFS) check boxes are selected (enabled). If this share is not to be accessed by Windows clients and servers, clear (disable) the Microsoft Windows (CIFS) check box. If this share is not to be accessed by UNIX clients and servers, clear the UNIX (NFS) check box.
8. If this share is to be accessed by:
   - Windows clients and servers, then click **CIFS Sharing** and specify the access permissions that you want. (Note that, by default, every user has full access to all files and directories under the shared folder.)
   - UNIX clients and servers, then click **NFS Sharing** and specify the access permissions that you want. (Note that by default, every user has full access to all files and directories under the shared folder.)
9. Click **OK**. The new share should appear in the list of shares.
10. Repeat Steps 4 through 9 for each additional share that you want to create.

A note on anonymous access: It is strongly recommended that you not disable anonymous access. If a client presents a UID that is not recognized, Server for NFS can still grant that client a very limited form of access as a special *nobody* user. This is known as anonymous access, and you can enable or disable on a per-share basis. This anonymous user will have very limited access to resources on the NAS: it has only the permissions that are granted to the *Everybody* group in Windows, which corresponds to the *other* (or *world*) bits in a POSIX permissions mode.

Allowing anonymous access is not a security risk, so disabling it might provide a false sense of security. (The real security risk is to grant everyone access to resources that should be protected.) And disabling anonymous access has one severe consequence: it is so unexpected by NFS clients that they might not be able to connect as NFS V3 clients at all, and might instead downgrade the connection to use the NFS V2 protocol.

Creating clustered file shares (CIFS and NFS)

Note: For HTTP and FTP clustering setup and file sharing, refer to the information at the following URL:

http://support.microsoft.com/default.aspx?scid=kb;EN-US;q248025

The creation of file shares on a cluster involves dependencies on a physical disk, a static IP address, and a network name. These dependencies allow resources that are defined to the same disk group to move as a group. The dependencies also assure necessary access for the given resource.

Note: You must configure Server for NFS before NFS file sharing can be used. See "Enabling Server for NFS" on page 52 for details.

Figure 5 illustrates the file share dependencies. Descriptions of the diagram components follow the figure.

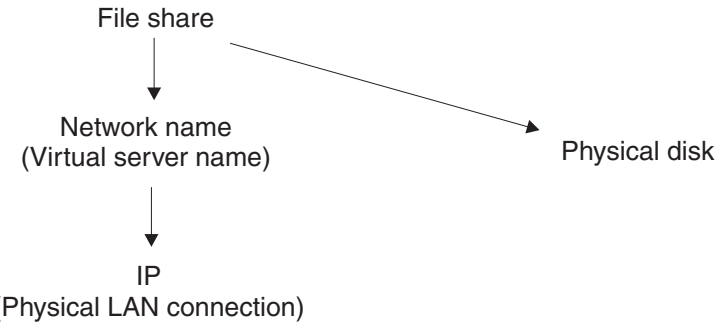

Figure 5. File share dependencies

Physical disk
The base resource in which to store user data. It is not dependent on any other resources except the physical disk that it defines. The disk resource must also have the same drive letters on both nodes so that the definitions of resources that depend on it do not change if the resource is moved to the other node.

Static IP address
A virtual address that binds onto an existing IP address on one of the cluster's public networks. This IP address provides access for clients and is not dependent on a particular node, instead a subnet that both nodes can access. Because this address is not the physical adapter's permanent address, it can bind and unbind to its paired adapter on the same network on the other node in the cluster. You can create multiple IP addresses using the Cluster Administrator on the same physical network. A unique static IP address is required for each virtual server.

Note: The cluster IP address should not be used for file shares. That address is reserved to connect to and manage the cluster through the network that it is defined on.

Network name
An alternate computer name for an existing named computer. It is physically dependent on an IP address of one of the public networks. When a disk group contains an IP address resource and a network name, it is a virtual

server and provides identity to the group, which is not associated with a specific node and can be failed over to another node in the cluster. Users access the groups using this virtual server. A virtual server can have multiple file shares.

In the creation of a basic file share that is publicized to the network under a single name, you must set it up to be dependent on the physical disk and network name in the same disk group you are creating the file share in. The network name is dependent on the IP address, so do not add that to the dependency list. You can set the share permissions and advanced share resources also.

Users will access the cluster resources using \\<*network_name*>\<*fileshare_name*>.

Clustered file share creation example

An example of how to create a clustered file share follows. For this example, assume that you are creating a file share in Disk Group 2.

1. Create the IP address resource:
   a. Right-click Disk Group 2, and select **New** → **Resource**.
   b. Enter an IP address name, for example *ipaddr2*, and change the resource type to IP Address.
   c. Select **Run this resource in a separate Resource Monitor** and click **Next**.
   d. A list of possible owners appears, and both nodes should remain as assigned. Click **Next**.
   e. There are no resource dependencies on this panel, so click **Next**.
   f. Enter your TCP/IP parameters. This will be the first virtual IP address. The value in the Network field identifies to the system the network on which the address is located. Click **Finish** to create the resource.
   g. Right-click the resource and select **Bring online**.
2. Create the network-name resource:
   a. Right-click Disk Group 2, and select **New** → **Resource**.
   b. Enter the virtual server name to use (for example, *NN2*), select **Network Name** as the resource type, and click **Next**.
   c. Both nodes are possible owners. Click **Next**.
   d. Add the IP address you created as a resource dependency in Step 1 and click **Next**.
   e. Type the virtual server name, NN2, into the Network Name Parameters field and click **Finish**.
   f. It takes a few moments to register the virtual server name with your name server. After this completes, bring the resource online.
3. Create the CIFS or NFS file share resource:
   a. Right-click Disk Group 2 and select **New** → **Resource**.
   b. Enter a file share name (for example, *FS2*) and select either **File Share** or **NFS Share**.
   c. Both nodes are possible owners. Click **Next**.
   d. Add the resource dependencies for the physical disk and network name that the file share will use and click **Next**.
   e. Enter the share name of FS2 and the path to the disk in this group, either drive or subdirectory. You can then set:
      - For CIFS shares properties:
        – User Limit
        – Permissions

- Advanced File Share
- For NFS shares properties:
  - Permissions
  - Share

    A note on anonymous access: When you create an NFS share, it is strongly recommended that you not disable anonymous access to avoid client-connection problems. See “Enabling Server for NFS” for more details.

f. Click **Finish** to create the resource.

g. Right-click the resource and select **Bring online**.

Enabling Server for NFS

To enable Server for NFS, you need to specify where User Name Mapping is running.

To specify where User Name Mapping is running, follow this path, **Services for UNIX → User Name Mapping**, and then enter the server name that is running User Name Mapping in the Computer Name field. For a cluster, this entry must be the clustered name or IP address, not that of an individual node.

When planning an NFS installation, consider which machines you want to have particular access-levels to NFS shares. Each class of access should be captured by defining a separate client group.

- To define a client group, click **Services for UNIX → Client Groups**, type the group name in the Group Name field, then click **New**.
- To add members to a client group, select a group name from the current groups list; then click **Advanced** and type the name of a client (a valid computer name).
- **A note on anonymous access:** It is strongly recommended that you not disable anonymous access. If a client presents a UID that is not recognized, Server for NFS can still grant that client a very limited form of access as a special *nobody* user. This is known as anonymous access, and you can enable or disable on a per-share basis. This anonymous user will have very limited access to resources on the NAS: it has only the permissions that are granted to the *Everybody* group in Windows, which corresponds to the *other* (or *world*) bits in a POSIX permissions mode.

  Allowing anonymous access is not a security risk, so disabling it might provide a false sense of security. (The real security risk is to grant everyone access to resources that should be protected.) And disabling anonymous access has one severe consequence: it is so unexpected by NFS clients that they might not be able to connect as NFS V3 clients at all, and might instead downgrade the connection to use the NFS V2 protocol.

Recovering from a corrupted Quorum drive

Attention: Restoring a Quorum rolls the cluster back in time to the backup date. There are impacts to performing this operation that include loss of data. You should perform this operation only when it is absolutely necessary.

Clustering relies on data stored on the Quorum disk to maintain resource synchronization between the two nodes in the cluster. In the event of a power loss to both nodes or a hardware failure that corrupts the Quorum data, the cluster service might not start, leading to the following event log error:

```
Event ID: 1147
Source: ClusSvc
Description: The Microsoft Clustering Service encountered a fatal error.
```

The Quorum drive data must be available so that the cluster service can confirm that the cluster configuration on the local node is up to date. If it cannot read the log, the cluster service does not start to prevent the loading of old configuration data.

To restore the Quorum disk, a Microsoft Windows Backup utility backup of the system state of the boot drive (C:) of one node must be available. Backing up the entire boot drive also saves the system state. Backing up the system state automatically saves the Quorum log and other cluster files.

A Microsoft tool is needed as part of the Quorum restore procedure. This tool is called Clusrest.exe and can be downloaded from the Microsoft Web site at the following URL:

http://download.microsoft.com/download/win2000platform/clusrest/1.0/NT5/EN-US/clusrest.exe

The Quorum restore procedure involves restoring the system state and cluster state to the node followed by execution of the Clusrest.exe tool. Upon completion of the restore, the node should rejoin the cluster and return to normal operation.

1. Restore the entire boot drive of the node if needed. Otherwise, restore the system state to the node.
2. Ensure that the cluster service is stopped on the other node.
3. Restore the Quorum/cluster information to that node by selecting to restore at least the system state. This creates a temporary folder under the Winnt\Cluster folder called Cluster_backup.
4. Run the Clusrest.exe tool to rebuild the Quorum drive. The tool moves the cluster information from the node's boot drive to the Quorum drive.
5. After you complete the process and the cluster service has started successfully on the newly restored node, restart the cluster service on the other node.

Note: If you do not follow this process, and another node with a more current database takes ownership of the Quorum before you update the database from the restored node, the restore does not work.

Before you add software

You have now completed networking and clustering setup and administration, and the NAS Gateway 300 is at a point where you can install software on it. But before you do, it is recommended that you take advantage of the Persistent Storage Manager (PSM) disaster recovery function, detailed in "Disaster Recovery" on page 66.

The PSM disaster recovery function enables you to restore the system drive from a single image, without having to go through the entire recovery procedure and then additionally having to restore a system drive backup. So, if any software you install creates unresolvable problems for your system, you can regain the stable system you had before you installed the software.

Chapter 6. Managing and protecting the network and storage

This chapter describes the additional administrative functions that you can use to manage and protect the network and storage on the NAS Gateway 300.

The following functions are available:

- "IBM Director", accessed through **Start → Programs**
- "NAS Backup Assistant" on page 60, accessed through IBM NAS Admin
- "Persistent Images" on page 62, accessed through the Windows 2000 for Network Attached Storage user interface
- "Tivoli SANergy" on page 75

IBM Director

Note: This section presents an overview of IBM Director functions. For more detailed information, consult the *IBM Director User's Guide* on the Documentation CD-ROM.

IBM Director is a systems-management solution that helps administrators manage single or large groups of IBM and non-IBM devices, NAS appliances, and workstations.

All of the functionality of IBM Director is contained in a simple GUI that enables single-click and drag-and-drop commands. IBM Director can manage up to 5,000 clients depending on configuration density. Powerful remote management functions include:

- Sophisticated discovery of network components
- Scheduled asset (hardware and software) inventories with persistent storage of data
- Proactive problem notification and tools for problem resolution
- Hardware system component monitors and thresholds to trigger alerts of impending problems
- Alert management with automated actions, manual intervention, or both
- Process scheduling to automate wide-scale client software maintenance (clean up temp files, restart tasks, backups, and so on) according to any timetable
- Help desk and routine maintenance functions such as remote control and file transfer
- Extensive security and authentication

IBM Director consists of three main components:

- Management Server
- Agent
- Console

The **Management Server** is a centralized systems manager and is the core of the IBM Director product. Management data, the server engine, and the management application logic reside there. Install the IBM Director Management Server on a dedicated server that has high-availability features. When installed on a Windows 2000 server or Windows NT 4.0 server system in the managed environment, the Management Server provides the management application logic and persistent data

storage of management information using an SQL database. The Management Server maintains a database of all Director Agents and their inventory. All alerts from the agents flow to the management server, which also acts as a central point of configuration for Event Action Plans and System Tasks.

The **Agent** resides on the NAS Appliances and other systems that IBM Director manages. IBM Director recognizes two types of managed systems: native agents (IBM Director Agent installed) and nonnative agents (SNMP agent installed). The Agent comes preinstalled on all IBM NAS appliances. It runs as a service that is automatically started at boot time. IBM Director Agent provides valuable information to IBM Director management server and other supported management applications. In addition to its native interface with the Director Management Console, it provides point-to-point remote management of client systems through a Web browser window.

You perform administrative tasks at the **Console**. It is a Java application that serves as the user interface to the Director-managed environment. The console provides comprehensive hardware management using a single click or drag-and-drop operation. You can install the console on a machine at a remote location from the server. Consoles are not licensed, so you can distribute them freely among unlimited number of machines. In addition, there is no limit to the number of IBM Director Consoles that can connect into the Management Server.

Dependencies

The IBM Director 3.1 Agent (the version included in this release) must be managed by an IBM Director 3.1 Management Server. If your Management Server is running an earlier version of IBM Director (V2.2 or earlier), then you must upgrade it to ensure proper operation. This includes Director Consoles as well. The IBM Director 3.1 Management Server contains an Agent software distribution package that you can use to upgrade pre-version 3.1 Agents. This allows easy and automated upgrading of the entire system to version 3.1. You can check the version of IBM Director Agent running on a NAS appliance by issuing: http://<system_name>:411/ on a local Web browser.

Hardware requirements

It is highly recommended that you install the IBM Director Server on a server separate from the IBM NAS appliance. The IBM Director Server running on an IBM NAS appliance will significantly reduce its performance. The server must meet these minimum requirements:

Hardware vendor	Must be IBM. The management tools of IBM Director and Director Extensions require IBM equipment.
CPU	A 733 MHz PIII processor is recommended. Standard PII processors can be functional, but these processors might not be sufficient during heavy usage.
Memory	512 MB RAM is recommended. During idle times, while using the standard JET database, the Management Console can consume 300+ MB RAM. The number of managed agents, active consoles, and amount of alerts being processed increases the amount of memory needed.

Disk	Because the Management Server software requires only 250 MB, and the JET database has a maximum size of 1 GB, 9 GB of disk space is sufficient. Use a 4 GB partition for the operating system (including the swap file).

All IBM NAS products exceed the minimum hardware requirements for operating an IBM Director Agent.

Director extensions

A portfolio of advanced management tools for IBM-specific hardware is provided by IBM Director as a set of optional enhancements. These tools integrate into IBM Director and provide management capabilities from a single console with a consistent look and feel. These extensions are provided as part of the preinstalled IBM Director Agent on the IBM NAS appliances:

- Management Processor
- Assistant Capacity Manager
- Cluster Systems Management
- Rack Manager
- Software Rejuvenation
- Systems Availability

To use these extensions, you must load them on the IBM Director Management Server during installation.

Naming conventions

All IBM Director Agents have a Director system name that it is known by the Management Server and Consoles. This Director System Name is defaulted to the computer name during the NAS appliance preinstallation process[1]. The Director system name does not have to be the same as the computer name. The Director system name is displayed on the IBM Director Console to identify the NAS Appliance under the Group Contents column. You can optionally change the Director System Name on an agent using the following procedure:

1. Open a command prompt window and enter the following IBM Director Agent command to open the GUI interface:

   ```
   twgipccf.exe
   ```

2. Type the new Director System Name and click **OK**.

The change takes place immediately.

Note: You might need to delete the NAS appliance from the Group Contents and have it rediscover the appliance by its new name.

Web-based access

IBM Director Agent uses an Apache Web Server for Web-based access. All traffic, even logon, is certificate-based encrypted. The Web server requires two ports. One port (411) accepts non-SSL HTTP requests and automatically redirects to the second port (423), which handles SSL requests.

---

1. Although you can do so, it is recommended that you not change the default computer name to avoid the chance of propagating misidentification through the system. And, if you are using IBM Director to manage your appliance, and you change the default name, the default name continues to appear in IBM Director.

Disaster recovery

It is important to provide adequate backup for key IBM Director Management Server files for restoration purposes. It is recommended that you regularly back up the IBM Director Management Server so that you can recover it in the event of a server disaster. You need to save customizations that you make to the IBM Director, including event action-plans, schedules, thresholds, and so on. Several commands are provided with IBM Director to accomplish this task:

twgsave

This command saves the complete settings to a directory named *Director.save.#*, where *#* shows the number of backups (for example, the third backup of the server will be saved in directory *Director.save.3*). You must stop the IBM Director Management Server service to execute this command. The command supports the following options:

```
twgsave -s
```

where the optional parameter *-s* specifies that software distribution packages not be saved. This helps reduce the size of the backup files.

twgrestore

This command restores the saved data from an IBM Director Management Server. Do not attempt to use this restore feature to replicate an IBM Director Server. The command supports the following options:

```
twgrestore -t directory
```

where the optional parameter *-t* specifies that the data is restored, but server ID and system name is not restored, and *directory* is where the saved data resides. IBM Director Management Server cannot be running when this command is issued.

twgreset

This command resets the Director Server system to the status after installing. You can use it if you want to clear all tables in the database and erase the system ID files. This command can be helpful to make sure that after a restore only the data from the saved directory will be in the Director System. The command supports the following options:

```
twgreset -d -i
```

Where *-d* means to clear the tables in the database, and *-i* means to erase the unique identification files for the system. You can save and restore data only when the Director Support Program and service are stopped. Agents running on IBM NAS appliances do not need to be explicitly backed up because the NAS recovery CD-ROM provides this feature. Applying the Recovery CD-ROM will reinstall the IBM Director Agent.

Software distribution

The Software Distribution task enables you to import and silently distribute predefined software distribution packages to an IBM Director Client system. These packages are prepared by IBM for IBM NAS products and include software fixes and release updates only. This includes upgrading the IBM Director client itself.

The basic delivery is a single file package that is signed with a unique IBM NAS key. Only IBM can create the signed packages that can be used by the IBM Director Software Distribution tool.

Software distribution using IBM Director can be deployed to a single IBM Director client, all IBM Director clients, or some combination in between. The administrator has complete control over which IBM Director clients receive any given package. By default, software distribution packages automatically install themselves immediately following delivery to the IBM client. Delivery of the package can be done manually or scheduled for a later, more convenient time.

Rack Manager and inventory enhancements

The Rack Manager task has been updated to include all of the IBM NAS components. A new component category, **NAS**, includes all of the IBM NAS appliance engines. All IBM NAS appliances are automatically discovered by the Rack Manager task for drag-and-drop rack construction. This enhancement is part of the IBM Director Server Service Pack 3.1.1; the service pack must be loaded on the IBM Director server before you can take advantage of this new category. The following component categories have been updated to include the new IBM NAS appliance components:

Racks Includes the new component, NAS Rack Model 36U

Storage

Includes these new components:

- NAS Storage Expansion Unit Model 0RU
- NAS Storage Expansion Unit Model 1RU

Fibre Channel

Includes these new components:

- NAS 8-port Fibre Channel Hub Model 1RU
- NAS Raid Storage Controller Model EXP
- NAS Raid Storage Controller Model 0RU
- NAS Raid Storage Controller Model 2RU
- NAS Raid Storage Controller Model EXU

NAS Is a new component category that includes these components:

- NAS 100 Engine Model R12
- NAS 100 Engine Model R18
- NAS 200 Engine Model 200
- NAS 200 Engine Model 201
- NAS 200 Engine Model 225
- NAS 200 Engine Model 226
- NAS 200 Engine Model 25T
- NAS 200i Engine Model 100
- NAS 200i Engine Model 110
- NAS 300 Engine Model 5RZ
- NAS 300 Engine Model 6RZ
- NAS 300G Engine Model 5RY
- NAS 300G Engine Model 6RY
- NAS Gateway 300 Engine Model 7RY

Dynamic NAS groups

Dynamic NAS groups are an IBM Director Management Server enhancement made specifically for IBM NAS appliances. You must install this enhancement on the IBM Director Management Server as well as all IBM Director Consoles. You can add dynamic NAS groups to the IBM Director Server and Consoles by downloading the InstallShield extension from the IBM Web site and invoking the executable file. This will create a new Group on all consoles that represent IBM NAS appliances in the managed network.

Dynamic groups are automatically populated and maintained based on queries to the database. These dynamic NAS groups must be added after the IBM Director Management Server has been installed on a dedicated server. IBM NAS appliances appear under the Groups column in the IBM Director Management Server. The Group Contents column will then contain all the IBM NAS devices that have been discovered on the network.

NAS Web UI task

NAS Web UI is an IBM Director Management Server enhancement made specifically for managed networks containing IBM NAS appliances. Install NAS Web UI on the IBM Director Management Server and all IBM Director Consoles to create a new task called **IBM NAS Appliances** with a subtask named **Launch UI Web**. You can apply this new console task to a NAS machine, causing a Web browser to be automatically launched with a URL pointing to the Web UI on the target NAS machine. The port specified in the URL is port 8099, which invokes Windows 2000 for NAS.

Predictive Failure Analysis

Predictive Failure Analysis (PFA) provides advanced notification of a pending failure so that corrective action can be taken to avoid unplanned downtime. The PFA alerts are sent to IBM Director, where a wide variety of Event Action Plans can be established, such as automatically notifying the administrator through e-mail, or executing tasks in response to the alert. When used in conjunction with the IBM electronic service agent, the PFA alerts are routed to an IBM support person, who responds to the customer about the alert. The alerts can also be forwarded to other management packages.

For more information

For more information on IBM Director, consult its user's manual contained on the Documentation CD-ROM.

NAS Backup Assistant

The NAS Backup Assistant is a preloaded utility that helps you create and schedule backup batch files, and maintain log files. It can be used for backing up either the NAS Gateway 300 operating system or user data.

If you want to back up selected folders, you can use NT Backup without the NAS Backup Assistant (which backs up an entire volume). However, if you use NT Backup, it is recommended that you select and back up the copy of the files in a previous persistent image, rather than the original data itself. When selecting the files for the NT Backup operation, you must select the specific folders in the persistent image. If you select the entire group of persistent images, the files in those images will not be selected for backup. For more information about persistent images see "Persistent Images" on page 62.

Because NAS Backup Assistant only creates and launches scripts, and is not a comprehensive backup application, it does not support interactive error messages. To check status of jobs, you must either view the Backup Logs or view the Windows Event Viewer.

You invoke the NAS Backup Assistant by clicking the **IBM NAS Admin** desktop icon to open the IBM NAS Administration console. Select **Backup and Restore** to expand that tree, then select **IBM NAS Backup Assistant**. When you select this

option, a logon prompt appears. Log on as a user who has backup operator privileges (an administrator or backup administrator). If a logon prompt does not appear, right-click the **IBM NAS Backup Assistant** link, and select refresh. When you log on, the main panel appears.

The four tabs on the main panel are:

Backup Operations
The main window where you create and schedule backup batch jobs.

Two backup methods you can select in the Backup Operations window are the standard NT Backup method and the Persistent Storage Manager (PSM) Persistent Image method. A standard NT Backup operation backs up only those files on the drive that are not in use. To guarantee a complete backup image using this method, you must ensure that no users are accessing any files on the drive, so this method is useful only for offline backup.

To do a complete online backup that includes files that are in use, choose the PSM Persistent Image backup method. This method creates a persistent image (mapped as an unused drive letter on the system), backs up a copy of that persistent image, and then deletes the original persistent image (drive letter). For more information about persistent images, see "Persistent Images" on page 62.

Scheduled Jobs
Displays a list of backup batch jobs that you scheduled.

Backup Logs
Displays a list of log files for each backup that has run.

Displayed Logs
Displays the text contained in the log files that you can select from the Backup Logs tab.

All of the options on each tab are described in detail in the online help. To access the online help:

1. Click the **IBM NAS Admin** icon.
2. Expand the Backup and Restore directory.
3. Select **IBM NAS Backup Assistant Help**.
4. Log in.

Restoring using the NT Backup panel

Note: If you are restoring a backup that you created using Persistent Images in the NAS Backup Assistant, the NT Backup file (*.BKF) was created for the persistent image virtual drive letter instead of the original drive letter. For example, if you selected drive C for backup, a persistent image was created on the next available drive letter in the system, and that drive was backed up instead of drive C. If you do not remember the original drive letter, you can view the backup log files in NAS Backup Assistant. The top section of the log file gives you the original drive letter, and the bottom section gives you the persistent image drive letter. When you have the original drive letter, perform the procedure below.

To restore backups, use the following procedure:

1. Click the **Restore using NT Backup** link in the Backup and Restore section of the IBM NAS Admin console to open the backup GUI.

2. Click **Restore Wizard**; then click **Next**. You are asked what you want to restore.
3. Select the appropriate media that you are restoring from.
4. If you are restoring from tape, expand the backup media pool name, and then double-click the media (this will normally be named *media created on {date - time}*. This action will read the set list from the tape.

   If you are restoring a file, select **Tools** → **Catalog a backup file**, then click **Browse** and find the backup file (.BKF) created for this backup.

   Note: If you do not know the .BKF file name, refer to the backup log in NAS Backup Assistant.
5. Click **OK**. You will now have a *Media created on {date - time}* listed under file.
6. Click the plus sign (**+**) to the left of this media to see the set list. You might be prompted to enter the path to the file that you want to catalog; if so, select the same file that you just imported. This will build a set list.
7. Select the files and directories to restore.
8. Select **Alternate Location** from the **Restore files to:** pull-down.
9. In the alternate location window, select the root directory of the original backup drive letter that you determined (see the note on page 61).
10. To change restore options, select **Tools** from the menu bar at the top of the window, and then select **Options**. Refer to NT Backup online help (see **Restore files from a file or a tape**) for use of these options.
11. After you select the files or directories for restore, the alternate location, and options, click **Start Restore**.
12. At the prompt, confirm that you want to begin the restore. Click **Advanced** to select advanced options (see the NT Backup online help for details); then click **OK** to begin the restore.

Persistent Images

A persistent image is a copy that you make of one or more file system volumes at a specific time. You can use the Persistent Images function to restore a file or volume to the state it was in at the time that you created the persistent image. Persistent images are maintained in a way that minimizes the storage required to keep multiple copies of the volume. This is done by using a copy-on-write technique that uses, for each volume, an area of pre-allocated storage (the PSM cache file) that keeps only those data blocks that have been written since the time you made a persistent image of the volume.

Persistent Storage Manager (PSM) allows you to create and preserve images of the NAS Gateway 300 drives. You can take a persistent image immediately or schedule persistent images as one-time events or regularly repeated events.

You can access the PSM tasks in the Disks/Persistent Storage Manager task group within the Windows 2000 for Network Attached Storage user interface in one of two ways:

- Open the IBM NAS Admin console on the appliance desktop and select Persistent Storage Manager. This automatically launches the Windows 2000 for Network Attached Storage user interface and brings up the Disks/Persistent Storage Manager page containing the PSM tasks.
- Start the Windows 2000 for Network Attached Storage user interface directly.

When you create a persistent image, it appears as a directory on the original drive. Access rights and permissions from the original drive are inherited by the persistent image. Persistent images are used in the same way as conventional drives. However, unlike conventional drives, persistent images are records of the content of the original drive at the time you created the persistent image. Persistent images are retained following shutdown and reboot.

There are six PSM tasks in the Disks/Persistent Storage Manager group:

- Global Settings
- Volume Settings
- Persistent Images
- Schedules
- Restore Persistent Images
- Disaster Recovery

Each of these tasks is described in the following sections. More detailed descriptions and instructions for each of the control panels and topics are covered in the online help.

Global Settings

On this panel, you can configure the persistent image system attributes shown in Table 4.

Table 4. Persistent image global settings

Attribute	Default value
Maximum number of persistent images	250
Inactive period	5 seconds
Inactive period wait timeout	15 minutes

Volume Settings

This panel displays statistics for each volume, such as total volume capacity, free space, and cache file size and usage. You can also select any volume and configure volume-specific PSM attributes for that volume, as shown in Table 5.

Table 5. Persistent image volume settings

Attribute	Default value
Cache-full warning threshold	80 percent full
Cache-full persistent image deletion threshold	90 percent full
Cache size	15 percent (of the total volume capacity)

Notes:

1. You cannot change the cache size for a volume while there are persistent images on that volume (the Cache size combination box will be disabled). You must delete all persistent images on the volume before changing the cache size for that volume.
2. Cache size (as a percent of volume size) and deletion threshold must be tuned to meet the heaviest load placed on the system. NAS Gateway 300 appliances that receive heavy write traffic for sustained periods will correspondingly generate more cached data per persistent image, as the system preserves old

data from being overwritten. Very high traffic systems can devote as much as 40% of a production volume to PSM cache, although 15% (the default) or 20% will meet the needs of most users. The cache-full persistent image deletion threshold must also be tuned to automatically delete persistent images in time to free cache space before the cache fills up. Management of the cache must be tuned carefully to avoid filling the cache completely, as any missed and uncached old data renders all the persistent images for a volume inconsistent, and PSM will automatically delete them.

Persistent Images

This panel lists all of the persistent images that exist on all volumes. On this panel you can:

- Create a new persistent image immediately (without scheduling it through the Schedules panel). When you create the persistent image, you can specify properties for the persistent image, including:

Volume(s)
The persistent image can contain a single volume or multiple volumes. To select multiple volumes, hold down the **Ctrl** key while clicking the volumes. For multi-volume persistent images, a virtual directory containing data for a volume appears under the persistent image directory in the top level of each volume in the persistent image (the name of the persistent image directory is configured in the Global Settings panel).

Name
You can name the persistent image. This becomes the name of the virtual directory containing the persistent image, underneath the persistent image directory in the top level of the volume (the name of the persistent image directory is configured in the Global Settings panel).

Read-only or read-write
A persistent image is read-only by default, so no modifications can be made to it. However, you can set the persistent image to read-write, which permits you to modify it. When a persistent image is written, the modifications made are also persistent (they survive a reboot of the system). Changing a persistent image from read-write to read-only resets the persistent image to its state at the time you took the persistent image, as does selecting **Undo Writes** for a read-write persistent image from the Persistent Images panel.

Retention value
A persistent image can be given a relative retention value or weight. This is important when PSM needs to delete some persistent images for a volume because the capacity of the cache file for that volume has reached a certain threshold, as described later in this section. If the volume cache file completely fills, then all persistent images for that volume are deleted regardless of

the retention values. By default, a new persistent image is assigned a “Normal” retention value (other higher and lower values can be selected).

- Delete an existing persistent image.
- Modify properties of an existing persistent image, including read-only or read-write, and retention value.

Schedules

Use this panel to schedule persistent images to be taken at specific times (this is independent of the scheduled backup function through NAS Backup Assistant described earlier). Each PSM schedule entry defines a set of persistent images to be taken starting at a specified time and at a specified interval, with each image having the set of properties defined in the entry. This allows you to customize scheduled persistent images on a per-volume basis. For instance, you could set a persistent image for one volume to occur every hour, and for another volume to occur only once a day.

The set of properties you define are the same properties described in the Persistent Images panel description assigned above; when you define these properties, all persistent images created according to this schedule entry will be given those properties. After a scheduled persistent image is created, certain properties of that persistent image can be modified through the Persistent Images panel, independently of other persistent images created according to the schedule.

After you create a schedule entry, it appears in the list of scheduled persistent images. Subsequently, you can modify the properties of an existing entry, such as start time, repetition rate, the volumes, and so on. For a schedule, you can name the persistent images based on a pattern that you configure. The following format specifiers allow you to customize variable portions of the name:

%M 3-letter month
%D Day
%Y Year
%h Hour in 12-hour format
%s Second
%i Instance
%a AM/PM
%H Hour in 24-hour format
%W Day of week (*M, T, W* ...)
%w 3-letter day of week (*Mon, Tue, Wed* ...)
%% Percent sign

As an example, the name pattern *%w_%M_%D_%Y_%h_%m_%a* would produce the persistent image name *Mon_Apr_1_2002_10_47_AM*.

Restore Persistent Images

On this panel, you can select an existing persistent image and quickly restore the volume contained in the image back to the state it was in when the selected persistent image was taken. This is useful if you need to recover an entire volume, as opposed to just a few files. This volume restore function is available for the data volumes, but not the system volume.

Disaster Recovery

PSM provides a disaster recovery solution for the system drive. This extends the volume restore function of PSM to provide disaster recovery in the event that the system drive is corrupted to the point where the file system is corrupt, or the operating system is unbootable. Note that while disaster recovery is also supported through the Recovery CD-ROM and backup and restore capability, it is a two-step process. In contrast, the method supported by PSM allows you to restore the system drive from a single image, without having to go through the entire recovery procedure and then additionally having to restore a system drive backup.

Use the Disaster Recovery panel to schedule and create backup images of the system drive, and to create a bootable diskette that will allow you to restore the system drive from a backup image (located on the maintenance partition, or network drive). The remainder of this section provides additional information on how to perform backup and recovery operations for the NAS Gateway 300.

Note: Restoration of a PSM backup image over the network is not supported for the Gigabit Ethernet adapter. If you have only Gigabit Ethernet adapters installed, it is recommended that you perform PSM backup of each node to its maintenance partition (D: drive), which would allow you to recover if the system volume is corrupt or unbootable. Should the hard disk drive fail completely, you would need to use the Recovery CD-ROM as described in Chapter 9, “Using the Recovery and Supplementary CD-ROMs” on page 111 to restore the node to its original (factory) configuration.

Backing up the system drive

The Disaster Recovery panel lists status information for backup operations, both scheduled and immediate, as well as buttons for starting and stopping a backup operation, for configuring backup, and for creating a recovery diskette.

Click **Modify Settings** to open the Disaster Recovery Settings page. Modify the settings that you want for backup. Do not include spaces in the Backup name field. When you have modified the settings, click **OK** to save the changes.

On the Disaster Recovery page, click **Start Backup** to begin the backup. The backup process will first create a persistent image of the system drive (C:), named *System Backup*. Then, it will create the backup images from that persistent image, and then delete that persistent image when the backup operation is complete.

Creating a PSM recovery diskette

You will now create a bootable PSM recovery diskette which, when used to boot up the node, will use the backup location settings that you configured on the Disaster Recovery Settings page to locate the backup image and restore it to the system drive of the node.

1. Insert a blank, formatted diskette in the diskette drive of the node.
2. On the Disaster Recovery page, click **Create Disk**.
3. Click **OK** on the Create Recovery Disk page. The diskette drive LED will turn off when the creation is complete. The diskette creation should take no more than two minutes.
4. The utility makes the disk DOS-bootable. From a command prompt, either through the desktop of the node itself (with the diskette still in the diskette drive of the node), or on another system with the diskette in its diskette drive, type **a:\fixboot.exe** and answer the prompts.

Note: When you run fixboot.exe on the diskette, the diskette remains bootable unless you reformat it; if you later erase files on the diskette, you do not need to run fixboot.exe again.

5. Remove the diskette from the appropriate diskette drive. Label the diskette appropriately and keep it in a safe place.

You can create additional copies of the diskette using the above procedure for each new copy.

Note: If you change the backup location or logon settings using the Disaster Recovery Settings page, you must rebuild the PSM recovery diskettes for that node to reflect the new settings for that node.

Static IP addressing

If you do not have a DHCP server on your network, and you must access a backup image that is accessible only through the network (for example, no backup image is located on the maintenance partition [D: drive] of the node to be recovered), then you must configure the recovery diskette so that it will use a static IP address and subnet mask when accessing the network.

On the PSM recovery diskette, edit the file a:\net_sets.bat. Set the IPAddress and SubnetMask environment variables as follows:

1. Uncomment the two lines that begin with *rem* (comment lines) by removing the *rem* from the beginning of both lines.
2. For each line, what follows the equals sign (=) is an IP address expressed as a set of four space-separated numbers (an IP address without the dots [.]). Change the SubnetMask value to match the subnet mask that your network uses. Change the IPAddress value to match the IP address that you want to assign to the node, during the recovery operation. Do not insert dots between the numbers (octets) in either value.

As an example, here is how the lines would look for a node using IP address 192.168.1.200, and subnet mask 255.255.255.0:

```
set SubnetMask=255 255 255 0
set IPAddress=192 168 1 200
```

If you later want to reconfigure the recovery diskette to use DHCP to obtain an IP address instead of static IP addressing, you must reinsert *rem* in front of the SubnetMask and IPAddress lines to disable static IP addressing, as follows (based on the previous example):

```
REM set SubnetMask=255 255 255 0
REM set IPAddress=192 168 1 200
```

Restoring the system drive using the PSM recovery diskette

To restore the system drive from a backup image created through the PSM Disaster Recovery panel as described above, you must use a PSM recovery diskette created through the Disaster Recovery panel. If you did not create a PSM recovery diskette, you must use the Recovery CD-ROM as described in Chapter 9, “Using the Recovery and Supplementary CD-ROMs” on page 111 to restore the system drive to its original (factory) configuration.

To restore the system drive:

1. Set the write-protection tab of the PSM recovery diskette to the write-protect position. This prevents accidental initiation of the recovery process (by booting the node with the PSM recovery diskette in the diskette drive).

2. Insert the PSM recovery diskette in the diskette drive of the node, and restart the node.
3. The recovery process begins. The PSM recovery diskette software locates the first backup image it can find, based on the backup locations specified when the diskette was created. When it locates a backup image, it begins restoring the system drive from the image. During the restore operation, the hard disk drive LED (on the front right of the node's hard disk drive) will flash green or stay nearly solid green; this indicates write activity to the system volume.

   Note: If the hard-disk drive LED stays off for at least 10 minutes since you restarted the node, then there is a problem with the recovery procedure and it will not be able to restore the system volume from a backup image. Should this occur, you will need to restore the system drive as described in Chapter 9, "Using the Recovery and Supplementary CD-ROMs" on page 111.
4. When the restore operation completes, the hard disk drive LED turns off, and a short song will play periodically (every 15 seconds). Remove the diskette, set the write-protection tab back to the write-enabled position, and reinsert the diskette. The log file RESULTS.HTM will be written to the diskette; this log file can be viewed with any Web browser to examine the results of the restore operation.
5. When the log file is written, another song will play (continuously). Remove the diskette and restart the node. If the restore was successful, the node will come back up in the state it was in at the time when you created the backup image used for the recovery operation.

   Note: The persistent image that was created on the system drive (named *System Backup*) by the backup process is restored by the restore process as it is preserved in the backup image. It is recommended that you now delete that persistent image as it is no longer needed. On the Persistent Images panel, select the persistent image named **System Backup** on drive C: from the list of persistent images, then click **Delete**, then click **OK** on the Delete Persistent Image panel that appears.

If the restore was unsuccessful, then you must use the Recovery CD-ROM as described in Chapter 9, "Using the Recovery and Supplementary CD-ROMs" on page 111.

Rebuilding the maintenance partition

If this is a new hard drive or if the Maintenance (D:) partition is unusable, you must rebuild the Maintenance partition by performing the following steps:

1. Start Disk Management on the node. You can do this in one of two ways:
   - Start a Terminal Services session to the node, then click the **IBM NAS Admin** icon, and then from the IBM NAS Administration console that appears, select **Computer Management**, then **Disk Management**.
   - Start a Windows 2000 for NAS user interface session to the node, then select **Disks and Volumes**, then select **Disks**, and then provide your administrator user name and password when prompted.
2. In the Disk Management window, right-click the unallocated area of Disk 0, and then click **Create Partition**.
3. In the Create Partition wizard, click **Next** and select **Primary Partition**.
4. Click **Next** and select **D:** as the drive letter.
5. Click **Next** and select **FAT32** as the file system and change the drive label to *Maintenance*.

6. Click **Finish** to close the wizard.

The partition will then be formatted. When formatting is complete, the status of the partition should appear as *Healthy*, and the other properties should appear as:

- Name: *Maintenance*
- Drive letter: *D:*
- File system: *FAT32*

Granting user access to persistent image files

You can give end-users access to files in the persistent images. For example, this would be helpful to a user who has accidentally corrupted a file and needs to get an uncorrupted copy of that file.

To enable end-user access to persistent image files:

1. Go into Terminal Services.
2. Click the **My Computer** icon.
3. Select the volume on which you want to enable persistent image access.
4. Go into the persistent images directory and right-click the mouse on the selected persistent image mount point, select **Sharing**, then specify sharing as appropriate. If you want to enable the same access to all persistent images on the volume, right-click the persistent images directory (from the top level of the volume), select **Sharing**, and then specify sharing as appropriate.

Note: The share settings are maintained in a persistent image. Therefore, granting access to all end-users only permits those users to access files and directories within the persistent image that they had permission to access originally on the actual drive.

PSM notes

- You can take and keep a maximum of 250 persistent images at one time. These can be taken on local drives, or drives on the external storage that are logically local.

  On various panels, such as the New Persistent Image Schedule panel, the *Keep the last:* field indicates the number of persistent images. The total number of persistent images that you enter in these fields does not override the maximum number of persistent images that you set in the Global Settings panel. For example, if the maximum number of persistent images is 10, and you enter numbers in other fields that add up to greater than 10, only 10 persistent images will be taken.
- You cannot take a persistent image of the maintenance drive (D:). Hence, you will not see it as a choice in either the New Persistent Image Schedule panel or the Create Persistent Image panel. Do not take a persistent image of the clustering Quorum disk. See “Recovering from a corrupted Quorum drive” on page 52 for information on how to recover from a corrupted Quorum drive.
- PSM stores the cache file for each drive on the drive itself. The first persistent image created on a particular drive will require a significant amount of time because the PSM cache file must be created (pre-allocated) for that drive.

  The time required for creation depends on the configured size of the cache file (15 percent of the total drive size by default). Creation takes roughly three to four minutes per gigabyte. For example, a 10-GB cache file would require 30 to 40 minutes to create. You should create a persistent image for a drive before

scheduling any persistent images for that drive, to build the cache file. You can then delete the persistent image that you just created if you do not need to keep it.

After the creation of the first persistent image on a volume, future persistent images on that volume will complete faster.

- The default size of the cache file per drive is 15 percent of the total drive capacity.

  In most cases, that should be sufficient. However, it is possible that it will not be enough to maintain the number of persistent images that you want to keep concurrently on the drive, given the amount of file-write activity to the drive. PSM automatically takes action to prevent the cache file from overflowing, because if that occurred, PSM would be forced to automatically delete all persistent images on the drive (when it cannot keep track of changes made to the drive, it cannot maintain a valid persistent image).

  PSM takes the following actions as the cache file usage approaches a full condition:

  – When the cache file usage exceeds the warning threshold (configured in the PSM Volumes panel for the drive; the default value is 80 percent), PSM generates a warning message to the system event log (viewable through the Windows 2000 Event Viewer in the IBM NAS Admin console), and to the alert log in the Microsoft Windows 2000 for Network Attached Storage user interface. The name of the source for the message is *psman5*. Additionally, while the cache file usage is above the warning threshold, PSM prohibits any attempt to create a new persistent image, and logs error messages (to the system log and alert log). The text of the error message that is logged in the system event log (from *psman5*) is “A persistent image could not be created due to error 0xe000102b”.

  – When the cache file usage exceeds the automatic deletion threshold (also configured in the PSM Volumes panel for the drive; the default value is 90 percent), PSM automatically selects a persistent image on the volume and deletes it to reduce the cache file usage. It selects the persistent image with the lowest retention value (as described in “Persistent Images” on page 64). If more than one persistent image has the same (lowest) retention value, then the oldest image will be selected for deletion. If this deletion does not reduce the cache file usage below the automatic deletion threshold, then it will continue to select and delete persistent images until the cache file usage is reduced below the automatic deletion threshold. For each deletion, PSM generates an error message to the system event log and to the Windows 2000 for Network Attached Storage alert log indicating that a persistent image was deleted.

  You should periodically check the system event log or Windows 2000 for Network Attached Storage alert log to ensure that the cache file usage is not consistently high, forcing existing persistent images to be deleted and preventing new persistent images from being created. If the cache file usage is high, you can increase the size of the cache file using the PSM Volumes page. However, because dynamic cache file resizing is not supported in this release, you must delete all persistent images currently on that volume first.

- When a shared volume performs a failover operation from one engine in the NAS Gateway 300 to the other engine, the persistent images for that volume move with the volume. The Persistent Images panel on a particular engine will display only those persistent images which are on volumes that the engine owns at a point in time. If persistent images are scheduled for a volume, on a particular

engine, a scheduled persistent image is created only as long as that engine owns the volume at the time the scheduled persistent image is to occur.

To ensure that a scheduled persistent image will take place regardless of which engine owns the volume, you must do the following:

1. Use the Schedules panel to create the schedule on the engine that currently owns the volume.
2. Use the Cluster Administrator to move the disk group that contains the volume to the other engine. You can create or edit a schedule only for a volume on the engine that currently owns the volume. If an engine does not own the volume, you cannot select the volume when creating a new schedule through the New Persistent Image Schedule panel (under Schedules).
3. Use the Schedules panel on the other engine to create the same schedule that you created on the original engine, with all of the same parameters (start time, frequency, number to keep, and so on).
4. Use the Cluster Administrator to move the disk group that contains the volume back to the original engine.

- Volume restore of the system volume (C: drive) is not supported. If you attempt to restore a persistent image containing the system volume, the restore operation will not take place.
- Volume restore of a data volume might require a reboot of the node. You will be notified by the Restore Persistent Images panel whether a reboot is required after a restore operation is initiated.
- When you restart the NAS Gateway 300 (“restart” in this case means that with both nodes down, the node that was shut down last is restarted first so that it initially owns all of the shared data volumes), Persistent Storage Manager (PSM) takes two actions:

  1. Loading
  2. Mapping

  During loading, PSM loads existing persistent images from the cache files on each of the volumes. The loading time depends on the amount of cache data there is to read. Cached data is used by PSM to maintain the persistent images, and the more cache data there is, the longer it takes to load the persistent images, and thus the longer it might take the NAS Gateway 300 to become fully operational after a restart.

  During mapping, PSM makes the loaded persistent images accessible through the file system by mounting each of them as a virtual volume underneath the persistent images directory on the real volume for which the persistent image was created. Mapping takes place five minutes after the real volume has been mounted. The mapping time varies with the number of persistent images, as well as the size of the volume.

  As an example, suppose that on your NAS Gateway 300, you defined a 1 TB volume with 50 percent of the volume allocated to the cache (500 GB cache), and that you had 20 persistent images on the volume, using 100 GB (20 percent) of the cache (based on the write activity to the volume since the first persistent image was created). You would observe an increase in the startup time of roughly 3 minutes, 20 seconds over what it would be without any persistent images on the volume. Then, once the NAS Gateway 300 has become fully operational, all 20 persistent images would become accessible within another 18 minutes (including the five minutes that PSM waits after the volume comes up to begin the mapping).

When a volume is moved between nodes during a failover operation, then PSM must perform persistent image loading and mapping on the node to which the volume is moving, just as it does when the “first node” is restarted.

In the failover scenario, loading must take place before the volume can be brought online on the node (when the clustered disk resource is shown as being *Online* in Cluster Administrator). Then, as in the restart case, mapping begins five minutes after the volume comes online.

Microsoft Cluster Server, which controls the disk resource failover, waits a certain period, called the pending timeout, for the disk to come online. (During the loading phase, the disk resource is shown as being in *Online Pending* state.) With a default value of 180 seconds (3 minutes) for the pending timeout, this time interval might be exceeded because of the time it takes to load the persistent images on the volume. If this occurs, the delay might cause Cluster Server to mark the disk as *Failed* and to not be available to either NAS Gateway 300 node. Other dependent resources (IP addresses, network names, file shares, and so on) might also fail.

For this reason, it is recommended that you increase the pending timeout value for all clustered resources to 1200 seconds (20 minutes). To do this, open Cluster Administrator, select **Resources** from the left pane to display all clustered resources in the right pane, and then for each resource listed in the right pane:

1. Right-click the resource name and select **Properties**.
2. Select the **Advanced** tab.
3. Change the Pending timeout value to 1200 (seconds).
4. Click **Apply**; then click **OK**.

- PSM imposes a limit of 1 terabyte (TB) of cached data, across all volumes on the NAS Gateway 300. For this reason, you should ensure that the total configured size of all cache files on the NAS Gateway 300 is not greater than 1 TB.

  You can do this by accessing Persistent Storage Manager, then going to the Volume Settings page, and making sure that the total of all values in the Cache Size column is 1 TB or less. (You can access Persistent Storage Manager through the Persistent Storage Manager link on the IBM NAS Admin console on the NAS Gateway 300 desktop, or by starting the Windows 2000 for Network Attached Storage user interface and then selecting **Disks**, then **Persistent Storage Manager**.)

  If the total is greater than 1 TB, you should reduce the size of the cache on one or more of the volumes by selecting the volume from the list, then clicking **Configure**, and then selecting a smaller value from the “Cache size” drop-down list and clicking **OK**.

  Note: You cannot change the size of the cache on a volume that has persistent images. You must delete all persistent images on the volume before changing the cache size. You should try to reduce the cache size on a volume that has no persistent images, if possible, before deleting any persistent images.

  If more than 1 TB of cache is configured on the NAS Gateway 300, the following can occur (note that a volume for which a persistent image has never been created is considered to have a cache size of zero, regardless of how large its cache is configured to be):

  – When the NAS Gateway 300 is restarted, PSM prevents a volume from being mounted on the file system (prevents it from being accessible) if that volume’s

PSM cache would increase the total size of all cache files (on all volumes mounted to that point) above 1 TB, and an error message is written to the system event log. The event source is psman5, and the text of the error message is:

```
There is insufficient memory available.
```

– When a volume is failed over between nodes, then PSM running on the "new" node will behave as it would if the volume were being mounted during a restart: if that volume's PSM cache would increase the total size of all cache files on that node above 1 TB, then PSM blocks the mount and writes the "insufficient memory available" error message to the system event log. (This will also cause the failover to fail, which means that either the volume will try to come online on the "original" node if it is up, or just simply fail to come online at all.)

– If you increase the size of any cache such that the total cache size of all volumes on the NAS Gateway 300 becomes greater than 1 TB, and if you do not restart the NAS Gateway 300 after you change the cache size, then no persistent images can be created on the volume for which the cache size increase was made. An attempt to create a persistent image on that volume will cause an error message to be written to the system event log. The event source is psman5, and the text of the error message is:

```
There is insufficient memory available.
```

- If you delete the last persistent image on a volume, and then immediately attempt to create a new persistent image on that volume, the creation of the new persistent image might fail, and an error message will be written to the system event log.

  The event source is psman5, and the text of the error message is:

  ```
  A persistent image could not be created due to error 0xc0000043.
  ```

  This message is generated because when PSM is reinitializing the PSM cache file on a particular volume (after you delete the last persistent image on that volume), a new persistent image cannot be created. If this error occurs, wait for a few minutes, and then try to create the persistent image again.
- If you use the Windows Powered Disk Defragmenter to attempt to defragment a volume containing persistent images, the volume will not be defragmented. If you select the volume and click the **Defragment** button, the Disk Defragmenter will run on the volume and then indicate that the volume was successfully defragmented. However, the Analysis display will appear the same as it did before you clicked **Defragment**, which indicates that defragmentation did not take place. You can defragment volumes without persistent images.
- PSM uses several system level files, one of which has a command line interface. Use of this interface is only supported for IBM-provided applications and services, as well as IBM support technician assisted debugging efforts. All PSM function including sophisticated scheduling and automation of remote management is provided by the Windows 2000 for NAS Web-based GUI.

Attention: The recovery process invalidates persistent images and leaves them in an inconsistent state. So, if you plan to use the Recovery CD-ROM, it is recommended that you first delete all persistent images to ensure a clean reload of the system software. For more information on using the Recovery CD-ROM, see Chapter 9, "Using the Recovery and Supplementary CD-ROMs" on page 111.

Storage Manager for SAK

The NAS Gateway 300 includes Storage Manager for SAK, a storage management tool that includes the following functions:

- Storage reports
- Directory quotas
- File screening

Storage reports address disk usage, wasted storage space, file ownership, security, and administration. Reports can run interactively, scheduled on a regular basis, or run as part of a storage resource management policy when disk-space utilization reaches a critical level.

Directory quotas allow the administrator to add, delete, monitor, and change disk-space limits for selected directories on the NAS appliance. Directory quotas provide disk-space monitoring and control in real time and supports active and passive limits with two real-time space alarms.

File screening allows the blocking of any file type such as MP3, graphics files, VBS viruses, and executables from writing to the NAS appliance.

Uninterruptible power supply support

The NAS Gateway 300 includes support for uninterrupted power supplies (UPS). UPS devices provide emergency backup power for a specific period of time when the local power fails. This power comes from batteries housed within the UPS. High-performance surge suppression helps protect your appliance from electrical noise and damaging power surges. During a power failure, the UPS is designed to instantly switch your appliance to emergency battery-backup power. After you have installed a UPS for your appliance, you can set options for its operation using the UPS task on the Maintenance page. The UPS task enables you to control how the UPS service works on your appliance. The available UPS settings depend on the specific UPS hardware installed on your system. Before you use your UPS device, type the following information on the UPS Configuration page:

- UPS device manufacturer
- UPS device model
- The serial port to which the UPS device is connected

To configure the UPS service, click **UPS** on the Maintenance page.

To help protect your server appliance from power failures, test it by simulating a power failure by disconnecting the main power supply from the UPS device. Do ***not*** perform this test during production use. Your appliance and peripherals connected to the UPS device should remain operational, messages should be displayed, and events should be logged. Wait until the UPS battery reaches a low level to ensure that a proper shutdown occurs. Restore the main power to the UPS device, and check the event log to verify that all actions were logged and there were no errors. All detected power fluctuations and power failures are recorded in the event log, along with UPS service start failures and appliance shutdown initiations. Critical events might change the status of the appliance.

Tivoli SANergy

Note: The NAS Gateway 300 is enabled for SANergy use. Although the SANergy component is included in the product, you will need to obtain additional licenses from Tivoli to use the SANergy client with this appliance.

Tivoli SANergy allows you to deliver shared data access at the speed of a SAN, using Fibre Channel, SCSI, or SSA. It gives multiple computers the power to dynamically share file and data access on SAN-based storage, using standard networks and file systems.

SANergy combines LAN-based file sharing with the very high data transfer speeds of the Fibre Channel, SCSI, and SSA storage networks. The result is high-speed, heterogeneous data sharing without the performance-limiting bottlenecks of file servers and traditional networking protocols.

SANergy extends standard file systems and network services provided by the operating systems that it supports. As an operating system extension built on standard systems interfaces, SANergy fully supports the user interface, management, access control, and security features native to the host platforms. Therefore, it provides you with all the file system management, access control, and security you expect in your network.

With SANergy, applications that you have configured into your network can access any file at any time, and multiple systems can transparently share common data. SANergy ensures maximum compatibility with existing and future operating systems, applications, and management utilities.

In addition to the SAN, SANergy also uses a standard LAN for all the metadata associated with file transfers. Because SANergy is based on standard file systems, if your SAN fails, you can continue to access your data through the LAN.

With SANergy, you can reduce or even eliminate the expense of redundant storage and the overhead of data synchronization in multi-host environments. These environments include large-scale Web, video, or file servers. Because each system has direct access to the SAN-based storage, SANergy can eliminate your file server as a single point of failure for mission-critical enterprise applications, reducing costly downtime. Also, with SANergy, you can readily manage all your data backup traffic over the storage network, while users have unimpeded LAN access to your existing file servers.

To set up SANergy on your network, use the following steps:

1. Make sure your network is correctly configured for LAN and SAN.
2. Configure your storage system, including disk formatting, partitioning, and volume configuration.
3. Enable SANergy bus management and device assignments. The Meta Data Controller will then be enabled for the appliance operating system.
4. Install additional SANergy licenses on properly configured hardware clients.

Further details on SANergy are contained in the online help.

Antivirus protection

You can perform antivirus scanning of NAS Gateway 300 storage from clients having the appropriate access permissions. Also, you can install Norton AntiVirus Version 7.5 or later on the NAS Gateway 300 engine using standard Windows 2000 software installation procedures.

Depending on configuration options, antivirus scanning might use substantial CPU or disk resources. Therefore, you should carefully select the scanning options and schedule.

Chapter 7. Managing adapters and controllers

This chapter describes the functions that you can use to manage various adapters and controllers installed in the NAS Gateway 300.

The following functions are available:

- "Managing Fibre Channel host bus adapters", accessed through the IBM NAS Admin
- "Enabling communication between system management adapters" on page 78

Managing Fibre Channel host bus adapters

The FAStT MSJ diagnostic utility allows you to manage and control Fibre Channel host bus adapters. With FAStT MSJ, you can:

- Retrieve and display general information about the adapters
- Request and display the real-time statistics of adapters
- Diagnose operations on the adapters and attached devices
- Display the NVRAM parameters of adapters (note that you cannot change the parameters)
- Monitor alarms and indications of the adapters

The primary purpose of FAStT MSJ in the NAS Gateway 300 is to obtain diagnostic information about the Fibre Channel connections.

To use FAStT MSJ:

1. Start FAStT MSJ by double-clicking the **IBM NAS Admin** icon.
2. Under the **NAS Management** icon, double-click **Storage**, and then **NAS Utilities**.
3. Select **FAStT MSJ**.
4. When the FAStT MSJ opens:
   - If you are connected locally with a monitor, keyboard, and mouse, select **localhost**; then click **Connect**.
   - If you are connected through Terminal Services, type the host name or IP address of the machine you are connected to through Terminal Services; then click **Connect**.

For further details on FAStT MSJ, see the online help.

Appendix E, "Fast!UTIL options" on page 155 provides detailed configuration information for advanced users who want to customize the configuration of the FAStT Host Adapter board and the connected devices, using Fast!UTIL to make changes.

Enabling communication between system management adapters

The two types of system management adapters are[2]:

- The Integrated System Management Processor (ISMP) integrated on the planar board of each engine of the NAS Gateway 300

  Provides basic operational status about key engine components, such as its processors, power supplies, fans, and so on.

- An optional Remote Supervisor Adapter (RSA) that can connect to up to twelve of the ISMPs

  The RSA allows you to connect through a LAN or modem from virtually anywhere for extensive remote management. The RSA works in conjunction with the ISMP of the NAS Gateway 300 and an interconnect cable that connects multiple engines to the ISMP. Remote connectivity and flexibility with LAN capability is provided by Ethernet connection. Along with ANSI terminal, Telnet, and IBM Director, the RSA enables more flexible management through a Web browser interface.

For more information, see “Using the RSA” on page 80.

Table 6 provides a summary of the features of the ISMP and the RSA.

The light-path diagnostics LED status that is available through the ISMP includes:

- Power-supply failure
- Insufficient power for power-supply redundancy
- Exceeded power-supply capabilities
- Non-maskable interrupt occurred
- Over heating
- Fan failure
- Memory error
- Microprocessor failure
- PCI-bus error
- VRM failure
- Planar SCSI failure for system disk or internal tape drive (if any)

Remote status includes information on power supply voltages, voltage-regulator module (VRM) readings, temperatures of system components, system power status, power-on hours, fan status, and system state.

Table 6. ISMP compared to the RSA

Feature	ISMP	RSA
Location	On planar board	Separate PCI adapter option
Light-path diagnostics	Remotely reports on	Remotely reports on
LED status of engine	Remotely reports on	Remotely reports on
LED status of HDD in engine	No	No
Remote update of system BIOS	Yes	Yes
Remote update of ISMP BIOS	No	Yes

2. A third might be referred to in some of the documentation that came with your system, but that adapter is not used in the NAS Gateway 300.

Table 6. ISMP compared to the RSA (continued)

Feature	ISMP	RSA
Immediate remote power on/off	Yes	Yes
Controlled remote power on/off using the OS	No	Yes
Remote POST (including all POST message IDs)	No	Yes
Remote access to engine vital product data (VPD) and serial number	No	Yes
Multiple login IDs	No	Yes
TELNET interface over IP	No	Yes (through a LAN connection)
Web-browser interface over IP	No	Yes
Forwarding of SNMP traps	Yes, to the RSA	Yes (through a LAN connection)
Automated server restart	Yes	Yes
Remote Alerts	No	Yes
Configuration	By DOS utility	By DOS utility/serial ports
Aggregate from other ISMP processors	No	Yes

Enabling ISMP to RSA communication on a single machine

You must follow one of two methods to enable communication between the ISMP and the RSA on a single machine:

- Using a single ISMP interconnect cable (with dual RJ-11 plugs):
  1. Connect one end of the *internal* ISMP interconnect cable to the J-54 connector on the system board.
  2. Connect the other end (the RJ-11 socket) of the *internal* ISMP interconnect cable to the knockout slot on the back panel of the machine until it locks into place.
  3. Connect one connector on the ISMP interconnect cable to the RJ-11 socket that you just installed on the back panel (in step 2).
  4. Connect the other connector to the RJ-11 socket on the RSA.
- Using two ISMP interconnect cables (each with a single RJ-11 plug):
  1. Connect one end of the *internal* ISMP interconnect cable to the J-54 connector on the system board.
  2. Connect the other end (with the RJ-11 socket) of the *internal* ISMP interconnect cable to the knockout slot on the back panel of the machine until it locks into place.
  3. Connect the first ISMP interconnect cable to the RJ-11 socket that you just installed on the back panel (in step 2).
  4. Connect the second ISMP interconnect cable to the RJ-11 socket on the RSA.
  5. Connect the two ISMP interconnect cables with a single Category 5 Ethernet cable (by plugging one end of the Ethernet cable into the “black box” on the first ISMP interconnect cable, and the other end into the “black box” on the second ISMP interconnect cable).

Using the RSA

The documentation CD-ROM that came with your system contains additional information and software for the RSA.

To use the RSA, complete the following steps:

1. Consult the RSA user's manual and the README file that is located on the documentation CD-ROM.
2. Run the executable to create a bootable floppy disk. The executable is located in:

   `C:\IBM\ASMP\UPDATES\33P2474.EXE`
3. Boot each node of the NAS Gateway 300 with floppy disk created in the previous step to configure the RSA.

Enabling Ethernet adapter teaming

This section describes how to enable adapter teaming on the Ethernet adapters.

Note: The integrated Ethernet controller on each NAS Gateway 300 node is dedicated to the clustering interconnection between it and another node and cannot be used for teaming.

The Ethernet adapters that you install in the PCI slots of the NAS Gateway 300 nodes support *adapter teaming* (also known as load balancing). With adapter teaming, two or more PCI Ethernet adapters can be physically connected to the same IP subnetwork and then logically combined into an adapter team.

The NAS Gateway 300 uses two different Ethernet vendors, Intel (PRO/1000 XT Server Adapter by Intel and IBM Gigabit Ethernet SX Server Adapter) and Alacritech (Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter and Alacritech 100x4 Quad-Port Server Accelerated Adapter). Each vendor offers different means of implementing teaming, and teaming cannot be done between adapters of different vendors. There might also be restrictions on what functions are supported with adapters from the same vendor.

Note: It is strongly recommended that you configure adapter teaming before you set up Microsoft Cluster Server (MSCS) clustering, as described in Chapter 5, "Completing networking, clustering, and storage access setup" on page 33. Additionally, for each team that you configure on one node, you must configure an identical team (same type of team, same set of adapters, and so on) on the other node.

Alacritech Ethernet adapter teaming

Alacritech uses SLIC (Session-Layer Interface Card) technology, which incorporates hardware assistance for TCP processing. Most, but not all, of the processing overhead for TCP/IP is removed from the NAS engine. This is an optional feature and can be disabled if required.

Alacritech offers four methods of teaming:

Cisco Fast EtherChannel (Fast EtherChannel and Gigabit EtherChannel compatible)

Fast EtherChannel (FEC) is a proprietary technology developed by Cisco. With FEC, you can create a team of two to four ports on an adapter to increase transmission and reception throughput. The FEC might also be

referred to as load balancing, port aggregation, or trunking. When you configure this feature, the adapter ports comprising the FEC team or group create a single high-speed, fault-tolerant link between the engine and the Ethernet switch sharing one IP address. With FEC, fault tolerance and load balancing is provided for both outbound and inbound traffic, unlike other load-balancing schemes that balance only outbound traffic. Fast EtherChannel and Gigabit EtherChannel (FEC/GEC) requires a FEC/GEC compatible switch. The same teaming also must be enabled on the connected switch ports.

Note: FEC requires an Ethernet switch with FEC capability. The FEC implementation on the Alacritech 100x4 Quad-Port Server Accelerated Adapter does not support the optional Port Aggregation Protocol (PAgP) feature of FEC-capable Ethernet switches. Likewise, The FEC/GEC implementation on the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter does not support the optional PAgP feature of FEC/GEC-capable Ethernet switches.

The following are the valid teaming configurations and restrictions for Cisco EtherChannel teaming with Alacritech adapters:

- Two Alacritech 1000x1 Single-Port Server and Storage Accelerated adapters can be teamed together.
- No Alacritech 1000x1 Single-Port Server and Storage Accelerated adapters can be teamed with any port of the Alacritech 100x4 Quad-Port Server Accelerated Adapters.
- One Alacritech 100x4 Quad-Port Server Accelerated Adapter can have two or more of its ports teamed.
- Two Alacritech 100x4 Quad-Port Server Accelerated Adapters can have any ports on any card teamed (limited to four ports per team). For example, two ports on one card can be teamed with two ports on a second card.

IEEE 802.3ad Link Aggregation Group

802.3ad is an IEEE industry-standard similar to the Cisco FEC/GEC. 802.3ad requires an Ethernet switch with 802.3ad capability. Alacritech does not support the optional Port Aggregation Protocol (PAgP) feature of some FEC switches or the 802.3ad LACP protocol. PAgP/LACP facilitates the automatic creation of link aggregation groups. All EtherChannel and Link Aggregation groups must be manually configured.

The following are the valid teaming configurations and restrictions for IEEE 802.3ad teaming with Alacritech adapters:

- Two Alacritech 1000x1 Single-Port Server and Storage Accelerated adapters can be teamed together.
- No Alacritech 1000x1 Single-Port Server and Storage Accelerated adapters can be teamed with any port of the Alacritech 100x4 Quad-Port Server Accelerated Adapters.
- One Alacritech 100x4 Quad-Port Server Accelerated Adapter can have two or more of its ports teamed.
- Two Alacritech 100x4 Quad-Port Server Accelerated Adapters can have any ports on any card teamed (limited to four ports per team). For example, two ports on one card can be teamed with two ports on a second card.

Send-Only Load Balancing

This is an inexpensive way to do load balancing when using a Ethernet switch that does not support FEC or 802.3ad. However, if TCP/IP acceleration is used with this method, all ports that are teamed must be on the same physical adapter. There is no load balancing when receiving. The following are valid teaming configurations and restrictions when doing send-only load balancing are:

- Two Alacritech 1000x1 Single-Port Server and Storage Accelerated adapters can be teamed together; however, acceleration is disabled.
- No Alacritech 1000x1 Single-Port Server and Storage Accelerated adapters can be teamed with any port of the Alacritech 100x4 Quad-Port Server Accelerated Adapters.
- One Alacritech 100x4 Quad-Port Server Accelerated Adapter can have two or more of its ports teamed, with acceleration enabled.
- Two Alacritech 100x4 Quad-Port Server Accelerated Adapters can have any ports on any card teamed (limited to four ports per team). For example, two ports on one card can be teamed with two ports on a second card. Acceleration will be disabled because the ports are on different adapter cards.

Hot Standby Failover

This technique does no load balancing but does allow failover and redundancy. One port is put online while the remaining ports in the team are offline. If the link for the online ports fails, it is taken offline and one of the other ports takes its place. It is not required that the ports in the team be on the same adapter. It is also not required that they be the same speed, although it is recommended. The following are valid teaming configurations:

- Two Alacritech 1000x1 Single-Port Server and Storage Accelerated adapters can be teamed together.
- No Alacritech 1000x1 Single-Port Server and Storage Accelerated adapters can be teamed with any port of the Alacritech 100x4 Quad-Port Server Accelerated Adapters.
- One Alacritech 100x4 Quad-Port Server Accelerated Adapter can have two or more of its ports teamed.
- Two Alacritech 100x4 Quad-Port Server Accelerated Adapters can have any ports on any card teamed (limited to four ports per team). For example, two ports on one card can be teamed with two ports on a second card.

To configure adapter teaming with the Alacritech adapters, perform the following steps:

1. Click **Control Panel**.
2. Click **Network and Dial-Up**.
3. Click **Adapter**.
4. Click **Properties**.
5. Click **Alacritech SLIC Team Configurator**.
6. Click **New Team**.

Intel Ethernet adapter teaming

Intel offers five teaming modes:

Adapter Fault Tolerance (AFT)

Adapter Fault Tolerance (AFT) is similar to Hot-Standby Failover for the Alacritech adapters. Only one adapter in the team is fully active on the Ethernet network (for example, sending and receiving data) at any point in time, while the other adapters are in standby mode (receiving data only). If that adapter detects a link failure or fails completely, another adapter in the team automatically and rapidly takes over as the active adapter, and all Ethernet traffic being handled by the failing adapter is seamlessly switched to the new active adapter, with no interruption to network sessions (for example, file transfers) in progress at the time of the failover.

An AFT team consists of between two and eight ports. In the NAS Gateway 300, the maximum number of ports is four, because all Intel adapters are single port and the total number of network cards is four. All adapters in the team should be connected to the same hub or switch with Spanning-Tree Protocol (STP) set to OFF. The team members can be different speeds or different adapters.

The following are valid teaming configurations for AFT with Intel adapters:

- Two IBM Gigabit Ethernet SX Server Adapters
- Two PRO/1000 XT Server Adapter by Intels
- One to two IBM Gigabit Ethernet SX Server Adapters with one to two PRO/1000 XT Server Adapter by Intels

Switch Fault Tolerance (SFT)

Two adapters connected to two switches to provide network availability of a second switch and adapter if the first switch, adapter, or cabling fails. STP must be set to ON. The following are valid teaming configurations for Intel SFT:

- Two IBM Gigabit Ethernet SX Server Adapters
- Two PRO/1000 XT Server Adapter by Intels
- One IBM Gigabit Ethernet SX Server Adapter with one PRO/1000 XT Server Adapter by Intel

Adapter Load Balancing (ALB)

Adapter Load Balancing (ALB) is similar to Send-Only Load Balancing for Alacritech adapters. All adapters in the team are active, increasing the total transmission throughput over the common IP subnetwork. If any adapter in the team fails (link failure or complete failure), the other adapters in the team continue to share the network transmission load, although total throughput is decreased. Load balancing is supported only for adapter teams consisting of only one type of adapter; different types of adapters cannot be combined in a load-balancing team.

Two to eight ports from any Intel adapters are combined in a team that allows increased network bandwidth when sending. AFT is also included. When receiving, only the port identified as primary receives data. There are no special switch requirements.

The following are valid teaming configurations for ALB with Intel adapters:

- Two IBM Gigabit Ethernet SX Server Adapters
- Two PRO/1000 XT Server Adapter by Intels
- One to two IBM Gigabit Ethernet SX Server Adapters with one to two PRO/1000 XT Server Adapter by Intels

Cisco Fast Etherchannel (FEC/GEC compatible)

FEC is a proprietary technology developed by Cisco. With FEC, you can create a team of two to four ports on an adapter to increase transmission and reception throughput. The FEC might also be referred to as load balancing, port aggregation, or trunking. When you configure this feature, the adapter ports comprising the FEC team or group create a single high-speed, fault-tolerant link between the engine and the Ethernet switch sharing one IP address. With FEC, fault tolerance and load balancing is provided for both outbound and inbound traffic, unlike other load-balancing schemes that balance only outbound traffic. FEC/GEC requires a FEC/GEC compatible switch. The same teaming must also be enabled on the connected switch ports.

The following are valid teaming configurations for Cisco FEC/GEC with Intel adapters:

- Two IBM Gigabit Ethernet SX Server Adapters
- Two PRO/1000 XT Server Adapter by Intels
- One to two IBM Gigabit Ethernet SX Server Adapters with one to two PRO/1000 XT Server Adapter by Intels

IEEE 802.3ad Link Aggregation Group

802.3ad is an IEEE industry-standard similar to the Cisco FEC/Gigabit Etherchannel (GEC). 802.3ad requires an Ethernet switch with 802.3ad capability. PAgP/LACP facilitates the automatic creation of link aggregation groups. All EtherChannel/Link Aggregation groups must be manually configured.

For the Intel adapters, there are two implementations of the standard. Static is equivalent to Etherchannel and requires a FEC/GEC, 802.3ad or Intel Link Aggregation capable switch. Dynamic requires 802.3ad dynamic capable switches.

The following are valid teaming configurations for IEEE 802.3ad with Intel adapters:

- Two IBM Gigabit Ethernet SX Server Adapters
- Two PRO/1000 XT Server Adapter by Intels
- One to two IBM Gigabit Ethernet SX Server Adapters with one to two PRO/1000 XT Server Adapter by Intels

To configure adapter teaming with the Intel adapters, use Intel PROSet II, which is preloaded on the NAS Gateway 300, as follows:

1. Physically connect the adapters that you want to team to the same IP subnetwork.
2. Access the NAS Gateway 300 desktop by directly attaching a keyboard, mouse, and monitor, or over the network by starting Terminal Services on another workstation (see “Terminal Services and the IBM NAS Administration console” on page 15).
3. From the NAS Gateway 300 desktop, click **Start** → **Settings** → **Control Panel**.
4. Double-click the Intel PROSet II icon in the Control Panel to start Intel PROSet II. You will see a list of all adapters for each slot and type supported under Network Components.
5. Under Network Components, you will see a list of resident and nonresident adapters for each slot and type supported. Drivers are preset for all supported adapter configurations but will be loaded only for resident adapters.

6. Identify which adapters you are going to team. Left-click the adapter under Network Components, and select one of the adapters that will be part of the teaming.
7. Right-click the adapter → **Add to Team** → **Create New Team...**.
8. Select the type of team to create.
9. Select the adapters to add to the team from the list, and then click **Next**.
10. Verify that these settings are correct, and then click **Finish**.
11. Perform Steps 1 on page 84 through 10 for the other node.

This procedure creates a device named Intel Advanced Network Services Virtual Adapter. It also binds all network protocols that were bound to the physical adapters that were added to the team to this virtual adapter and unbinds those protocols from the physical adapters. If you delete the team, the settings will return to the state prior to creating the team.

For complete help on adapter teaming, from Intel PROSet II, click **Network Components**, and then select **Help** from the Help menu.

RAID-1 mirroring

The NAS Gateway 300 hardware has a RAID-1 mirroring option using the onboard SCSI adapter. The System and Maintenance partitions are mirrored using two 36-GB hard drives to provide increased reliability and failover capability. This feature provides physical mirroring of the boot volume through firmware, thus providing extra reliability for the system's boot volume without burdening the host CPU.

To enable RAID-1 mirroring:

1. Power OFF the appliance (see "Shutting down and powering on the NAS Gateway 300" on page 87).
2. Attach a monitor, keyboard, and mouse to the first engine.
3. Ensure that there are two hard disk drives in the appliance engine.
4. Power ON the appliance.
5. When the LSI Logic BIOS starts and displays *Press CTRL-C to start LSI Logic Configuration Utility*, press **CTRL** and **C**.
6. Press **Enter** to select channel 1.
7. Select *Mirroring Properties* and press **Enter**.
8. Press the space bar to to change *No* to *Primary* in the column labeled *Mirrored Pair*.
9. Press **Esc**.
10. Select *Save changes then exit this menu* and press **Enter**. The drives will begin to synchronize.
11. Press **Esc**.
12. Select *Exit the Configuration Utility* and press **Enter**.
13. The engine will reboot automatically.
14. Repeat this process for the other engine.

Memory notes

The following sections contain information on adding memory.

Adding more engine memory to increase performance

You can enhance the performance of the NAS Gateway 300 in an NFS environment by adding more RAM to its processor. To do this:

1. Purchase either of the 5187 memory field-upgrade feature codes from your IBM representative:

   0301 1 GB memory upgrade

   0302 2 GB memory upgrade

2. Follow the instructions in Chapter 3, section “Replacing memory modules,” of the *Installation Guide*.
3. Before rebooting the appliance, attach a keyboard and display directly to the rear connectors of the product. During the first IPL, you will have to read and answer questions about the additional memory you have installed.

Using the Recovery CD-ROM if you have added more processor memory

If you have installed more processor memory, and later use the Recovery CD-ROM (see Chapter 9, “Using the Recovery and Supplementary CD-ROMs” on page 111), you will have to attach a keyboard and display and answer questions about the additional memory that you have installed.

Chapter 8. Troubleshooting

This chapter provides basic troubleshooting information to help you resolve some common problems that might occur with your appliance. Use Table 7 as an index to this information.

Table 7. Troubleshooting index

Shutting down and powering on the NAS Gateway 300

The clustering function requires special considerations when you need to shut down and power on the NAS Gateway 300. This section gives the details for those considerations.

Shutting down the NAS Gateway 300 when clustering is active

CAUTION:
<2–19> The power control button on the device and the power switch on the power supply do not turn off the electrical current supplied to the device. The device also might have more than one power cord. To remove all electrical current from the device, ensure that all power cords are disconnected from the power source.

Note: For translations of this safety notice, refer to *IBM TotalStorage Network Attached Storage Translated Safety Notices*, which can be found on the Documentation CD-ROM that was shipped with your appliance.

Attention: Powering off the appliance using the power button can result in a loss of data. Instead, it is recommended you use the following procedure to shut down.

1. Make note of the order in which you shut down the nodes.

   You shut the nodes down one at a time, and in the powering on procedure you start the nodes in the opposite order in which you shut them down.
2. On the node that you want to shut down last (the second node), click **Cluster Administration**, located in IBM NAS Admin, in the Cluster Tools folder. If prompted for a cluster name, enter the name of the cluster, and then click **Open**. Make sure that all resources are in the online state.

3. With all clustered resources in the online state, on the node that you want to shut down first (the first node), go to **Start → Shut Down** and select **Shut down** from the drop-down menu. Click **OK**.
4. On the second node, in Cluster Administrator, wait for all resources to failover to that node and return to the online state.
5. When all resources are in the online state, and the first node has shut down, on the second node go to **Start → Shutdown** and select **Shut down** from the drop-down menu. Click **OK**.
6. You can power off any network hubs or switches that are used exclusively by the NAS Gateway 300. If they are used by other network attached devices, do not power these off.
7. You can also power off any uninterruptible power supplies (UPS) that regulate power for the NAS Gateway 300, provided that no other equipment that you want to keep powered on is plugged into the same UPS.

Powering on the NAS Gateway 300 when clustering is active

1. Power on any UPS that you powered off in the powering off procedure, and allow it to return to normal operation.
2. Power on any network hubs or switches that you powered off in the powering off procedure.
3. Power on the node that you shut down ***last*** in the powering off procedure.
4. After the node comes up, start Cluster Administrator on that node and make sure that all resources are in an online state or shortly return to that state.
5. If no problems exist and all clustered resources are online, power on the node that you shut down ***first*** in the powering off procedure. Each resource for which that node is the preferred owner will fail back to that node and return to an online state.

   If a problem does exist, see the appropriate section of this chapter or Appendix D, “Symptom-to-part index” on page 123 for more information on identifying and resolving problems.

Diagnostic tools overview

The following tools are available to help you identify and resolve hardware-related problems:

- **Light-path diagnostics**

  LEDs help you identify problems with NAS Gateway 300 components. These LEDs are part of the light-path diagnostics that are built into your NAS Gateway 300. By following the *path of lights*, you can quickly identify the type of system error that occurred. See “Identifying problems using LEDs” on page 89 for more information.

- **Diagnostic programs and error messages**

  The diagnostic programs are stored in upgradable read-only memory (ROM) on the system board. These programs are the primary method of testing the major components of your NAS Gateway 300. See “Diagnostic programs” on page 93 for more information.

  Note: You must connect a keyboard, mouse, and monitor to your appliance in order to see error messages. If the engine does not recognize the monitor, keyboard, and mouse, reboot the engine while they are connected. If a Remote Supervisor Adapter is used for system management, the logs can

be accessed remotely. See “Shutting down and powering on the NAS Gateway 300” on page 87 for more information on shutting down and powering on the appliance.

- **POST beep codes, error messages, and error logs**

  The POST generates beep codes and messages to indicate successful test completion or the detection of a problem. See “POST” on page 91 for more information.

Identifying problems using LEDs

Each NAS Gateway 300 has LEDs to help you identify problems with some engine components. These LEDs are part of the light-path diagnostics built into the engine. By following the *path of lights*, you can identify the type of system error that occurred. There are three sources of LED information:

- Operator information panel
- Power supply
- Diagnostics panel

Operator information panel

The operator information panel on the front of the appliance contains status LEDs.

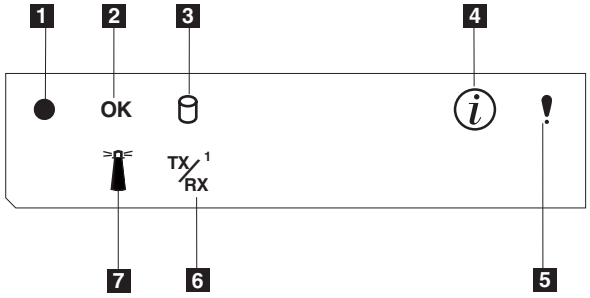

Figure 6. Operator information panel

1 **Power-on LED**: This green LED is on when system power is present in the appliance. When this LED flashes, the appliance is in standby mode (the system power supply is turned off and ac power is present). If this LED is not on, the power cord is not connected, the power supply has failed, or this LED has failed.

2 **POST complete (OK) LED**: This green LED is on when when POST completes without any errors.

3 **Hard disk drive activity LED**: This green LED is on when there is activity on the internal hard disk drive.

4 **Information LED**: This amber LED is on when the information log contains information about certain conditions in your appliance that might affect performance. For example, the LED is on if your appliance does not have functioning redundant power. An LED on the diagnostic panel on the system board will also be on.

5 **System error LED**: This amber LED is on when a system error occurs. An LED on the diagnostic panel on the system board will also be on to further isolate the error.

6 **Ethernet transmit/receive activity (TX/RX[1]) LED**: This green LED is on when there is transmit or receive activity to or from the appliance on the integrated Ethernet port.

7 **System locator LED**: This blue LED can be turned on remotely to identify a specific engine.

Power-supply LEDs

The AC Power LED and DC Power LED on the power supply provide status information about the power supply. Figure 7 show the location of the power-supply LEDs.

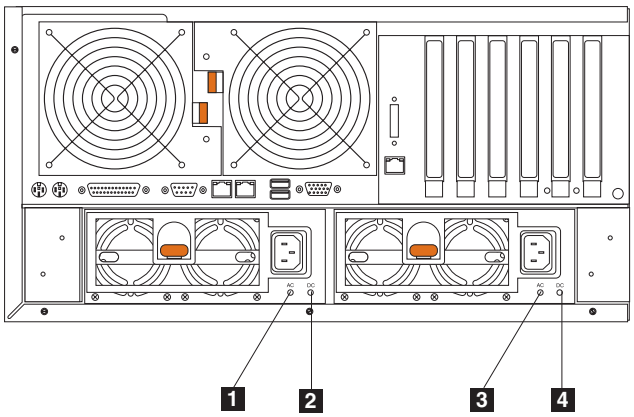

Figure 7. Location of the power-supply LEDs

1 Power supply 1 AC power LED (green)
2 Power supply 1 DC power LED (green)
3 Power supply 2 AC power LED (green)
4 Power supply 2 DC power LED (green)

Table 8. Power supply LED errors

AC Power LED	DC Power LED	Description and recommended actions
Off	Off	No power to system or ac problem **Action:** Check ac power to appliance.
On	Off	Standby mode or dc problem **Action:** 1. Check processor-board cable connector, J33. by moving the jumper on the J32 extension cable to pins 2 and 3 to bypass power control. If the DC Good LED comes on, press Ctrl-Alt-Delete. Watch the screen for any POST errors. 2. Check the System Error log for any listed problems. 3. If the appliance powers on with no errors: check the power switch assembly or system board. 4. Remove the adapters, and disconnect the cables and power connectors to all internal and external devices. Power on the system. If the DC Power LED comes on, replace the adapters and devices one at a time until you isolate the problem.
On	On	Power is OK

Diagnostics panel LEDs

You can use the light-path diagnostics to quickly identify the type of system error that occurred. The diagnostics panel is under the cover. Each engine is designed so that any LEDs that are on, remain on when the engine shuts down if the ac power

source is good and the power supplies can supply +5 V dc current to the engine. This feature helps isolate the problem if an error causes the engine to shut down.

Figure 8 shows the LEDs on the diagnostics panel.

Note: You might have to remove the cover to view these LEDs. (For more information on removing the cover, refer to “Removing the cover” in Chapter 3 of the hardware installation guide). See Table 22 on page 130 for more information on the LEDs.

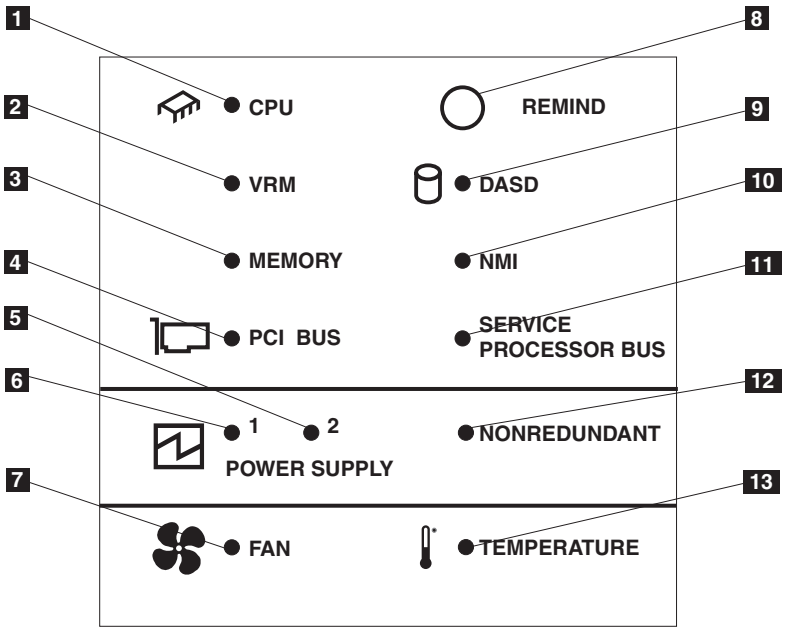

Figure 8. LED diagnostics panel

1. CPU LED
2. VRM LED
3. Memory LED
4. PCI Bus LED
5. Power supply 1 LED
6. Power supply 2 LED
7. Fan LED
8. Remind button
9. DASD LED
10. NMI LED
11. Service Processor Bus LED
12. Nonredundant LED
13. Temperature LED

POST

When you power on the NAS Gateway 300, it performs a series of tests to check the operation of its components and some of the options installed in the NAS Gateway 300. This series of tests is called the power-on self-test (POST).

One beep after the POST completes successfully indicates that the NAS Gateway 300 is operating correctly.

If the POST detects a problem, a series of beeps sound. See “Beep symptoms” on page 123 and “Error messages” for more information.

Notes:

1. If you have a power-on password or administrator password set, you must type the password and press **Enter**, when prompted, before POST will continue.
2. A single problem might cause several error messages. When this occurs, work to correct the cause of the first error message. After you correct the cause of the first error message, the other error messages usually will not occur the next time you run the test.

Error messages

Note: You must connect a keyboard, mouse, and monitor to your appliance in order to see error messages. If the engine does not recognize the monitor, keyboard, and mouse, reboot the engine while they are connected. If a Remote Supervisor Adapter is used for system management, the logs can be accessed remotely.

The table, “POST error codes” on page 140, provides information about the POST error messages that can appear during startup.

Event/error logs

The POST error log contains the three most recent error codes and messages that the system generated during POST. The System Event/Error Log contains all error messages issued during POST and all system status messages from the Remote Supervisor Adapter. In the event of a POST error, check the System Event/Error Log, as it will contain the most recent errors. This log often provides an indicator as to the nature of key failures.

To view the contents of the error logs, start the Configuration/Setup Utility program; then select **Event/Error Logs** from the main menu (see “Starting the Configuration/Setup Utility program” in Chapter 6 of the hardware installation guide that came with this appliance).

SCSI messages

If you receive a SCSI error message while using the SCSISelect Utility, use the following list to determine the cause of the error and the action to take.

One or more of the following conditions might be causing the problem:

- A failing SCSI adapter or drive
- An incorrect SCSI configuration
- Duplicate SCSI IDs in the same SCSI chain
- An incorrectly installed SCSI terminator
- A defective SCSI terminator
- An incorrectly installed cable
- A defective cable

Verify that:

- The external SCSI devices are powered on. External devices must be powered on before powering on the NAS Gateway 300.
- The cables or all external SCSI devices are connected correctly.
- The last device in each SCSI chain is terminated correctly.
- The SCSI devices are configured correctly.

You will receive these messages only when running the SCSISelect Utility.

Diagnostic programs

The NAS Gateway 300 diagnostic programs are stored in upgradable read-only memory (ROM) on the system board, on CD-ROMs, and in the software. These programs are the primary method of testing the major components of your NAS Gateway 300. Diagnostic error messages indicate that a problem exists; they are not intended to be used to identify a failing part.

The error messages and codes are listed in “Diagnostic error codes” on page 131.

Sometimes the first error to occur causes additional errors. In this case, the NAS Gateway 300 displays more than one error message. Always follow the suggested action instructions for the ***first*** error message that appears.

Error codes appear in the detailed test log and summary log when running the diagnostic programs.

Format of an error code

The error code format is as follows:

```
fff-ttt-iii-date-cc-text message
```

where:

fff is the three-digit function code that indicates the function being tested when the error occurred. For example, function code 089 is for the microprocessor.

ttt is the three-digit failure code that indicates the exact test failure that was encountered.

iii is the three-digit device ID.

date is the date that the diagnostic test was run and the error recorded.

cc is the check digit that is used to verify the validity of the information.

text message
is the diagnostic message that indicates the reason for the problem.

The diagnostic text message format is as follows:

```
Function Name: Result (test specific string)
```

where:

Function Name
is the name of the function being tested when the error occurred. This corresponds to the function code (*fff*) given in the previous list.

Result
can have one of these values:

Passed
The diagnostic test completed without any errors.

Failed The diagnostic test discovered an error.

User aborted
You stopped the diagnostic test before it was complete.

Not applicable
You specified a diagnostic test for a device that is not present.

Aborted
The test could not proceed because of the system configuration.

Warning
A possible problem was reported during the diagnostic test, such as a tested device was not installed.

Test specific string
Additional information that you can use to analyze the problem.

Starting the diagnostic programs

To start the diagnostic programs:

Note: You must connect a keyboard, mouse, and monitor to your appliance in order to see error messages. If the engine does not recognize the monitor, keyboard, and mouse, reboot the engine while they are connected. If a Remote Supervisor Adapter is used for system management, the logs can be accessed remotely.

1. Ensure you have connected a monitor, keyboard, and mouse to each NAS Gateway 300. When you have a monitor, keyboard, and mouse attached and POST completes successfully, one beep sounds. If the NAS Gateway 300 fails POST, a series of beeps sound (see “Beep symptoms” on page 123 for more details) and an error message appears on the monitor screen.

   Note: You might receive a configuration change informational message about a newly discovered mouse. Accept the message and continue.

2. Power on the NAS Gateway 300 and watch the screen.
3. When the message `F2 for Diagnostics` appears, press **F2**.
4. Type in the appropriate password; then, press **Enter**.

   Note: To run the diagnostic programs, you must start the NAS Gateway 300 with the highest-level password that is set. That is, if an administrator password is set, you must enter the administrator password, rather than the power-on password, to run the diagnostic programs.

5. Select either **Extended** or **Basic** from the top of the screen. (`PC-Doctor 2.0` with a copyright statement appears at the bottom of this screen.)
6. When the Diagnostic Programs screen appears, select the test you want to run from the list that appears; then, follow the instructions on the screen.

   Notes:

   a. Press **F1** while running the diagnostic programs to obtain Help information. Also press **F1** from within a help screen to obtain online documentation from which you can select different categories. To exit from Help and return to where you left off, press **Esc**.

   b. If the NAS Gateway 300 stops during testing and you cannot continue, restart the NAS Gateway 300 and run the diagnostic programs again.

   c. If you run the diagnostic programs with no mouse attached to your NAS Gateway 300, you will not be able to navigate between test categories using the Next Cat and Prev Cat buttons. All other functions provided by mouse-selectable buttons are also available using the function keys.

d. To view NAS Gateway 300 configuration information (such as system configuration, memory contents, interrupt request (IRQ) use, direct memory access (DMA) use, device drivers, and so on), select **Hardware Info** from the top of the screen.

e. You cannot use the Base System Diagnostics program to test adapters; use the procedure outlined in "Running adapter diagnostics" on page 103.

When the tests have completed, view the Test Log by selecting **Utility** from the top of the screen.

If the hardware checks out OK but the problem persists during normal NAS Gateway 300 operations, a software error might be the cause.

Viewing the test log

The test log will contain no information until the diagnostic tests have completed.

Note: If you already are running the diagnostic programs, begin with step 4.

To view the test log:

1. Ensure that a monitor, keyboard, and mouse are connected to each NAS Gateway 300.
2. Power on the NAS Gateway 300 and watch the screen.

   If the NAS Gateway 300 is on, shut down your operating system and restart the NAS Gateway 300.
3. When the message `F2 for Diagnostics` appears, press **F2**.

   If a power-on password or administrator password is set, the NAS Gateway 300 prompts you for it. Type in the appropriate password; then, press **Enter**.
4. When the Diagnostic Programs screen appears, select **Utility** from the top of the screen.
5. Select **View Test Log** from the list that appears; then, follow the instructions on the screen.

   The system maintains the test-log data while the NAS Gateway 300 is powered on. When you power off the power to the NAS Gateway 300, the test log is cleared.

You can also view the test log using the Remote Supervisor Adapter interface. See the *IBM Remote Supervisor Adapter User's Guide* provided on the Documentation CD-ROM.

Troubleshooting the Ethernet controller

This section provides troubleshooting information for problems that might occur with the Ethernet controller.

Network connection problems

If the Ethernet controller cannot connect to the network:

- Ensure that the NAS Gateway 300 is correctly connected to the Ethernet with a verified cable that has been built to the related Category 3, 4, or 5 unshielded twisted pair (UTP) standards.

  The network cable must be securely attached at all connections. If the cable is attached but the problem persists, try a different cable.

  If you set the Ethernet controller to operate at 100 Mbps, you must use Category 5 or better cabling.

If you directly connect two workstations (without a hub) or if you are not using a hub with X ports, use a crossover cable.

Note: To determine whether a hub has an X port, check the port label. If the label contains an "X," the hub has an X port.

- If you are connecting through an Ethernet hub or repeater, make sure that the signal LEDs are operational while the device is on and connected to the LAN.
- Determine whether the hub supports auto-negotiation. If not, configure the integrated Ethernet controller manually to match the speed and duplex mode of the hub.
- Check the Ethernet controller LEDs on the operator information panel.

  These LEDs indicate whether a problem exists with the connector, cable, or hub.
  - The Ethernet Link Status LED is on when the built-in Ethernet controller receives a LINK pulse from the hub. If the LED is off, there might be a bad connector or cable, or a problem with the hub.
  - The Ethernet Transmit/Receive Activity LED is on when the built-in Ethernet controller sends or receives data over the Ethernet network. If the Ethernet Transmit/Receive Activity LED is off, make sure that the hub and network are operating and that the correct device drivers are loaded.
- Make sure that you are using the correct device drivers, supplied with your NAS Gateway 300.
- Check for causes for the problem that are specific to the operating system.
- Make sure that the device drivers on the client and NAS Gateway 300 are using the same protocol.
- Test the Ethernet controller by running the NAS Gateway 300 diagnostic programs, as described in "Diagnostic programs" on page 93.

Gigabit Ethernet controller troubleshooting chart

Use the following troubleshooting chart to find solutions to Ethernet controller problems that have definite symptoms.

Table 9. Ethernet controller troubleshooting chart

Ethernet controller problem	Suggested action
Ethernet-link status light is not on.	• Ensure that the engine is powered on. • Check all connections at the Ethernet controller. • Check the cable. • If you manually configured the duplex mode, ensure that you also manually configured the speed. • Run Base System Diagnostics on the LEDs. If the problem remains, go to "Starting the diagnostic programs" on page 94 to run all of the diagnostic programs.
The Ethernet transmit/receive activity light is not on.	**Note:** The Ethernet Transmit/Receive Activity LED is on only when data is sent to or by this Ethernet controller. • Ensure that you have loaded the network device drivers. • The network might be idle. Attempt to send data from this workstation. • Run Base System Diagnostics on the LEDs. • The function of this LED can be changed by device-driver load parameters. If necessary, remove any LED parameter settings when you load the device drivers.
Data is incorrect or sporadic.	• Ensure that you are using Category 5 cabling when operating the NAS Gateway 300 at 100 Mbps or 1000 Mbps. • Make sure that the cables do not run close to noise-inducing sources, such as fluorescent lights.

Table 9. Ethernet controller troubleshooting chart (continued)

Ethernet controller problem	Suggested action
The Ethernet controller stopped working when another adapter was added to the NAS Gateway 300.	• Ensure that the cable is connected to the Ethernet controller. • Ensure that your PCI system BIOS is current. • Reseat the adapter. • Ensure that the adapter you are testing is supported by the NAS Gateway 300. If the problem remains, go to “Starting the diagnostic programs” on page 94 to run the diagnostic programs.
The Ethernet controller stopped working without apparent cause.	• Run Base System Diagnostics for the Ethernet controller. • Reinstall the device drivers. Refer to your operating-system documentation. If the problem remains, go to “Starting the diagnostic programs” on page 94 to run the diagnostic programs.

Troubleshooting adapters

This section explains how to troubleshoot adapters.

Ethernet adapters

Refer to these sections when troubleshooting your Ethernet adapters.

IBM Gigabit Ethernet SX Server Adapter troubleshooting chart

Use the troubleshooting chart in Table 10 to find solutions to the IBM Gigabit Ethernet SX Server Adapter problems that have definite symptoms.

Table 10. IBM Gigabit Ethernet SX Server Adapter troubleshooting chart

IBM Gigabit Ethernet SX Server Adapter problem	Suggested action
No Link or TX/RX activity.	If you cannot link to your switch: 1. Check the following LED lights on the adapter: **TX — On** The adapter is sending data. **RX — On** The adapter is receiving data. **Link — On** The adapter is connected to a valid link partner and is receiving link pulses. **Link — Off** Link is inoperative. • Check all connections at the adapter and link partner. • Make sure that the link partner is set to 1000 Mbps and full-duplex. • Ensure that the required drivers are loaded. **PRO — Programmable LED** Identifies the adapter by blinking. Use the Identify Adapter push-button in Intel PROSet II to control blinking. 2. Ensure that the cable is installed correctly. The network cable must be securely attached at all connections. If the cable is attached but the problem persists, try a different cable.
The appliance engine cannot find the IBM Gigabit Ethernet SX Server Adapter.	1. Verify that the adapter is seated firmly in the slot. 2. Try a different IBM Gigabit Ethernet SX Server Adapter.

Table 10. IBM Gigabit Ethernet SX Server Adapter troubleshooting chart (continued)

IBM Gigabit Ethernet SX Server Adapter problem	Suggested action
Diagnostics pass but the connection fails.	Ensure that the network cable is securely attached.
Another adapter stopped operating correctly after you installed the IBM Gigabit Ethernet SX Server Adapter.	1. Verify that the cable is connected to the IBM Gigabit Ethernet SX Server Adapter and not to another adapter. 2. Check for a resource conflict. 3. Ensure that both adapters are seated firmly in the slot. 4. Check all cables.
The adapter stopped working without apparent cause.	1. Reseat the adapter. 2. The network driver files might be damaged or deleted. Reinstall the drivers. 3. Use a different IBM Gigabit Ethernet SX Server Adapterr.
LINK LED is not on.	1. Ensure that you have loaded the adapter driver. 2. Check all connections at the adapter and the buffered repeater or switch. 3. Use another port on the buffered repeater or switch. 4. Ensure that the buffered repeater or switch port is configured for 1000 Mbps and full-duplex. 5. Change the auto-negotiation setting on the link partner, if possible.
RX or TX LED is not on.	1. Ensure that you have loaded the adapter driver. 2. The network might be idle; log in from a workstation. 3. The adapter is not transmitting or receiving data; use another adapter.

PRO/1000 XT Server Adapter by Intel troubleshooting chart

Use the troubleshooting chart in Table 11 to find solutions to the PRO/1000 XT Server Adapter by Intel problems that have definite symptoms.

Table 11. PRO/1000 XT Server Adapter by Intel troubleshooting chart

PRO/1000 XT Server Adapter by Intel problem	Suggested action
The NAS Gateway 300 cannot find the PRO/1000 XT Server adapter.	1. Verify that the adapter is seated firmly in the slot. 2. Reboot the NAS Gateway 300. 3. Try a different PRO/1000 XT Server Adapter by Intel.
Diagnostics pass but the connection fails.	1. Verify that the responding link is operating correctly. 2. Ensure that the network cable is securely attached. 3. Try a different cable.
Another adapter stopped operating correctly after you installed the PRO/1000 XT Server Adapter by Intel.	1. Verify that the cable is securely connected to the PRO/1000 XT Server Adapter by Intel and not to another adapter. 2. Check for a resource conflict as indicated by problem icons in Device Manager. To access Device Manager: a. Right-click on My Computer, and and select *Manage*. b. Select *Device Manager*. 3. Reload all PCI device drivers. 4. Ensure that both adapters are seated firmly in the slot. 5. Check all cables.

Table 11. PRO/1000 XT Server Adapter by Intel troubleshooting chart (continued)

PRO/1000 XT Server Adapter by Intel problem	**Suggested action**
PRO/1000 XT Server Adapter by Intel is unable to connect at 1000 Mbps, instead it connects at 100 Mbps.	1. Ensure that the cable is the correct type. 2. Try another cable.
The adapter stopped working without apparent cause.	1. Reseat the adapter. 2. The network driver files might be damaged or deleted. Reinstall the drivers using Device Manager. To access Device Manager: a. Right-click on My Computer, and and select *Manage*. b. Select *Device Manager*. c. Select *Update Drivers*. d. Reload the appropriate driver found in c:\drivers. 3. Reboot the NAS Gateway 300. 4. Try a different cable. 5. Use a different PRO/1000 XT Server Adapter by Intel.
LINK LED is off.	1. Ensure that you have loaded the adapter driver. 2. Check all connections at the adapter and the buffered repeater or switch. 3. Use another port on the buffered repeater or switch. 4. Ensure that the cable is securely attached. 5. Change the auto-negotiation setting on the link partner, if possible.
The link light is on, but communications are not correctly established.	1. Ensure that the latest drivers are loaded. 2. Ensure that the adapter and its link partner are set to either auto-negotiate or set to the same speed and duplex settings.
The ACT light is off.	1. Ensure that the drivers are loaded. 2. Try accessing a server. 3. Try another PRO/1000 XT Server Adapter by Intel. 4. Ensure that the cable is securely attached.
Windows 2000 message: Unable to remove PROSet in SAFE mode.	If the NAS Gateway 300 hangs after configuring the adapters with the PROSet utility, perform the following steps: 1. Start Windows in Safe mode. 2. Access the Device Manager and disable the network adapters and teams. 3. Restart the NAS Gateway 300. 4. Windows should operate normally if the disabled adapters were causing the problem.

Table 11. PRO/1000 XT Server Adapter by Intel troubleshooting chart (continued)

PRO/1000 XT Server Adapter by Intel problem	Suggested action
LED indicators	• ACT/LNK **Green on** The adapter is connected to a valid link partner. **Green flashing** Data activity is detected. **Off** No link is detected. **Yellow flashing** There is an identity problem. Use the Identify Adapter button in Intel PROSet II to control the blinking. See the PROSet online help for more information. • The color identifies the adapter data rate: 10=OFF 100=GREEN 1000=YLW **Off** The adapter is operating at a 10-Mbps data rate. **Green on** The adapter is operating at a 100-Mbps data rate. **Yellow on** The adapter is operating at a 1000-Mbps data rate

Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter troubleshooting chart

Use the troubleshooting chart in Table 12 to find solutions to the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter problems that have definite symptoms.

Table 12. Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter troubleshooting chart

Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter problem	Suggested action
Error message: “Setup cannot find OEMSETUP.INF or OEMSETNT.INF”	• Make sure that you are installing from the Adapters tab of the Network Control Panel Applet. • Make sure that you are specifying the correct drive letter for your CD-ROM drive.
Error message: “No SLIC adapters found”	Make sure that the adapter has been correctly installed in the PCI slot and that the PCI slot is enabled.
Event log reports “SLIC <*x*> Has determined that the adapter is not functioning properly”	• Run the diagnostics Network Control Panel Applet. • Replace the adapter.
Event log reports “SLIC <*x*>: Could not find an adapter”	• Ensure that the card was not moved from its original slot. • Uninstall the adapter from the adapters tab of the Network Control Panel Applet, reboot, and then reinstall.

Table 12. Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter troubleshooting chart (continued)

Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter problem	**Suggested action**
Event log reports "The Alacritech 100 or 1000 Series Server Adapter service failed to start due to the following error: The system cannot find the file specified."	Update the drivers with the drivers found at: www.ibm.com/storage/support
There is no connectivity with other computers on the network. Pinging does not work.	• Make sure that the cables are attached securely at both RJ-45 connections (adapter and switch) and that the network cable is functional. • Check the LEDs on the adapter. These LEDs help indicate if there are problems with the adapter, switch, or cable.
TCP connections can be established to other systems on the same subnet, but connections cannot be established to systems on the other side of a router.	If you have a Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter and an adapter from another manufacturer installed on your system, ***and*** your network is configured such that the route from your system to the remote system is different than the route from the remote system back to your system (a routing loop), then it might not be possible to establish a TCP connection. To solve this problem, perform one of the following solutions: • Replace the existing adapter with an Alacritech 100 or 1000 Series Server Adapter, or • Reconfigure the network to eliminate the routing loop, or • Disable TCP/IP Offload on the interfaces associated with the routing loop by completing this procedure: **Note:** Disabling SLI TCP/IP Offload will disable the performance benefits of the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter on that interface. Complete these steps only when required. 1. From the start menu, open the Network & Dial-up Connections folder located in the *Settings* menu. 2. Double-click the LAN connection of the interface you want to disable TCP Offload. 3. Click the Properties button. Clear the box labeled *Alacritech TCP Fast-path driver*. Click **OK**. 4. Restart the engine.
IPSEC does not function through the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter interface.	In order to use IPSEC with the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter, disable SLIC TCP/IP Offload on the interface through which you want to establish the IPSEC connection. To disable the SLIC TCP/IP Offload, complete these steps: 1. From the start menu, open the Network & Dial-up Connections folder located in the *Settings* menu. 2. Double-click the LAN connection of the interface you want to disable TCP Offload. 3. Click the Properties button. Clear the box labeled *Alacritech TCP Fast-path driver*. Click **OK**. 4. Restart the engine.

Table 12. Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter troubleshooting chart (continued)

Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter problem	**Suggested action**
QOS does not work through the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter interface	In order to use QOS with the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter, disable SLIC TCP/IP Offload on the interface through which you want to establish the QOS connection. To disable the SLIC TCP/IP Offload, complete these steps: 1. From the start menu, open the Network & Dial-up Connections folder located in the *Settings* menu. 2. Double-click the LAN connection of the interface you want to disable TCP Offload. 3. Click the Properties button. Clear the box labeled *Alacritech TCP Fast-path driver*. Click **OK**. 4. Restart the engine.
Network monitoring does not function through the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter.	Network monitoring applications require packet access to the network, and are incompatible with the session layer interface provided by the SLIC TCP/IP Offload. In order to use network monitoring with the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter, disable SLIC TCP/IP Offload on the interface through which you want to use network monitoring. To disable the SLIC TCP/IP Offload, complete these steps: 1. From the start menu, open the Network & Dial-up Connections folder located in the *Settings* menu. 2. Double-click the LAN connection of the interface you want to disable TCP Offload. 3. Click the Properties button. Clear the box labeled *Alacritech TCP Fast-path driver*. Click **OK**. 4. Restart the engine.
Firewall software does not work through the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter interface.	Firewall applications require packet access to the network, and are incompatible with the session layer interface provided by the SLIC TCP/IP Offload. In order to use network monitoring with the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter, disable SLIC TCP/IP Offload on the interface through which you want to use for a firewall. To disable the SLIC TCP/IP Offload, complete these steps: 1. From the start menu, open the Network & Dial-up Connections folder located in the *Settings* menu. 2. Double-click the LAN connection of the interface you want to disable TCP Offload. 3. Click the Properties button. Clear the box labeled *Alacritech TCP Fast-path driver*. Click **OK**. 4. Restart the engine.
Microsoft Network Load Balancing does not work through the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter interface.	In order to use Microsoft Network Load Balancing with the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter, disable SLIC TCP/IP Offload on the interfaces which you want to use for Microsoft Network Load Balancing. To disable the SLIC TCP/IP Offload, complete these steps: 1. From the start menu, open the Network & Dial-up Connections folder located in the *Settings* menu. 2. Double-click the LAN connection of the interface you want to disable TCP Offload. 3. Click the Properties button. Clear the box labeled *Alacritech TCP Fast-path driver*. Click **OK**. 4. Restart the engine.

Table 12. Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter troubleshooting chart (continued)

Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter problem	**Suggested action**
Point-to-Point Tunneling Protocol (PPTP) does not work through the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter interface.	PPTP connections cannot be established through a SLIC TCP/IP Offload-enabled interface. In order to use PPTP with the Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter, disable SLIC TCP/IP Offload on the interfaces through which you want to establish a PPTP connection. To disable the SLIC TCP/IP Offload, complete these steps: 1. From the start menu, open the Network & Dial-up Connections folder located in the *Settings* menu. 2. Double-click the LAN connection of the interface you want to disable TCP Offload. 3. Click the Properties button. Clear the box labeled *Alacritech TCP Fast-path driver*. Click **OK**. 4. Restart the engine.
When uninstalling the Alacritech TCP Fast-path driver, an error message appears with the message: "Could not uninstall the Alacritech TCP Fast-path driver component. The error is 0x800F020B"	No corrective action is necessary.

Alacritech 100x4 Quad-Port Server Accelerated Adapter

Table 13 displays the LED definitions for the Alacritech 100x4 Quad-Port Server Accelerated Adapter LED definitions.

Table 13. Alacritech 100x4 Quad-Port Server Accelerated Adapter LED definitions

LED	**Indication**	**Meaning**
LNK	Off	Either the adapter or the switch (or both) are not receiving power, or the cable connection between them is faulty.
	Green	The adapter and switch are receiving power. The cable connection between them is good. A 100 Mbps link has been established.
	Amber	The adapter and switch are receiving power. The cable connection between them is good. A 10 Mbps link has been established.
ACT	Off	The adapter is not sending or receiving network data.
	Flashing amber	The adapter is sending or receiving network data.

Running adapter diagnostics

This section describes how to test the adapter using the diagnostics.

Note: Running the adapter diagnostics will disrupt the network connection.

Ethernet adapters

This section explains how to test Ethernet adapters.

IBM Gigabit Ethernet SX Server Adapter, PRO/1000 XT Server Adapter by Intel: The NAS Gateway 300 is equipped with the Intel PROSet II utility for:

- Monitoring the status of the Ethernet adapter PCI cards
- Testing the Ethernet adapter to see if there are any problems with the adapter hardware, the cabling, or the network connection
- Isolating problems during troubleshooting

To access the PROSet utility, you must first go into Terminal Services. For instructions on how to invoke Terminal Services, see "Terminal Services and the IBM NAS Administration console" on page 15. Within Terminal Services perform the following steps:

1. Double-click the **Intel PROSet II** icon in the Control Panel to start the Intel PROSet II utility.
2. In the Intel PROSet II utility, select the Ethernet adapter you want to test (IBM Gigabit Ethernet SX Server Adapter or PRO/1000 XT Server Adapter by Intel).
3. Select the **Diagnostics** tab. A list of available tests is displayed.
4. Select **Run Tests**.

   You can also select or clear individual tests with the check boxes. If an error is detected, information about the error is displayed.
5. Repeat Steps 2 through 4 for each Ethernet adapter installed in the NAS Gateway 300.

Alacritech 100x4 Quad-Port Server Accelerated Adapter and Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter: Note that running these diagnostic tests will disrupt all ports of the adapter.

Use the following procedure to run diagnostic tests on this adapter:

1. Open the Control Panel.
2. In the Network dialog box, select the *Adapters* tab.
3. Select one of the Alacritech 100x4 Quad-Port Server Accelerated Adapter or Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter ports.
4. Click **Properties**.
5. Select the *Diagnostics* tab.
6. Click **Run**. The adapter will complete the self-test of the IPP and the selected port of the adapter.
7. Close the Network Control Panel by clicking **OK**.

Fibre Channel adapter

Note: The Fibre Channel adapter diagnostics utility that was used in earlier versions was FAStT Check.

Each NAS Gateway 300 has FAStT MSJ for viewing the status of your Fibre Channel connection as well as testing the adapter. To use FAStT MSJ utility, you must first go into Terminal Services. For instructions on how to invoke Terminal Services, see "Terminal Services and the IBM NAS Administration console" on page 15.

Access FAStT MSJ by going into the IBM NAS administrator console and selecting **NAS Management** > **Storage** > **NAS Utilities** > **FAStT MSJ**. Then select **Connect**. A diagnostic panel displays the following general information related to your Fibre Channel adapter, which can be useful if you need to place a support call:

- Node name (the worldwide name or MAC address of the Fiber Channel adapter)

- Loop ID
- BIOS version
- Firmware version number
- Device-driver version number
- PCI slot number

Note: Ensure that there is no adapter activity before running the test or data can be lost.

To test the Fibre Channel adapter, select the adapter and then select **Diagnostic**.

Remote Supervisor Adapter

1. Insert the Remote Supervisor Adapter Support CD-ROM into the CD-ROM drive and restart the NAS Gateway 300. If the NAS Gateway 300 does not boot from the CD-ROM, use POST/BIOS setup to configure the CD-ROM drive as a boot device.

   After the NAS Gateway 300 boots, the main menu displays.
2. Use the Up and Down arrow keys to select **Hardware Status and Information** and press **Enter**. The Hardware Status and Information menu reports on the Advanced System Management devices in the NAS Gateway 300 in the following manner:

```
       System Management Processor Communication   :  Passed
 ->    Built in  Self Test Status  ......    : Passed
       Boot Sector Code Revision ...         :16, Build ID: CNET15A
       Main Application Code Revision        :16, Build ID: TUET15A
```

Testing the connection between two NAS Gateway 300s

Verify that the LINK OK lights for both NAS Gateway 300s are on. The LINK OK light is on the operator information panel.

Power checkout

Power problems can be difficult to troubleshoot. For example, a short circuit can exist anywhere on any of the power-distribution busses. Usually a short circuit causes the power subsystem to shut down because of an overcurrent condition.

A general procedure for troubleshooting power problems is as follows:

1. Shut down (see “Shutting down and powering on the NAS Gateway 300” on page 87) the system and disconnect the ac cords.
2. Check for loose cables in the power subsystem. Also check for short circuits; for example, a loose screw could cause a short circuit on a circuit board.
3. Remove adapters and disconnect the cables and power connectors to all internal and external devices until the NAS Gateway 300 is at minimum configuration required for power-on (see “Minimum operating requirements” on page 152).
4. Reconnect the ac cord and power on the NAS Gateway 300. If it powers up successfully, replace adapters and devices one at a time until the problem is isolated.

For specific power problems, see “Power error messages” on page 148.

Replacing the battery

IBM has designed this product with your safety in mind. The lithium battery must be handled correctly to avoid possible danger. If you replace the battery, you must adhere to the following safety guidelines.

CAUTION:
<2–16> When replacing the lithium battery, use only IBM Part Number 33F8354 or an equivalent type battery recommended by the manufacturer. If your system has a module containing a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of.

Do not:
- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble

Dispose of the battery as required by local ordinances or regulations.

Note: For translated versions of this safety notice, refer to *Translated Safety Notices*, which is on the publications CD-ROM that was shipped with your appliance.

Note: In the U.S., call 1 800-IBM-4333 for information about battery disposal.

If you replace the original lithium battery with a heavy-metal battery or a battery with heavy-metal components, be aware of the following environmental consideration. Batteries and accumulators that contain heavy metals must not be disposed of with normal domestic waste. They will be taken back free of charge by the manufacturer, distributor, or representative, to be recycled or disposed of in a correct manner. To order replacement batteries, call 1 800-772-2227 within the United States.

Before you begin:
- Review the information in the section “Before you begin” in Chapter 3 of the hardware installation guide that came with this appliance, and any special handling and installation instructions supplied with the replacement battery.
- Attach a monitor and keyboard.

  Note: After you replace the battery, you must reconfigure your appliance and reset the system date and time.

To replace the battery:
1. Shut down (see “Shutting down and powering on the NAS Gateway 300” on page 87) the appliance and peripheral devices, and disconnect all external cables and power cords.
2. Remove the cover (see “Removing the cover” in Chapter 3 of the hardware installation guide that came with this appliance).

3. Lift the adapter-retention clip on top of the adapter-support bracket (see “Adapter-support bracket” in Chapter 3 of the hardware installation guide that came with this appliance).
4. Remove all of the full-length adapters and plastic dividers (refer to “Adapters” in Chapter 3 of the hardware installation guide that came with this appliance).

   Attention: Note the location of the adapters. You must replace each adapter in the same slot from which it was removed.

   Note: You do not need to unplug the internal adapter cables.
5. Locate the battery (connector BH1) on the system board using Figure 9.

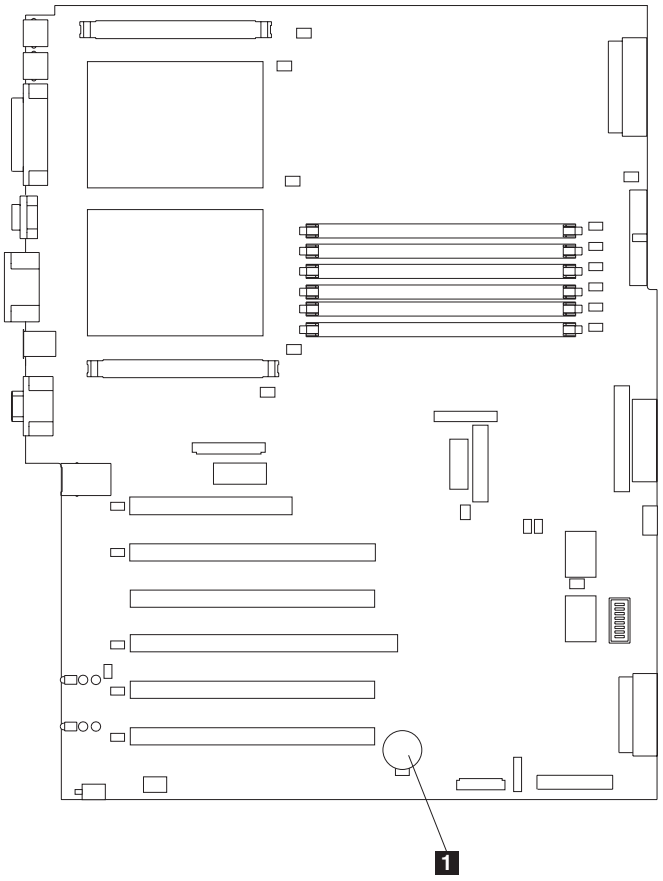

1 Battery

Figure 9. Replacing the battery

6. Remove the plastic cover on the system board.
7. Remove the battery:
   a. Use one finger to press the top of the battery clip away from the battery until the battery releases upward from the socket as shown in Figure 10 on page 108.
   b. Lift and remove the battery from the socket.

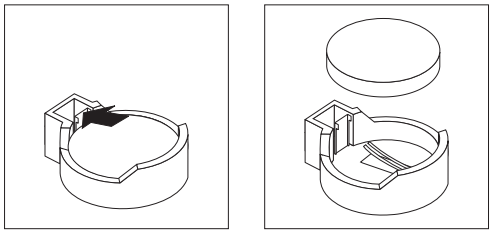

Figure 10. Releasing the battery

8. Insert the new battery:

   Note: Ensure that the polarity of the battery is correct. The positive (+) side must face up.

   a. Tilt the battery so that you can insert it into the socket on the side opposite the battery clip.
   b. Press the battery down into the socket until it clicks under the battery clip, as shown in Figure 11.

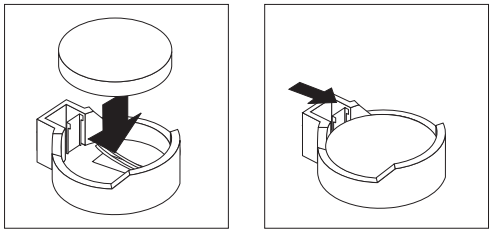

Figure 11. Inserting the new battery

9. Reinstall the adapters and plastic dividers in the same slots from where they were removed, and reconnect any internal cables that were disconnected (see "Adapters" in Chapter 3 of the hardware installation guide that came with this appliance).
10. Replace the adapter-support bracket (see "Adapter-support bracket" in Chapter 3 of the hardware installation guide that came with this appliance).

    Note: You must reinstall the air-baffle assembly when you reinstall the adapter-support bracket. Make sure that no cable is under the adapter-support bracket or interferes with the center fans.

    Important: To ensure correct cooling and engine operation, you must keep the air-baffle cover closed.
11. Reinstall the engine cover, and connect the power cords and external cables.

    Note: Wait approximately 20 seconds after you plug the power cords of your appliance into electrical outlets for the power button to become active.
12. Power on the appliance.
13. Start the Configuration/Setup Utility program and set configuration parameters:
    - Set the system date and time
    - Set the power-on password
    - Reconfigure your appliance

Temperature checkout

Cooling of the NAS Gateway 300 is important for correct operation and reliability. Ensure that:

- Each of the drive bays has either a drive or a filler panel installed.
- Each of the power-supply bays has either a power supply or a filler panel installed.
- The top cover is in place during normal operation.
- There is at least 50 mm (2 in.) of ventilated space at the sides of the NAS Gateway 300 and 100 mm (4 in.) at the rear of the NAS Gateway 300.
- The top cover is removed for no longer than 30 minutes while the NAS Gateway 300 is operating.
- The processor housing cover covering the processor and memory area is removed for no longer that 10 minutes while the NAS Gateway 300 is operating.
- A removed hot-swap drive is replaced within 2 minutes of removal.
- Cables for optional adapters are routed according to the instructions provided with the adapters (ensure that cables are not restricting air flow).
- The fans are operating correctly and the air flow is good.
- A failed fan is replaced within 48 hours.

In addition, ensure that the environmental specifications for the NAS Gateway 300 are met. Refer to “Products and features” in Chapter 2 of the hardware installation guide.

For more information on specific temperature error messages, see “Temperature error messages” on page 150.

Recovering BIOS

If your BIOS has become corrupted, such as from a power failure during a flash update, you can recover your BIOS using the recovery boot block and a BIOS flash diskette.

Note: You can obtain a BIOS flash diskette from one of the following sources:

- Download a BIOS flash diskette from the Web site at:
  www.ibm.com/storage/support/nas
- Contact your IBM service representative.

The flash memory of the NAS Gateway 300 contains a protected area that cannot be overwritten. The recovery boot block is a section of code in this protected area that enables the NAS Gateway 300 to start up and to read a flash diskette. The flash utility recovers the system BIOS from the BIOS recovery files on the diskette.

To recover the BIOS:

1. Shut down (see “Shutting down and powering on the NAS Gateway 300” on page 87) the NAS Gateway 300 and peripheral devices. Disconnect all external cables and power cords. Remove the cover.
2. Locate the boot-block jumper block (J28) on the system board.
3. Place a jumper on pins 2 and 3 to enable BIOS backup page.
4. Insert the BIOS flash diskette into the diskette drive.
5. Reconnect all external cables and power cords.

6. Restart the NAS Gateway 300.
7. After the NAS Gateway 300 completes the POST, select **1 -- Update POST/BIOS** from the menu that contains various flash (update) options.
8. When you are asked if you would like to move the current POST/BIOS image to the backup ROM location, type **N**.

   Attention: Typing **Y** will copy the corrupted BIOS into the secondary page.
9. When you are asked if you would like to save the current code on a diskette, select **N**.
10. You will be asked to select the language that you want to use. Select your language (**0 - 7**) and press **Enter** to accept your choice. You will be prompted to remove the diskette and press **Enter** to restart the system. Remove the flash diskette from the diskette drive.
11. Shut down (see “Shutting down and powering on the NAS Gateway 300” on page 87) the NAS Gateway 300.
12. Remove the jumper on the boot-block jumper block, or move it to pins 1 and 2 to return to normal startup mode.
13. Restart the NAS Gateway 300. The NAS Gateway 300 should start up normally.

Chapter 9. Using the Recovery and Supplementary CD-ROMs

Attention: Changing the preloaded software configuration of this product, including applying or installing unauthorized service packs or updates to preinstalled software, or installing additional software products that are not included in either the preloaded image or on the Supplementary CD-ROM might not be supported and could cause unpredictable results. For updated compatibility information, see:

www.ibm.com/storage/support/nas

To correct problems with a preloaded software component, back up your user and system data. Then, use the Recovery CD-ROM Set to restore the preloaded software image.

This chapter describes the applications included on the Supplementary and Recovery CD-ROMs, and how and when you should use them.

As an alternative to using the Recovery CD-ROM Set, you can use the restore portion of the disaster recovery solution provided by Persistent Storage Manager (PSM) to recover the node. The restore function allows you to restore the node to the state it was in at the time of the PSM backup in one step. You will not have to revert back to the original (factory) configuration, which would require you to subsequently reconfigure clustering and other components. See “Restoring the system drive using the PSM recovery diskette” on page 67 to determine whether you meet the requirements to use PSM. If you have not met the requirements for using the PSM recovery method, or if the PSM recovery fails, then you must use the Recovery CD-ROM Set as described in this chapter.

Using the Recovery Enablement Diskette and Recovery CD-ROM

The Recovery CD-ROM Set (four CD-ROMs, labeled as “Recovery CD-ROM 1,” “Recovery CD-ROM 2,” “Recovery CD-ROM 3,” and “Recovery CD-ROM 4”) contains the preload image for this appliance and is used to recover the preloaded image on either node. You must start the (failed) node using the Recovery Enablement Diskette before you can boot from Recovery CD-ROM 1.

Attention: The Model G27 does not have a monitor, keyboard, or mouse attached to it under normal operating conditions. Because of this, you cannot interact with the preload-image restore process using a monitor. Starting Recovery CD-ROM 1 will, without visually prompting the user, automatically destroy all data on the system drive. Use the Recovery Enablement Diskette and Recovery CD-ROM Set only when it is absolutely necessary to restore the preloaded system image.

To recover the preloaded image on a (failed) node, perform the following steps:

1. Delete all persistent images to ensure a clean reload of the system software.

   Note: The recovery process invalidates persistent images and leaves them in a state that is inconsistent with their pre-recovery state.
2. On the other (operational) node of the Model G27, select **Cluster Administration**, located in the Cluster Tools folder in the IBM NAS Admin. If prompted for a cluster name, enter the name of the cluster, and then click **Open**.
3. The cluster name appears in the left panel. Underneath it, locate the name of the failed node, right-click the failed node machine name, and select **Evict**

Node. The name of the failed node will be removed from the left pane, and the cluster will now contain only the operational node of the Model G27.

4. Attach a keyboard and display to the failed node.
5. Insert the Recovery Enablement Diskette into the diskette drive and place Recovery CD-ROM 1 into the CD-ROM drive of the failed node.

> **Important**
> The Recovery Enablement Diskette enables the Model G27 to start from the CD-ROM drive. You will not be able to restore the preload image from the Recovery CD-ROM Set without first restarting the appliance using the Recovery Enablement Diskette.

6. Restart the appliance.
7. If you installed additional processor memory on the appliance, the BIOS configuration program will appear. Click **Continue** on the first panel, click **Continue** again, click **Exit Setup**, and finally, click **Yes, save and exit Setup**.
8. When the diskette loads, you are prompted with a message which asks if you want to proceed. Type *Y* to proceed. If you type *N*, you will be returned to a command prompt.
9. The recovery process will begin automatically. Follow the instructions provided by the image restoration software, and the original manufacturing preload will be restored. During the restoration of the preload, you are prompted to insert the other recovery CD-ROMs into the CD-ROM drive. When the preload image is restored, the Model G27 restarts automatically. You may now remove Recovery CD-ROM 3 from the CD-ROM drive.
10. If you installed additional processor memory, the BIOS configuration program will now appear a second time. Click **Continue** on the first panel, click **Continue** again, click **Exit Setup**, and finally, click **Yes, save and exit Setup**. You can now detach the keyboard and display from the failed node and allow the recovery process to complete automatically.

> **Important**
> After the node restarts, a series of configuration and system preparation programs that finish configuring the node run automatically. These programs must finish running before you use any included applications (such as the IBM Advanced Appliance Configuration Utility or the Terminal Services Client) to connect to or configure the Model G27. Do not connect to or configure the node for at least 15 minutes after system restart. This notice applies only to the first time the Model G27 is started after using the Recovery CD-ROM Set.
>
> Logical Disk 0 will be configured to have a 6-GB NTFS boot partition. Any other previously configured logical disk drives, as well as the remainder of Logical Disk 0 (which, on the original hard disk drive of the node, contains the Maintenance partition, but for a replacement hard disk drive would not contain any other partitions), will be unchanged.

11. Reinstall all software updates that you had installed on the node. Or, if the Recovery CD-ROM Set that you used in this procedure is a newer version than the one you received with the Model G27, reinstall only those software updates that are newer than those on the Recovery CD-ROM Set.

12. If you are using the recovery procedure to restore the failed node after replacing the internal hard disk drive, continue with this step. Otherwise, go to Step 13. You must now rebuild the Maintenance (D) partition on the new hard disk drive, as the recovery process only rebuilds the System (C) partition.

    Start Disk Management on the failed node. You can do this in one of two ways:

    a. Start a Terminal Services session to the node and click the **IBM NAS Admin** icon. From the IBM NAS Administration console that appears, select **Storage** and then **Disk Management**.
    b. Start a Windows 2000 for NAS user interface session to the node and select **Disks and Volumes**. Select **Disks and Volumes** again, and then provide your administrator user name and password when prompted.

    When Disk Management has started, perform these steps:

    a. In the Disk Management window, right-click the unallocated area of Disk 0, and then click **Create Partition**.
    b. In the Create Partition wizard, click **Next** and select **Primary Partition**.
    c. Click **Next** and select **D:** as the drive letter.
    d. Click **Next** and select **FAT32** as the file system. Change the volume label to *Maintenance*.
    e. Click **Finish** to close the wizard. The partition is then formatted. When formatting is complete, the status of the partition should appear as *Healthy*, and the other properties should appear as:
        - Name: *Maintenance*
        - Drive letter: *D*
        - File system: *FAT32*

13. On the failed (now recovered) node, follow the procedures for configuring a joining node located in “Roadmap for setting up and configuring the NAS Gateway 300” on page 3. The recovered node rejoins the cluster that already contains the other (operational) node. You will also need to reconfigure any cluster resource balancing you had set up prior to recovery, so that the recovered node will once again be the preferred owner of any resources for which it had been the preferred owner prior to the recovery.

Using the Supplementary CD-ROM

The Supplementary CD-ROM contains documentation and copies of key software applications that are preinstalled on the NAS Gateway 300. Table 14 and Table 15 include the names of the directories found on the Supplementary CD-ROMs and a description of the contents of the directory.

Table 14. Supplementary CD-ROM 1 directories

Directory name	Contents
DB2	• EnableDB2Support.exe • DisableDB2Support.exe These files enable and disable support for Linux- and Solaris-based DB2® clients using NFS shares.
DiskImages	This directory contains a diskette image for the Recovery Enablement Diskette and a diskette image for a bootable diskette that automatically configures the ServeRAID controller and drives. To create the Recovery Enablement Diskette, run enablement_diskette25.exe and insert a HD 1.44 floppy diskette into drive A: when prompted. To create the bootable diskette to automatically configure the ServeRAID controller and drives, run IBM_NAS_AutoRAID_diskette_2.5.EXE and insert a HD 1.44 floppy diskette into drive A: when prompted. **Note:** The contents of this directory are not used on the NAS Gateway 300.
diskpart Samples	This directory contains an example script for use with the DiskPart utility. This script will clean disk 2, convert it to dynamic, partition it, and assign drive letters to the partitions. This script is unsupported and should be used with extreme caution.
IBM Advanced Appliance Configuration	Run Setup.exe on the machine from which you will administer the appliance. The agent is preinstalled on the appliance.
IBM NAS Extensions For IBM Director	The IBM NAS extensions to IBM Director provide capabilities to IBM Director that are specific to the IBM NAS appliances.
Zip Tools	This directory contains compression tools used for sending information to IBM technical support.
readme.txt	This text file describes the contents of the Supplementary CD-ROMs.

Table 15. Supplementary CD-ROM 2 directories

Directory name	Contents
AoP	This is the add-on pack for the Server Appliance Kit.
SFU_2073.1	Microsoft® Services for UNIX (SFU) Version 2.2 support files: • QFE 320175 for performance enhancements • QFE 321096 for SAK and SFU performance enhancements readme_SFN5.txt: Instructions for installing Microsoft File and Print Services for NetWare 5.0.
Terminal Services Client	Microsoft Terminal Services Client install files.
w2ksp2	Windows 2000 Service Pack 2.

Appendix A. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering the subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION ″AS IS″ WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or program(s) described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Trademarks

IBM, the IBM logo, ServeRAID, DB2, ServerGuide, TotalStorage, NetView, SecureWay, Tivoli and WebSphere are trademarks or registered trademarks of International Business Machines Corporation or Tivoli Systems Inc. in the United States, other countries, or both.

Alacritech and SLIC Technology are registered trademarks of Alacritech, Inc. in the United States, other countries, or both.

Intel, LANDesk, MMX, Pentium, Pentium II Xeon, Itanium, and Pentium III Xeon are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Lotus and Domino are trademarks of Lotus Development Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

NetWare is a trademark of Novell, Inc.

Persistent Storage Manager is a trademark of Columbia Data Products, Inc.

UNIX is a registered trademark in the United States, other countries, or both, and is licensed exclusively through X/Open Company Ltd.

Other company, product, and service names may be trademarks or service marks of others.

Appendix B. Getting help, service, and information

If you need help, service, technical assistance, or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you.

IBM maintains pages on the World Wide Web where you can get information about IBM products and services and find the latest technical information.

Table 16 lists some of these pages.

Table 16. IBM Web sites for help, services, and information

www.ibm.com	Main IBM home page
www.ibm.com/storage	IBM Storage home page
www.ibm.com/storage/support	IBM Support home page

You might also want to visit the Web pages of other companies for information about other operating systems, software, and accessories. The following are some other Web sites you might find helpful:

www.tivoli.com
www.cdpi.com

Services available and telephone numbers listed are subject to change without notice.

Service support

With the original purchase of an IBM hardware product, you have access to extensive support coverage. During the IBM hardware product warranty period, you may call the IBM Support Center (1 800 426-7378 in the U.S.) for hardware product assistance covered under the terms of the IBM hardware warranty.

The following services are available during the warranty period:

- Problem determination: Trained personnel are available to assist you with determining if you have a hardware problem and deciding what action is necessary to fix the problem.
- IBM hardware repair: If the problem is determined to be caused by IBM hardware under warranty, trained service personnel are available to provide the applicable level of service.
- Engineering change management: Occasionally, there might be changes that are required after a product has been sold. IBM or your reseller, if authorized by IBM, will make Engineering Changes (ECs) available that apply to your hardware.

Be sure to retain your proof of purchase to obtain warranty service.

Please have the following information ready when you call:

- Machine type and model
- Serial numbers of your IBM hardware products
- Description of the problem
- Exact wording of any error messages
- Hardware and software configuration information

If possible, be at your NAS device when you call.

A compatible monitor, keyboard, and mouse are required for many service activities. Before you have the NAS device serviced, be sure to attach these components to the device, either directly or indirectly through a console switch.

The following items are not covered:

- Replacement or use of non-IBM parts or nonwarranted IBM parts

  Note: All warranted parts contain a 7-character identification in the format IBM FRU XXXXXXX.
- Identification of software problem sources
- Configuration of BIOS as part of an installation or upgrade
- Changes, modifications, or upgrades to device drivers
- Installation and maintenance of network operating systems (NOSs)
- Installation and maintenance of application programs

Refer to your IBM hardware warranty for a full explanation of IBM's warranty terms.

Before you call for service

Some problems can be solved without outside assistance, by using the online help, by looking in the online or printed documentation that comes with your network-attached storage appliance, or by consulting the support Web page noted in Table 16 on page 117. Also, be sure to read the information in any README files that come with your software.

Your network-attached storage appliance comes with documentation that contains troubleshooting procedures and explanations of error messages. The documentation that comes with your appliance also contains information about the diagnostic tests you can perform.

If you receive a POST error code or beep code when you turn on your Network Attached Server appliance, refer to the POST error-message charts in your hardware documentation. If you do not receive a POST error code or beep code, but suspect a hardware problem, refer to the troubleshooting information in your hardware documentation or run the diagnostic tests.

If you suspect a software problem, consult the documentation (including any README files) for the operating system or application program.

Getting customer support and service

Purchasing an IBM network-attached storage appliance entitles you to standard help and support during the warranty period. If you need additional support and services, a wide variety of extended services are available for purchase that address almost any need; see Appendix C, "Purchasing additional services" on page 121 for information.

Getting help online: www.ibm.com/storage/support

Be sure to visit the support page that is specific to your hardware, complete with FAQs, parts information, technical hints and tips, technical publications, and downloadable files, if applicable. This page is at: www.ibm.com/storage/support.

Getting help by telephone

With the original purchase of an IBM hardware product, you have access to extensive support coverage. During the IBM hardware product warranty period, you may call the IBM Support Center (1 800 426-7378 in the U.S.) for hardware product assistance covered under the terms of the IBM hardware warranty. Expert technical-support representatives are available to assist you with questions you might have on the following:

- Setting up your network-attached storage appliance
- Arranging for service
- Arranging for overnight shipment of customer-replaceable parts

In addition, if you purchased a network-attached storage appliance, you are eligible for IBM up and running support for 90 days after installation. This service provides assistance for:

- Setting up your network-attached storage appliance
- Limited configuration assistance

Please have the following information ready when you call:

- Machine type and model
- Serial numbers of your IBM hardware products, or your proof of purchase
- Description of the problem
- Exact wording of any error messages
- Hardware and software configuration information

If possible, be at your computer when you call.

In the U.S. and Canada, these services are available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9:00 a.m. to 6:00 p.m. In all other countries, contact your IBM reseller or IBM marketing representative.[3]

3. Response time will vary depending on the number and complexity of incoming calls.

Appendix C. Purchasing additional services

During and after the warranty period, you can purchase additional services, such as support for IBM and non-IBM hardware, operating systems, and application programs; network setup and configuration; upgraded or extended hardware repair services; and custom installations. Service availability and name might vary by country.

Warranty and repair services

You can upgrade your standard hardware warranty service or extend the service beyond the warranty period.

Warranty upgrades in the U.S. include:

- On-site service to premium on-site service

  If your warranty provides for on-site service, you can upgrade to premium on-site service (4-hour average on-site response, 24 hours a day, 7 days a week).

You also can extend your warranty. Warranty and Repair Services offers a variety of post-warranty maintenance options. Availability of the services varies by product.

For more information about warranty upgrades and extensions:

- In the U.S., call 1-800-426-4343.
- In Canada, call 1-800-465-7999.
- In all other countries, contact your IBM reseller or IBM marketing representative.

Appendix D. Symptom-to-part index

The symptom-to-part index lists symptoms, errors, and the possible causes. The most likely cause is listed first.

The POST BIOS displays POST error codes and messages on the screen.

Note: You must connect a keyboard, mouse, and monitor to your appliance in order to see error messages. If the engine does not recognize the monitor, keyboard, and mouse, reboot the engine while they are connected. If a Remote Supervisor Adapter is used for system management, the logs can be accessed remotely.

Use Table 17 to locate a type of symptom.

Table 17. Error symptoms index

Beep symptoms

Beep symptoms are short tones or a series of short tones separated by pauses (intervals without sound). Table 18 on page 124 shows an example of beep symptoms. See Table 19 on page 124 for a listing of the beep symptoms.

Note: One beep after successfully completing POST indicates the appliance engine is functioning correctly.

Table 18. Examples of beep symptoms

Beeps	Description
1-2-3	• One beep • A pause (or break) • Two beeps • A pause (or break) • Three Beeps
4	Four continuous beeps

Table 19. Beep symptoms

Beep/symptom	Description	Part/action
1-1-2	Microprocessor register test failed.	1. Optional microprocessor (if installed) 2. Microprocessor 3. System board
1-1-3	CMOS write/read test failed.	1. Battery 2. System board
1-1-4	BIOS EEPROM checksum failed.	1. Recover BIOS 2. System board
1-2-1	Programmable Interval Timer failed.	System board
1-2-2	DMA initialization failed.	System board
1-2-3	DMA page register write/read failed.	System board
1-2-4	RAM refresh verification failed.	1. DIMM 2. System board
1-3-1	First 64K RAM test failed	DIMM
2-1-1	Secondary DMA register failed.	System board
2-1-2	Primary DMA register failed.	System board
2-1-3	Primary interrupt mask register failed.	System board
2-1-4	Secondary interrupt mask register failed.	System board
2-2-1	Interrupt vector loading failed.	System board
2-2-2	Keyboard controller failed.	1. System board 2. Keyboard
2-2-3	CMOS power failure and checksum checks failed.	1. Battery 2. System board
2-2-4	CMOS configuration information validation failed.	1. Battery 2. System board
2-3-1	Screen initialization failed.	System board
2-3-2	Screen memory failed.	System board

Table 19. Beep symptoms (continued)

Beep/symptom	Description	Part/action
2-3-3	Screen retrace failed.	System board
2-3-4	Search for video ROM failed.	System board
2-4-1	Video failed; screen believed operable.	System board
3-1-1	Timer tick interrupt failed.	System board
3-1-2	Interval timer channel 2 failed.	System board
3-1-3	RAM test failed above address X'OFFFFH'.	1. DIMM 2. System board
3-1-4	Time-Of-Day clock failed.	1. Battery 2. System board
3-2-1	Serial port failed.	System board
3-2-2	Parallel port failed.	System board
3-2-3	Math coprocessor test failed.	1. Microprocessor 2. System board
3-2-3	Failure comparing CMOS memory size against actual.	1. DIMM 2. Battery
3-3-1	Memory size mismatch occurred.	1. DIMM 2. Battery
3-3-2	Critical SMBUS error occurred.	1. Disconnect the engine power cord from outlet, wait 30 seconds and retry. 2. System board. 3. DIMMs. 4. DASD backplane. 5. Power supply. 6. Power cage assembly. 7. I^2C cable.
3-3-3	No operational memory in system.	1. Install or reseat the memory modules, and then do a 3 boot reset. 2. DIMMs. 3. Memory board. 4. System board.
4-4-4	Optional Remote Supervisor Adapter not installed in slot 1 or not functioning correctly.	1. Verify that the adapter is installed in slot 1. 2. Adapter. 3. System board.
Two short beeps	Information only, the configuration has changed.	1. Run Base System Diagnostics. 2. Run the Configuration/Setup Utility program.
Three short beeps		1. DIMM 2. System board

Table 19. Beep symptoms (continued)

Beep/symptom	Description	Part/action
One continuous beep		1. Microprocessor 2. Optional microprocessor (if installed) 3. System board
Repeating short beeps		1. Keyboard 2. System board
One long and one short beep		System board
One long and two short beeps		System board
One long and three short beeps		1. Monitor 2. System board
Two long and two short beeps		System board

No beep symptoms

Table 20. No beep symptoms

No-beep symptom	Part/action
No beep and the system operates correctly.	1. Verify that the speaker cables are securely connected. 2. Speaker. 3. System board.
No beeps occur after successfully completing POST (the power-on status is disabled)	1. Run the Configuration/Setup Utility program, and set the Start Options Power-On Status to enable. 2. Check the speaker connection. 3. System board.
No ac power (power supply ac LED is off)	1. Check the power cord. 2. Power supply. (If two are installed, swap them to determine if one is defective.) 3. Power cage assembly. 4. Hot-swap power ac inlet box.
No beep and no video	See “Undetermined problems” on page 151.
System will not start (power supply ac LED is on)	See “Power-supply LED errors” on page 146.

Information-panel system error LED

The system error LED is turned on when an error is detected. If the system error LED is on, remove the cover and check the diagnostic panel LEDs. Table 21 on page 127 contains a complete list of diagnostic panel LEDs followed by the part or action required to correct the problem. This table is valid only when the system error LED is on.

Notes:

1. If a diagnostic panel LED is on and the information-LED panel system error LED is off, there is probably an LED problem. Run LED diagnostics.

2. To locate the LEDs on the system board see Figure 12 on page 129.
3. Check the System Event/Error Log for additional information before replacing a FRU.
4. The DIMM-error LEDs, processor-error LEDs, and VRM-error LEDs turn off when the system is shut down.

Table 21. Errors diagnosed by the diagnostic panel LEDs

Diagnostics panel LED	Part/action
All LEDs off (Check System Error Log for error condition, then clear System Error Log when the problem is found.)	1. System Error Log is 75% full; clear the log. 2. PFA alert; check log for failure; clear PFA alert; remove ac power for at least 20 seconds, reconnect, and then turn on the system. 3. Run light path diagnostics.
MEMORY LED on (The LED next to the failing DIMM is on.)	1. Failing DIMM 2. System board
CPU LED on (The LED next to the failing CPU is on.)	1. Microprocessor 1 or 2 2. System board
PCI BUS LED on	1. Remove all PCI adapters from slots on affected bus (refer to “Adapters” in Chapter 3 of the hardware installation guide). 2. System board.
VRM LED on (The LED next to the failing VRM is on.)	1. Voltage regulator module indicated by the lit VRM LED. 2. Microprocessor indicated by the microprocessor LED.
DASD LED on (The LED located next to the drive bay that the failing drive is installed in is lit. Check the amber drive LED for the failing hard drive.)	1. Be sure that the fans are operating correctly and the airflow is good. 2. If installed, reseat I^2C cable between DASD backplane and DASD I^2C on the system board (J10). 3. Failing drive. SCSI channel A has failed. (This is the SCSI channel for the hot-swap hard disk drives). 4. SCSI backplane.
SERVICE PROCESSOR BUS LED	1. Unplug the power supplies from the engine for 30 seconds, and then plug in and retry. 2. Reflash or update firmware for ISMP, BIOS. 3. System board.
POWER SUPPLY 1 LED on	1. Check the dc good LED on power supply 1. If it is off, replace power supply 1. 2. Power cage assembly.
POWER SUPPLY 2 LED on	1. Check the dc good LED on power supply 2. If it is off, replace power supply 2. 2. Power cage assembly.
NONREDUNDANT LED on	1. Check the PS1 and PS2 LEDs and replace any indicated power supply. 2. Install an additional power supply or remove optional devices from the engine.
NMI LED on	1. Restart the engine. 2. Check the System Error Log.

Table 21. Errors diagnosed by the diagnostic panel LEDs (continued)

Diagnostics panel LED	Part/action
TEMPERATURE LED on	1. Ambient temperature must be within normal operating specifications. Refer to “Specifications” in Chapter 1 of the hardware installation guide. 2. Ensure filler panels are installed in all empty bays. 3. Ensure fans are operating correctly by checking the fan LEDs. 4. Examine System Error Log. a. System over recommended temperature • Information LED panel b. DASD over recommended temperature (DASD LED also on) 1) Overheating hard drive 2) DASD backplane c. System over recommended temperature for CPU *x* (where *x* is 1 or 2) (CPU LED is also on) 1) CPU *x* 2) System board 5. If the CPU LED on the diagnostics panel is also lit, one of the microprocessors has caused the error.
FAN LED on	1. Check individual fan LEDs. 2. Replace respective fan. 3. Fan cable. 4. System board. 5. Power cage assembly.

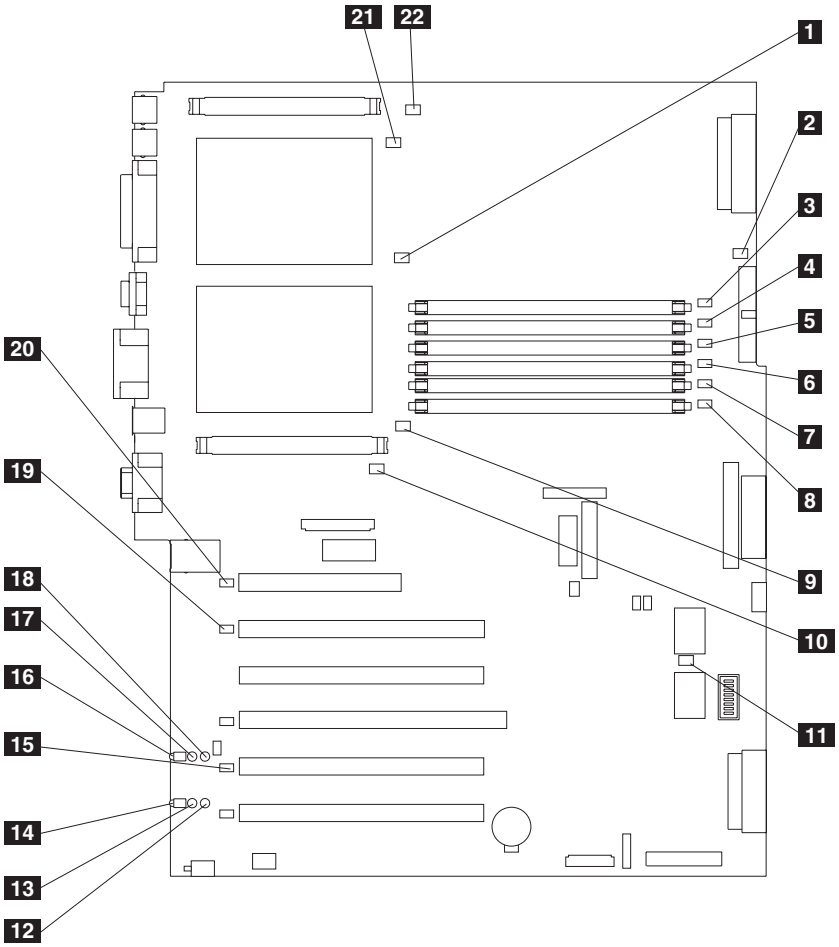

Figure 12. System-board LED locations

1 CPU mismatch error LED (CR14)
2 Power error LED (CR15)
3 DIMM 1 error LED (CR16)
4 DIMM 2 error LED (CR17)
5 DIMM 3 error LED (CR18)
6 DIMM 4 error LED (CR20)
7 DIMM 5 error LED (CR22)
8 DIMM 6 error LED (CR23)
9 CPU 1 error LED (CR24)
10 VRM 1 error LED (CR33)
11 Service processor activity LED (CR67)
12 PCI-X slot 6 power LED (CR79)
13 PCI-X slot 6 internal attention LED (CR78) (Disabled)
14 PCI-X slot 6 external attention LED (CR77) (Disabled)
15 PCI-X bus C error LED (CR76)
16 PCI-X slot 5 external attention LED (CR74) (Disabled)
17 PCI-X slot 5 internal attention LED (CR73) (Disabled)
18 PCI-X slot 5 power LED (CR75)
19 PCI-X bus B error LED (CR68)
20 PCI bus A error LED (CR66)
21 CPU 2 error LED (CR4)
22 VRM 2 error LED (CR1)

Diagnostics-panel LEDs viewed with the cover off:

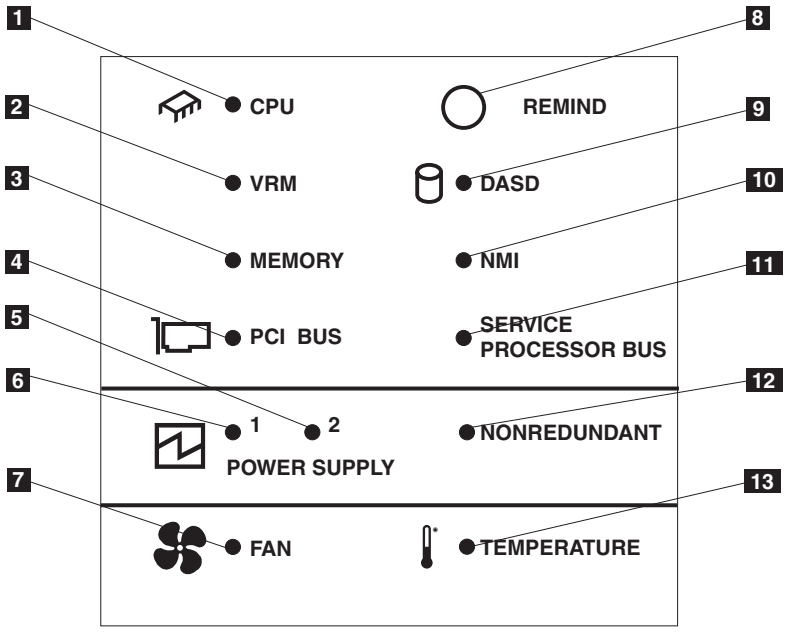

Figure 13. Diagnostics panel LEDs (viewed with the cover off)

Table 22. Diagnostics-panel LED descriptions

Index	Name	Meaning
1	CPU	Microprocessor failure. One or both microprocessors have failed.
2	VRM	Error on VRM or on integrated voltage regulator. The LED next to the affected VRM will also be on.
3	Memory	Memory failure. One or more memory DIMMs have failed.
4	PCI BUS	Error on the PCI bus or system board.
5	Power supply 2	Power supply 2 failure.
6	Power supply 1	Power supply 1 failure.
7	Fan	A fan failed or is operating slowly.
8	Remind button	Press this button to temporarily reset the LEDs on the diagnostics panel.
9	DASD	A hot-swap hard disk drive, backplane, or other part of SCSI channel A has failed. The amber LED next to the drive bay with the failing drive will also be on.
10	NMI	A non-maskable interrupt occurred.
11	Service Processor bus	The system environmental monitor detected an error.
12	Nonredundant	Nonredundant power.
13	Temperature	The operating temperature inside the engine was exceeded.

Diagnostic error codes

Note: In the following error codes, if *XXX* is *000*, *195*, or *197*, ***do not*** replace a part. The description for these error codes are:

000 The test passed.

195 The Esc key was pressed to abort the test.

197 This is a warning error and might not indicate a hardware failure.

For all error codes, replace/follow the Part/action indicated.

Attention: If diagnostic error messages appear that are not listed in the tables, make sure that your appliance engine has the latest levels of BIOS, Advanced System.

Table 23 describes diagnostic error codes and their suggested actions.

Table 23. Diagnostic error codes

Error code/symptom	Meaning	Part/action
001-XXX-000	Failed core tests.	System board
001-XXX-001	Failed core tests.	System board
001-250-000	Failed system board ECC.	System board
001-250-001	Failed system board ECC.	System board
005-XXX-000	Failed video test.	System board
011-XXX-000	Failed COM1 serial port test.	System board
011-XXX-001	Failed COM2 serial port test.	System board
014-XXX-000	Failed parallel port test.	System board
015-XXX-001	USB interface not found, board damaged.	System board
015-XXX-015	Failed USB external loopback test.	No action required. USB is not required for normal functioning.
015-XXX-198	USB device connected during USB test.	No action required. USB is not required for normal functioning.
020-XXX-000	Failed PCI interface test.	System board
020-XXX-001	Failed hot-swap slot 1 PCI latch test.	1. PCI hot-swap latch assembly 2. System board
020-XXX-002	Failed Hot-swap slot 2 PCI latch test.	1. PCI hot-swap latch assembly 2. System board
020-XXX-003	Failed hot-swap slot 3 PCI latch test.	1. PCI hot-swap latch assembly 2. System board
020-XXX-004	Failed hot-swap slot 4 PCI latch test.	1. PCI hot-swap latch assembly 2. System board
030-XXX-000	Failed internal SCSI interface test.	System board
035-XXX-099		1. No adapters were found. 2. If adapter is installed recheck connection.

Table 23. Diagnostic error codes (continued)

Error code/symptom	Meaning	Part/action
035-XXX-S99	Failed RAID test on PCI slot *S*, where *S* = number of failing PCI slot. Check System Error Log before replacing a part.	1. Adapter 2. SCSI backplane 3. Cable
035-XXX-SNN	Check System Error Log before replacing a part. *s* = number of failing PCI slot, *nn* = SCSI ID of failing fixed disk.	Hard disk drive with SCSI ID *nn*on RAID adapter in PCI slot *s*.
035-253-S99	RAID adapter initialization failure.	1. ServeRAID adapter in slot *s* is not configured correctly. Obtain the basic and extended configuration status and see the *ServeRAID User's Reference* on the Documentation CD-ROM for more information. 2. Cable. 3. SCSI backplane. 4. Adapter.
075-XXX-000	Failed power supply test.	Power supply
089-XXX-001	Failed microprocessor test.	1. VRM 1 for microprocessor 1 2. Microprocessor 1
089-XXX-002	Failed optional microprocessor test.	1. VRM 2 for optional microprocessor 2 2. Optional microprocessor 2
166-198-000 System Management: Aborted	Unable to communicate with ASM. It might be busy. Run the test again.	1. Run the diagnostic test again. 2. Correct other error conditions and retry. These include other failed system management tests and items logged in the System Error Log of the optional Remote Supervisor Adapter. 3. Disconnect all engine and option power cords from the engine, wait 30 seconds, reconnect, and retry. 4. Remote Supervisor Adapter, if installed. 5. System board.
166-201-001 System Management: Failed	I^2C bus error(s). See SERVPROC and DIAGS entries in event log.	1. If installed, reseat the I^2C cable between the Remote Supervisor Adapter (in PCI slot 1/J32) and the system board (J27). 2. Reseat memory DIMMs. 3. Memory DIMMs. 4. System board.
166-201-002 System Management: Failed	I^2C bus error(s) See SERVPROC and DIAGS entries in event log.	1. Reseat I^2C cable between the operator information panel and the system board (J24). 2. Reseat I^2C cable between the diagnostics panel and the system board (J23). 3. Operator information panel. 4. Diagnostics panel. 5. System board.

Table 23. Diagnostic error codes (continued)

Error code/symptom	Meaning	Part/action
166-201-003 System Management: Failed	I^2C bus error(s) See SERVPROC and DIAGS entries in event log.	1. Reseat cables between the system board and the power supply or power cage assembly. 2. Power cage assembly. 3. System board.
166-201-004 System Management: Failed	I^2C bus error(s) See SERVPROC and DIAGS entries in event log.	1. DASD backplane 2. System board
166-201-005 System Management: Failed	I^2C bus error(s) See SERVPROC and DIAGS entries in event log.	1. Reseat Memory DIMMs. 2. Reseat microprocessors. 3. Memory DIMMs. 4. Microprocessors. 5. System board.
166-250-000 System Management: Failed	I^2C cable is disconnected. Reconnect I^2C cable between Remote Supervisor Adapter and system board.	1. Reseat I^2C cable between the Remote Supervisor Adapter (in PCI slot 1/J32) and the system board (J27). 2. I^2C cables. 3. Replace the Remote Supervisor Adapter. 4. System board.
166-260-000 System Management: Failed	Restart Remote Supervisor Adapter Error. After restarting, Remote Supervisor Adapter communication was lost. Unplug and cold boot to reset the Remote Supervisor Adapter.	1. Disconnect all engine and option power cords from the engine, wait 30 seconds, reconnect, and retry. 2. Reseat the Remote Supervisor Adapter (in PCI slots 1/J32). 3. Replace the Remote Supervisor Adapter.
166-342-000 System Management: Failed	ASM adapter BIST indicate failed tests.	1. Ensure that the latest firmware levels are installed for Remote Supervisor Adapter and BIOS. 2. Disconnect all engine and option power cords from engine, wait 30 seconds, reconnect, and retry. 3. Remote Supervisor Adapter.
166-400-000 System Management: Failed	ISMP self test result failed tests: *x* where *x* = Flash, RAM, or ROM.	1. Reflash or update firmware for ISMP. 2. System board.
180-XXX-000	Diagnostics panel LED failure.	Run diagnostics panel LED test for the failing LED.
180-XXX-001	Failed front LED panel test.	1. Operator information panel 2. System board
180-XXX-002	Failed diagnostics LED panel test.	1. Diagnostics panel 2. System board
180-361-003	Failed fan LED test.	1. Fan(s) 2. System board
180-XXX-003	Failed system board LED test.	System board
180-XXX-005	Failed SCSI backplane LED test.	1. SCSI backplane 2. SCSI backplane cable 3. System board

Table 23. Diagnostic error codes (continued)

Error code/symptom	Meaning	Part/action
201-XXX-0NN	Failed memory test.	1. DIMM Location slots 1-6 where ***nn* = DIMM location.** **Note:** *nn* : 1=DIMM 1 2=DIMM 2 3=DIMM 3 4=DIMM 4 5=DIMM 5 6=DIMM 6. 2. System board
201-XXX-999	Multiple DIMM failure, see error text.	1. See error text for failing DIMMs. 2. System board.
202-XXX-001	Failed system cache test.	1. VRM 1 2. Microprocessor 1
202-XXX-002	Failed system cache test.	1. VRM 2 2. Microprocessor 2
206-XXX-000	Failed diskette drive test.	1. Cable 2. Diskette drive 3. System board
215-XXX-000	Failed IDE CD-ROM drive test.	1. CD-ROM drive cables 2. CD-ROM drive 3. System board
217-198-XXX	Could not establish drive parameters.	1. Check cable and termination. 2. SCSI backplane. 3. Hard disk.
217-XXX-000	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array.	Logical drive 1
217-XXX-001	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array.	Logical drive 2
217-XXX-002	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array.	Logical drive 3
217-XXX-003	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array.	Logical drive 4
217-XXX-004	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array.	Logical drive 5

Table 23. Diagnostic error codes (continued)

Error code/symptom	Meaning	Part/action
217-XXX-005	Failed BIOS hard disk test. **Note:** If RAID is configured, the hard disk number refers to the RAID logical array.	Logical drive 6
264-XXX-0NN	Failed tape drive test.	1. Tape cartridge, if user executed the Read/Write Tape Drive test (failure code of XXX = 256) 2. SCSI or power cable connected to tape drive with SCSI ID *nn* 3. Tape drive with SCSI ID *nn* (refer to the Help and Service Information appendix of the tape drive's User Guide) 4. System board or SCSI controller (run SCSI controller diagnostic to determine if the SCSI bus is functioning properly.)
264-XXX-999	Errors on multiple tape drives, see error text for more information.	See error messages/text in the Base System Diagnostics error log for detailed information on each individual tape drive error.
301-XXX-000	Failed keyboard test.	Keyboard
405-XXX-000	Failed Ethernet test on controller on the system board.	1. Verify that Ethernet is not disabled in BIOS. 2. System board.
405-XXX-00N	Failed Ethernet test on adapter in PCI slot *n*.	1. Adapter in PCI slot ***n*** 2. System board
415-XXX-000	Failed Modem test.	No action required. A modem is not required for normal functioning.

Error symptoms

You can use the error symptom table to find solutions to problems that have definite symptoms.

If you cannot find the problem in the error symptom charts, go to "Starting the diagnostic programs" on page 94 to test the engine.

If you have just added new software or a new option and your engine is not working, do the following before using the error symptom charts:

- Remove the software or device that you just added.
- Run the diagnostic tests to determine if your engine is running correctly.
- Reinstall the new software or new device.

Table 24 on page 136 contains information on error symptoms and suggested actions. In the table, if the entry in the Part/action column is a suggested action, perform that action; if it is the name of a component, reseat the component and replace it if necessary. The most likely cause of the symptom is listed first.

Table 24. Error symptoms and suggested actions

CD-ROM drive problems	
Symptom	**Part/action**
CD-ROM drive is not recognized.	1. Verify that: • The primary IDE channel to which the CD-ROM drive is attached is enabled in the Configuration/Setup Utility program. **Note:** On an engine with a single IDE channel, only the primary channel can be used. • All cables and jumpers are installed correctly. • The correct device driver is installed for the CD-ROM drive. 2. Run CD-ROM drive diagnostics. 3. CD-ROM drive.
Diskette drive problems	
Symptom	**Part/action**
Diskette drive activity LED stays on, or the system bypasses the diskette drive.	1. If there is a diskette in the drive, verify that: • The diskette drive is enabled in the Configuration/Setup utility program. • The diskette is good and not damaged. (Try another diskette if you have one.) • The diskette is inserted correctly in the drive. • The diskette contains the necessary files to start the engine. • The software program is working properly. • The cable is installed correctly (in the proper orientation). 2. Run diskette drive diagnostics. 3. Cable. 4. Diskette drive. 5. System board.
Hard disk drive problems	
Symptom	**Part/action**
Not all drives are recognized by the hard disk drive diagnostic test (Fixed Disk test).	1. Remove the first drive not recognized and try the hard disk drive diagnostic test again. 2. If the remaining drives are recognized, replace the drive you removed with a new one.
System stops responding during hard disk drive diagnostic test.	1. Remove the hard disk drive being tested when the engine stopped responding and try the diagnostic test again. 2. If the hard disk drive diagnostic test runs successfully, replace the drive you removed with a new one.
General problems	
Symptom	**Part/action**
Problems such as broken cover locks or indicator LEDs not working	Broken part
Intermittent problems	
Symptom	**Part/action**

Table 24. Error symptoms and suggested actions (continued)

A problem occurs only occasionally and is difficult to detect.	1. Verify that: • All cables and cords are connected securely to the rear of the engine and attached options. • When the engine is turned on, air is flowing from the rear of the engine at the fan grill. If there is no airflow, the fan is not working. This causes the engine to overheat and shut down. • Ensure that the SCSI bus and devices are configured correctly and that the last external device in each SCSI chain is terminated correctly. 2. Check the system error log from Setup or through the diagnostics in the optional Remote Supervisor Adapter.
Keyboard, mouse, or point-device problems	
Symptom	**Part/action**
All or some keys on the keyboard do not work.	1. Verify that: • The keyboard cable is securely connected to the system, and the keyboard and mouse cables are not reversed. • The engine and the monitor are turned on. 2. Keyboard. 3. System board.
The mouse or pointing device does not work.	1. Verify that the mouse or pointing-device cable is securely connected, and that the keyboard and mouse cables are not reversed. 2. Mouse or pointing device. 3. System board.
Memory problems	
Symptom	**Part/action**
The amount of system memory displayed is less than the amount of physical memory installed.	1. Verify that: • The memory modules are seated properly. • You have installed the correct type of memory. • If you changed the memory, you updated the memory configuration with the Configuration/Setup Utility program. • All banks of memory on the DIMMs are enabled. The engine might have automatically disabled a DIMM bank when it detected a problem or a DIMM bank could have been manually disabled. 2. Check POST error log for error message 289: • If the DIMM was disabled by a system-management interrupt (SMI), replace the DIMM. • If the DIMM was disabled by the user or by POST: a. Start the Configuration/Setup Utility program. b. Enable the DIMM. c. Save the configuration and restart the engine. 3. DIMM. 4. System board.
Microprocessor problems	
Symptom	**Part/action**
The engine emits a continuous tone during POST. (The startup (boot) microprocessor is not working properly.)	1. Verify that the startup microprocessor is seated properly. 2. Startup microprocessor.
Monitor problems	
Symptom	**Part/action**

Table 24. Error symptoms and suggested actions (continued)

Testing the monitor.	See the information that comes with the monitor for adjusting and testing instructions. (Some IBM monitors have their own self-tests.)
The screen is blank.	1. Verify that: • The engine power cord is plugged into the engine and a working electrical outlet. • The monitor cables are connected properly. • The monitor is turned on and the Brightness and Contrast controls are adjusted correctly. • If the engines are C2T chained together, verify that: – The C2T chain cables are securely connected to the engines. – The C2T breakout cable is connected correctly. – A engine that is turned on is selected. **Important:** In some memory configurations, the 3-3-3 beep code might sound during POST followed by a blank display screen. If this occurs and the **Boot Fail Count** feature in the Start Options of the Configuration/Setup Utility program is set to **Enabled** (its default setting), you must restart the engine three times to force the system BIOS to reset the CMOS values to the default configuration (memory connector or bank of connectors enabled). 2. If you have verified these items and the screen remains blank, replace: a. Monitor b. System board
Only the cursor appears.	See “Undetermined problems” on page 151.
The monitor works when you turn on the engine but goes blank when you start some application programs.	1. Verify that: • The application program is not setting a display mode higher than the capability of the monitor. • You installed the necessary device drivers for the applications. 2. If you have verified these items and the screen remains blank, replace the monitor.
The screen is wavy, unreadable, rolling, distorted, or has screen jitter.	1. If the monitor self-tests show the monitor is working properly, consider the location of the monitor. Magnetic fields around other devices (such as transformers, appliances, fluorescent lights, and other monitors) can cause screen jitter or wavy, unreadable, rolling, or distorted screen images. If this happens, turn off the monitor. (Moving a color monitor while it is turned on might cause screen discoloration.) Then move the device and the monitor at least 305 mm (12 in.) apart. Turn on the monitor. **Notes:** a. To prevent diskette drive read/write errors, be sure the distance between monitors and diskette drives is at least 76 mm (3 in.). b. Non-IBM monitor cables might cause unpredictable problems. c. An enhanced monitor cable with additional shielding is available for the 9521 and 9527 monitors. For information about the enhanced monitor cable, contact your IBM reseller or IBM marketing representative. 2. System board.
Wrong characters appear on the screen.	1. If the wrong language is displayed, update the BIOS code with the correct language. 2. System board.
Option problems	
Symptom	**Part/action**

Table 24. Error symptoms and suggested actions (continued)

An IBM option that was just installed does not work.	1. Verify that: • You followed the installation instructions that came with the option. • The option is installed correctly. • You have not loosened any other installed options or cables. • You updated the configuration information in the Configuration/Setup Utility program. Whenever memory or an option is changed, you must update the configuration. 2. Option you just installed.
An IBM option that used to work does not work now.	1. Verify that all of the option hardware and cable connections are secure. 2. If the option comes with its own test instructions, use those instructions to test the option. 3. If the failing option is a SCSI option, verify that: • The cables for all external SCSI options are connected correctly. • The last option in each SCSI chain, or the end of the SCSI cable, is terminated correctly. • Any external SCSI option is turned on. You must turn on an external SCSI option before turning on the engine. 4. Failing option.
Power problems	
Symptom	**Part/action**
The engine does not turn on.	1. Verify that: • The power cables are properly connected to the engine. • The electrical outlet functions properly. • The type of memory installed is correct. • If you just installed an option, remove it, and restart the engine. If the engine now turns on, you might have installed more options than the power supply supports. 2. If LEDs for CPUs or VRMs are on, verify that: a. A VRM is populated for each microprocessor. b. Override front panel pushbutton by turning on switch 7 of SW1; if power comes on: 1) Service processor error. 2) Power reset card. 3. See “Undetermined problems” on page 151.
The engine does not turn off.	1. Verify whether you are using an ACPI or non-ACPI operating system. If you are using a non-ACPI operating system: a. Press Ctrl+Alt+Delete. b. Turn off the system by holding the power-control button for 4 seconds. c. If engine fails during BIOS POST and power-control button does not work, remove the AC power cord. 2. If the problem remains or if you are using an ACPI-aware operating system, suspect the system board.
Serial port problems	
Symptom	**Part/action**
The number of serial ports identified by the operating system is less than the number of serial ports installed.	1. Verify that: • Each port is assigned a unique address by the Configuration/Setup Utility program and none of the serial ports is disabled. • The serial-port adapter, if you installed one, is seated properly. 2. Failing serial port adapter.

Table 24. Error symptoms and suggested actions (continued)

A serial device does not work.	No corrective action. A serial device is not required for normal functioning.
Software problem	
Symptom	**Part/action**
Suspected software problem.	1. To determine if problems are caused by the software, verify that: • Your engine has the minimum memory needed to use the software. For memory requirements, see the information that comes with the software. **Note:** If you have just installed an adapter or memory, you might have a memory address conflict. • The software is designed to operate on your engine. • Other software works on your engine. • The software that you are using works on another system. If you received any error messages when using the software program, see the information that comes with the software for a description of the messages and suggested solutions to the problem. 2. If you have verified these items and the problem remains, contact your place of purchase.
Universal Serial Bus (USB) port problems	
Symptom	**Part/action**
A USB device does not work.	No corrective action. A USB device is not required for normal functioning.

POST error codes

Note: You must connect a keyboard, mouse, and monitor to your appliance in order to see error messages. If the engine does not recognize the monitor, keyboard, and mouse, reboot the engine while they are connected. If a Remote Supervisor Adapter is used for system management, the logs can be accessed remotely.

Table 25 describes POST error codes and their suggested actions. *X* can be any number or letter.

In the following error codes, *X* can be any number or letter.

Table 25. POST error codes

Error code/symptom	Meaning	Part/action
062	Three consecutive startup failures using the default configuration.	1. Run the Configuration/Setup Utility program. 2. Battery. 3. System board. 4. Microprocessor.
101, 102	System and processor error.	System board
106	System and processor error.	System board
111	Channel check error.	1. Memory DIMM 2. System board
114	Adapter read-only memory error.	1. Failing adapter. 2. Run Base System Diagnostics.

Table 25. POST error codes (continued)

Error code/symptom	Meaning	Part/action
129	Internal cache error.	1. Microprocessor 2. Optional microprocessor (if installed)
151	Real time clock error.	1. Run Base System Diagnostics. 2. Battery. 3. System board.
161	Real time clock battery error.	1. Run the Configuration/Setup Utility program. 2. Battery. 3. System board.
162	Device configuration error. **Note:** Be sure to load the default settings and any additional desired settings; then, *save the configuration*.	1. Run the Configuration/Setup Utility program. 2. Battery. 3. Failing device. 4. System board.
163	Real-time clock error.	1. Run the Configuration/Setup Utility program. 2. Battery. 3. System board.
164	Memory configuration changed.	1. Run the Configuration/Setup Utility program. 2. DIMM. 3. System board.
175	Hardware error.	System board
176	Computer cover or cable cover was removed without a key being used.	1. Run the Configuration/Setup Utility program. 2. System board.
177, 178	Security hardware error.	1. Run the Configuration/Setup Utility program. 2. System board.
184	Power-on password damaged.	1. Run the Configuration/Setup Utility program. 2. System board.
185	Drive startup sequence information corrupted.	1. Run the Configuration/Setup Utility program. 2. System board.
186	Security hardware control logic failed.	1. Run the Configuration/Setup Utility program. 2. System board.
187	VPD serial number not set.	1. Set serial number in the Configuration/Setup Utility program. 2. System board.
188	Bad EEPROM CRC #2.	1. Run the Configuration/Setup Utility program. 2. System board.
189	An attempt was made to access the engine with invalid passwords.	1. Run the Configuration/Setup Utility program, and type the administrator password.

Table 25. POST error codes (continued)

Error code/symptom	Meaning	Part/action
201	Memory test error. If the engine does not have the latest level of BIOS installed, update the BIOS to the latest level and run the diagnostic program again.	1. DIMM 2. System board
229	Cache error.	1. Microprocessor 2. Optional microprocessor (if installed)
262	DRAM parity configuration error.	1. Run the Configuration/Setup Utility program. 2. Battery. 3. System board.
289	DIMM disabled by POST or user.	1. Run the Configuration/Setup Utility program, if the DIMM was disabled by the user. 2. Disabled DIMM, if not disabled by user.
301	Keyboard or keyboard controller error.	1. Keyboard 2. System board
303	Keyboard controller error.	System board
602	Invalid diskette boot record.	1. Diskette 2. Diskette drive 3. Cable 4. System board
604	Diskette drive error.	1. Run the Configuration/Setup Utility program and Base System Diagnostics. 2. Diskette drive. 3. Drive cable. 4. System board.
605	Unlock failure.	1. Diskette drive 2. Drive cable 3. System board
662	Diskette drive configuration error.	1. Run the Configuration/Setup Utility program and Base System Diagnostics. 2. Diskette drive. 3. Drive cable. 4. System board.
762	Coprocessor configuration error.	1. Run the Configuration/Setup Utility program. 2. Battery. 3. Microprocessor.
962	Parallel port error.	No corrective action. No parallel device is required for normal functioning.
11XX	System board serial port 1 or 2 error.	1. Disconnect the external cable on the serial port. 2. Run the Configuration/Setup Utility program. 3. System board.

Table 25. POST error codes (continued)

Error code/symptom	Meaning	Part/action
1301	I^2C cable to front panel not found.	1. Cable 2. Front panel 3. Power switch assembly 4. System board
1302	I^2C cable from system board to power on and reset switches not found.	1. Cable 2. Power switch assembly 3. System board
1303	I^2C cable from system board to power backplane not found.	1. Cable 2. Power cage assembly, if installed 3. System board
1304	I^2C cable to diagnostic LED board not found.	1. Power switch assembly 2. System board
1600	The system management processor is not functioning. Do the following before replacing a part: 1. Ensure that a jumper is not installed on J34. 2. Remove the ac power to the engine, wait 20 seconds; then, reconnect the ac power. Wait 30 seconds; then, turn on the engine.	System board
1601	The system is able to communicate to the system management processor, but the system management processor failed to respond at the start of POST. Do the following before replacing a part: 1. Remove the ac power to the engine, wait 20 seconds; then, reconnect the ac power. Wait 30 seconds; then, turn on the engine. 2. Flash update the Remote Supervisor Adapter. See the *Remote Supervisor Adapter User's Guide* provided on the Documentation CD-ROM for more information.	1. Remote Supervisor Adapter, if installed 2. System board
1602	Cable for optional service processor adapter not installed.	Disconnect all engine and option power cords from engine, wait 30 seconds, reconnect, and retry.

Table 25. POST error codes (continued)

Error code/symptom	Meaning	Part/action
1762	Hard disk configuration error.	1. Hard disk drive. 2. Hard disk cables. 3. Run the Configuration/Setup Utility program. 4. Hard disk adapter. 5. SCSI backplane. 6. System board.
178X	Fixed disk error.	1. Hard disk cables. 2. Run Base System Diagnostics. 3. Hard disk adapter. 4. Hard disk drive. 5. System board.
1800	No more hardware interrupt available for PCI adapter.	1. Run the Configuration/Setup Utility program. 2. Failing adapter. 3. System board.
1962	Drive does not contain a valid boot sector.	1. Verify that a startable operating system is installed. 2. Run Base System Diagnostics. 3. Hard disk drive. 4. SCSI backplane. 5. Cable. 6. System board.
2400	Video controller test failure.	System board
2462	Video memory configuration error.	System board
5962	IDE CD-ROM drive configuration error.	1. Run the Configuration/Setup Utility program and load the default values. 2. CD-ROM drive. 3. CD-ROM power cable. 4. IDE cable. 5. System board. 6. Battery.
8603	Pointing-device error.	1. Pointing device 2. System board
0001200	Machine check architecture error.	1. Microprocessor 1 2. Optional microprocessor 2
00012000	Microprocessor machine check.	1. Microprocessor 2. System board
00019501	Microprocessor 1 is not functioning - check VRM and microprocessor LEDs.	1. VRM 1 2. Microprocessor 1 3. System board
00019502	Microprocessor 2 is not functioning - check VRM and microprocessor LEDs.	1. VRM 2 2. Microprocessor 2

Table 25. POST error codes (continued)

Error code/symptom	Meaning	Part/action
00019701	Microprocessor 1 failed.	1. Microprocessor 1 2. System board
00019702	Microprocessor 2 failed.	1. Microprocessor 2 2. System board
00180100	A PCI adapter has requested memory resources that are not available.	1. Ensure that the PCI adapter and all other adapters are set correctly in the Configuration/Setup Utility program Utility program. If the memory resource settings are not correct, change the settings. 2. If all memory resources are being used, you might need to remove an adapter to make memory available to the PCI adapter. Disabling the adapter BIOS on the adapter might correct the error. (See the documentation provided with the adapter.)
00180200	No more I/O space available for PCI adapter.	1. Run the Configuration/Setup Utility program. 2. Failing adapter. 3. System board.
00180300	No more memory (above 1MB for PCI adapter).	1. Run the Configuration/Setup Utility program. 2. Failing adapter. 3. System board.
00180400	No more memory (below 1MB for PCI adapter).	1. Run the Configuration/Setup Utility program. 2. Failing adapter. 3. System board.
00180500	PCI option ROM checksum error.	1. Remove failing PCI card. 2. System board.
00180600	PCI to PCI bridge error.	1. Run the Configuration/Setup Utility program. 2. Failing adapter. 3. System board.
00180700, 00180800	General PCI error.	1. System board 2. PCI card
00181000	PCI error.	1. Adapter 2. System board
01295085	ECC checking hardware test error.	1. System board 2. Microprocessor
01298001	System BIOS installed on this engine does not support level of processor.	Microprocessor 1.
01298002	System BIOS installed on this engine does not support level of processor.	Microprocessor 2.
01298101	System BIOS installed on this engine does not support level of processor.	Microprocessor 1.

Table 25. POST error codes (continued)

Error code/symptom	Meaning	Part/action
01298102	System BIOS installed on this engine does not support level of processor.	Microprocessor 2.
I9990301	Hard disk sector error.	1. Hard disk drive 2. SCSI backplane 3. Cable 4. System board
I9990305	Hard disk sector error, no operating system installed.	Install operating system to hard disk.
I9990650	AC power has been restored.	1. Check cable. 2. Check for interruption of power. 3. Power cable.

Fan error messages

Table 26. Fan error messages

Message	Action
Fan "*X*" failure (level-critical; fan "*X*" had a failure)	1. Check connections to fan "*X*" 2. Replace fan "*X*"
Fan "*X*" fault (level-critical; fan "*X*" beyond recommended RPM range)	1. Check connections to fan "*X*" 2. Replace fan "*X*"
Fan "*X*" outside recommended speed	Replace fan "*X*"

Power-supply LED errors

Use the information in Table 27 to troubleshoot power-supply problems.

Note: The minimum configuration required to enable the DC power LED is:

- Power supply
- Power backplane
- System board (with pins 2 and 3 on the J23 extension cable connected to bypass the power switch; see Figure 14 on page 147).

Table 27. Power-supply LED errors

AC good LED	DC good LED	Description	Part/action
Off	Off	No power to system or ac problem.	1. Check ac power to the system. 2. Power supply.

Table 27. Power-supply LED errors (continued)

AC good LED	DC good LED	Description	Part/action
On	Off	Standby mode or dc problem.	1. Check system board cable connectors J4 and J10. Move switch 7 of SW 1 to bypass power control. If the dc good LED is lit, press Ctrl+Alt+Delete. Watch the screen for any POST errors. Check the System Error Log for any listed problems. If the system starts with no errors: a. Power switch assembly b. System board 2. Remove the adapters and disconnect the cables and power connectors to all internal and external devices. Turn on the system. If the dc good LED is lit, replace the adapters and devices one at a time until you isolate the problem. 3. Power supply. 4. Power cage assembly, if installed. 5. System board.
On	On	Power is working properly.	N/A

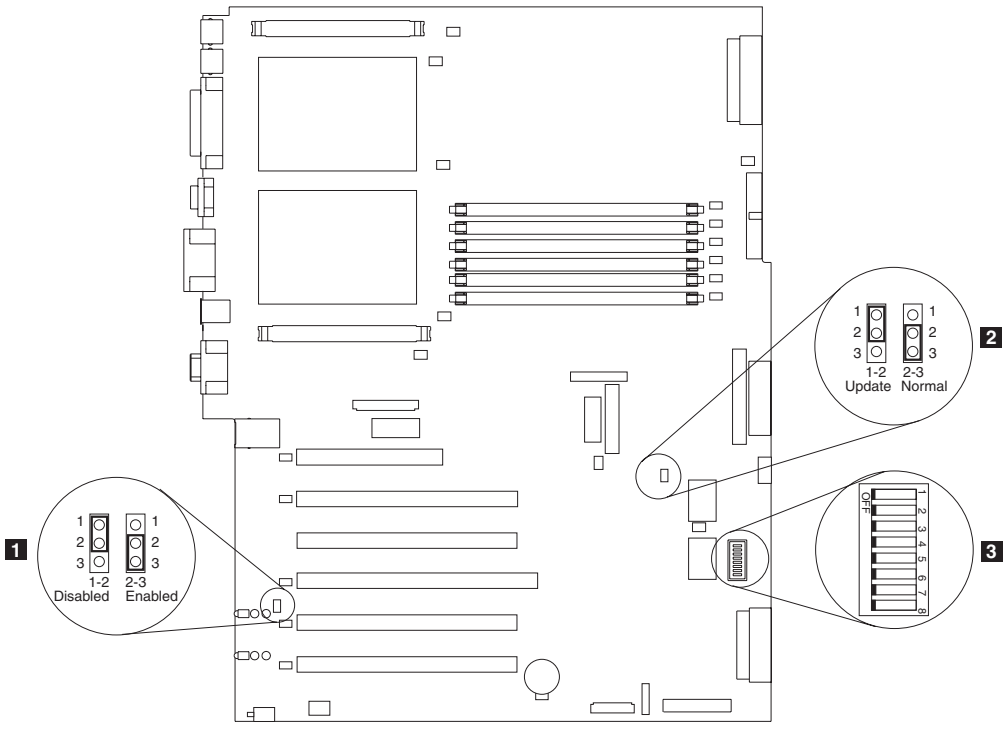

Figure 14. System-board switches and jumpers

1 Adapter jumper (J42)
2 Boot block recovery jumper (J28)
3 System board switch block (SW1)

Power error messages

Table 28. Power error messages

Message	Action
Power supply *x* current share fault (level-critical; excessive current demand on power supply *x*)	Replace power supply ***x***.
Power supply *x* DC good fault (level-critical; power good signal not detected for power supply *x*)	Replace power supply ***x***.
Power supply *x* temperature fault	Replace power supply ***x***.
Power supply *x* removed	No action required - information only.
Power supply *x* fan fault (level-critical; fan fault in power supply *x*)	Replace power supply ***x***.
Power supply *x* 12 V fault (level-critical; overcurrent condition detected)	See “Power checkout” on page 105.
Power supply *x* 3.3 V fault (level-critical; 3.3 V power supply *x* had an error)	See “Power checkout” on page 105.
Power supply *x* 5 V fault (level-critical; 5 V power supply *x* had an error)	See “Power checkout” on page 105.
System running non-redundant power (level-noncritical; system does not have redundant power)	Replace power supply. **Note:** System can continue to operate without redundancy protection.
System under recommended voltage for *x* (level-warning; indicated voltage supply under nominal value; value for *x* can be +12, -12, or +5)	1. Check connections to the power subsystem. 2. Power supply. 3. Power cage assembly.

SCSI error codes

Table 29 describes SCSI error codes and their suggested actions.

Table 29. SCSI error codes and actions

Error code	Meaning	Part/action
All SCSI Errors	You might be experiencing one or more of the following problems: • A failing SCSI device (adapter, drive, or controller) • An incorrect SCSI configuration or SCSI termination jumper setting • Duplicate SCSI IDs in the same SCSI chain • A missing or incorrectly installed SCSI terminator • A defective SCSI terminator • An incorrectly installed cable • A defective cable	1. Ensure that the external SCSI devices were powered-on before the appliance was powered-on. 2. Ensure that the cables for all external SCSI devices are connected correctly. 3. If you have attached an external SCSI device to the appliance, make sure the external SCSI termination is set to automatic. 4. Ensure that the last device in each SCSI chain is terminated correctly. 5. Ensure that the SCSI devices are configured correctly.

Bus fault messages

Table 30. Bus fault messages

Message	Action
Failure reading I²C device (Check devices on bus 0)	1. If installed, reseat the I²C cable between Remote Supervisor Adapter (in PCI slot 1/J32) and system board (J45). 2. Memory DIMMs. 3. System board.
Failure reading I²C device (Check devices on bus 1)	1. Reseat the I²C cable between the operator information panel and system board (J24). 2. Operator information panel. 3. System board.
Failure reading I²C device (Check devices on bus 2)	1. Reseat the cable between system board and the power supply (power cage assembly) (J10). 2. Power cage assembly. 3. Power supply. 4. System board.
Failure reading I²C device. (Check devices on bus 3)	1. Reseat the cable between the DASD backplane and connector (J10) of system board. 2. DASD backplane. 3. System board.
Failure reading I²C device (Check devices on bus 4)	System board

DASD checkout

Table 31. DASD checkout messages

Message	Action
Hard disk drive "*X*" removal detected (level: critical; hard disk drive "*X*" has been removed)	Information only; take action as appropriate.

Engine shutdown

Refer to Table 32 and Table 33 on page 150 when experiencing appliance engine shutdown related to voltage or temperature problems.

Voltage-related appliance engine shutdown

Table 32. Voltage-related shutdown

Message	Action
System shutoff due to *x* current over max value (level-critical; system drawing too much current on voltage *x* bus)	See "Power checkout" on page 105.
System shutoff due to *x* V over voltage (level-critical; system shutoff due to *x* supply over voltage)	1. Check the power-supply connectors. 2. Power supply. 3. Power cage assembly.

Table 32. Voltage-related shutdown (continued)

Message	Action
System shutoff due to *x* V under voltage (level-critical system shutoff due to *x* supply under voltage)	1. Check the power-supply connectors. 2. Power supply. 3. Power cage assembly.
System shutoff due to VRM x over voltage	Replace VRM ***x***.
System shutoff due to excessive (< 240 VA) loading	1. See “Power checkout” on page 105. 2. Cycle ac on/off.

Temperature-related appliance engine shutdown

Table 33. Temperature related shutdown

Message	Action
System shutoff due to board over temperature (level-critical; board is over temperature)	1. Ensure that the system is being correctly cooled; refer to “Specifications” in Chapter 1 of the hardware installation guide. 2. Replace board.
System shutoff due to CPU x over temperature (level-critical; CPU *x* is over temperature)	1. Ensure that the system is being correctly cooled; refer to “Specifications” in Chapter 1 of the hardware installation guide. 2. Replace CPU *x*.
System shutoff due to CPU *x* under temperature (level-critical; CPU *x* is under temperature)	Ambient temperature must be within normal operating specifications; refer to “Specifications” in Chapter 1 of the hardware installation guide.
System shutoff due to DASD temperature (sensor *x*) (level-critical; DASD area reported temperature outside recommended operating range)	Ensure that the system is being correctly cooled; refer to “Specifications” in Chapter 1 of the hardware installation guide.
System shutoff due to high ambient temperature (level-critical; high ambient temperature)	Ambient temperature must be within normal operating specifications; refer to “Specifications” in Chapter 1 of the hardware installation guide.
System shutoff due to system board under temperature (level-critical; system board is under temperature)	Ambient temperature must be within normal operating specifications; refer to “Specifications” in Chapter 1 of the hardware installation guide.

Temperature error messages

Table 34. Temperature error messages

Message	Action
DASD Over Temperature (level-critical; direct access storage device bay *x* was over temperature)	Ensure that the system is being correctly cooled; refer to “Specifications” in Chapter 1 of the hardware installation guide.
DASD Over recommended Temperature (sensor *x*) (level-warning; DASD bay *x* had over temperature condition)	Ensure that the system is being correctly cooled; refer to “Specifications” in Chapter 1 of the hardware installation guide.
DASD under recommended temperature (sensor *x*) (level-warning; direct access storage device bay *x* had under temperature condition)	Ambient temperature must be within normal operating specifications; refer to “Specifications” in Chapter 1 of the hardware installation guide.

Table 34. Temperature error messages (continued)

Message	Action
DASD Over Temperature (level-critical; sensor for DASD1 reported temperature over recommended range)	Ensure that the system is being correctly cooled; refer to "Specifications" in Chapter 1 of the hardware installation guide.
Power supply *x* Temperature Fault (level-critical; power supply *x* had over temperature condition)	1. Ensure that the system is being correctly cooled; refer to "Specifications" in Chapter 1 of the hardware installation guide. 2. Replace power supply *x*
System board is over recommended temperature (level-warning; system board is over recommended temperature)	1. Ensure that the system is being correctly cooled; refer to "Specifications" in Chapter 1 of the hardware installation guide. 2. Replace the system board.
System board is under recommended temperature (level-warning; system board is under recommended temperature)	Ambient temperature must be within normal operating specifications; refer to "Specifications" in Chapter 1 of the hardware installation guide.
System over temperature for CPU *x* (level-warning; CPU x reporting over temperature condition)	Ensure that the system is being correctly cooled; refer to "Specifications" in Chapter 1 of the hardware installation guide.
System under recommended CPU *x* temperature (level-warning; system reporting under temperature condition for CPU x)	Ambient temperature must be within normal operating specifications; refer to "Specifications" in Chapter 1 of the hardware installation guide.

Host Built-In Self Test

Table 35. Host built-in self test messages

Note: If the replacement procedure in not in the *IBM TotalStorage NAS 300G Hardware Installation Guide*, then the component should be replaced by a field service technician.	
Message	**Action**
Host fail (level-informational; built-in self-test for the host failed)	1. Reseat the microprocessor. 2. Reseat the VRM. 3. Replace the microprocessor CPU.

Undetermined problems

Use the information in this section if the diagnostic tests did not identify the failure, the devices list is incorrect, or the system is inoperative. Make sure that the default settings are loaded in the Configuration/Setup Utility.

Notes:

1. Damaged data in CMOS can cause undetermined problems.
2. Damaged data in BIOS code can cause undetermined problems.

Check the LEDs on all the power supplies. If the LEDs indicate the power supplies are working correctly, complete the following steps:

1. Power off the engine.
2. Be sure that the engine is cabled correctly.
3. Remove or disconnect the following devices (one at a time) until you find the failure (turn on the engine and reconfigure each time):
   Any external devices
   Surge suppressor device (on the engine)

Mouse or non-IBM devices
Each adapter
Drives
Memory modules (minimum requirement = 512 MB [2 banks of 256 MB DIMMs])

Note: Minimum operating requirements are:
a. One power supply
b. Power cage assembly
c. System board
d. One microprocessor and VRM
e. Memory module (with a minimum of two 256 MB DIMMs)

4. Power on the engine. If the problem remains, suspect the following parts in the order listed:
   Power supply
   Power cage assembly
   System board

Notes:

1. If the problem goes away when you remove an adapter from the system and replacing that adapter does not correct the problem, suspect the system board.
2. If you suspect a networking problem and all the system tests pass, suspect a network cabling problem external to the system.

Problem determination tips

Due to the variety of hardware and software combinations that can be encountered, use the following information to assist you in problem determination. If possible, have this information available when requesting assistance from Service Support and Engineering functions.

- Machine type 5196 and Model G27
- Microprocessor or hard disk upgrades
- Failure symptom
  - Do Base System Diagnostics fail?
  - What, when, where, single, or multiple systems?
  - Is the failure repeatable?
  - Has this configuration ever worked?
  - If it has been working, what changes were made prior to it failing?
  - Is this the original reported failure?
- Base System Diagnostics version
  - Type and version level
- Hardware configuration
  - Print (print screen) configuration currently in use
  - BIOS level
- Operating system software
  - Type and version level

Note: To eliminate confusion, identical systems are considered identical only if they:
1. Are the exact machine type and models
2. Have the same BIOS level

3. Have the same adapters/attachments in the same locations
4. Have the same address jumpers/terminators/cabling
5. Have the same software versions and levels
6. Have the same diagnostics code (version)
7. Have the same configuration options set in the system
8. Have the same setup for the operation system control files

Comparing the configuration and software set up between working and nonworking systems will often lead to problem resolution.

Appendix E. Fast!UTIL options

This appendix provides detailed configuration information for advanced users who want to customize the configuration of the FAStT Host Adapter and the connected devices. However, customizing the configuration is not recommended, as the adapter was configured specifically for the NAS Gateway 300.

The board can be configured using Fast!UTIL. Access Fast!UTIL by pressing **Alt + Q** during the adapter BIOS initialization (it may take a few seconds for the Fast!UTIL menu to appear). If you have more than one FAStT Host Adapter, Fast!UTIL prompts you to select the adapter you want to configure. After changing the settings, Fast!UTIL restarts your system to load the new parameters.

Attention: If the configuration settings are incorrect, your FAStT Host Adapter board might not function properly.

Configuration settings

This is the first selection on the Fast!UTIL Options menu. These settings configure the Fibre Channel (FC) devices and the FAStT Host Adapter to which they are attached.

Host adapter settings

From the Configuration Settings menu in Fast!UTIL, select **Host Adapter Settings**. The default settings for the FAStT Host Adapter are listed in Table 36 and described in the following paragraphs.

Table 36. Host adapter settings

Setting	Options	Default
Host Adapter BIOS	Enabled or Disabled	Disabled
Frame Size	512, 1024, 2048	2048
Loop Reset Delay	0-15 seconds	5 seconds
Adapter Hard Loop ID	Enabled or Disabled	Disabled
Hard Loop ID	0-125	125

Host Adapter BIOS

When this setting is disabled, the ROM BIOS on the FAStT Host Adapter is disabled, freeing space in upper memory. This setting must be enabled if you are booting from an FC disk drive attached to the FAStT Host Adapter. The default is Disabled.

Frame Size

This setting specifies the maximum frame length supported by the FAStT Host Adapter. The default size is 1024. If using F-Port (point-to-point) connections, change this setting to 2048 for maximum performance.

Loop Reset Delay

After resetting the loop, the firmware does not initiate any loop activity for the number of seconds specified in this setting. The default is 5 seconds.

Adapter Hard Loop ID

This setting forces the adapter to attempt to use the ID specified in the Hard Loop ID setting. The default is Disabled.

Hard Loop ID

If the Adapter Hard Loop ID setting is enabled, the adapter attempts to use the ID specified in this setting. The default ID is 0.

Selectable boot settings

Though you can access this option from the Configuration Settings menu, do not change the settings as booting from Fibre Channel is not supported.

Restore default settings

You can access this option from the Configuration Settings menu. It restores the FAStT Host Adapter default settings.

Raw NVRAM data

This option displays the adapter's NVRAM contents in hexadecimal format. This is a QLogic troubleshooting tool; you cannot modify the data.

Advanced adapter settings

You can access this option from the Configuration Settings menu. The default settings for the FAStT Host Adapter are listed in Table 37 and described in the following paragraphs.

Table 37. Advanced adapter settings

Setting	Options	Default
Execution Throttle	1-256	256
Fast Command Posting	Enabled or Disabled	Enabled
>4 GByte Addressing	Enabled or Disabled	Disabled
LUNs per Target	0, 8, 16, 32, 64, 128, 256	0
Enable LIP Reset	Yes or No	No
Enable LIP Full Login	Yes or No	Yes
Enable Target Reset	Yes or No	Yes
Login Retry Count	0-255	30
Port Down Retry Count	0-255	30
Drivers Load RISC Code	Enabled or Disabled	Enabled
Enable Database Updates	Yes or No	No
Disable Database Load	Yes or No	No
IOCB Allocation	1-512 buffers	256 buffers
Extended Error Logging	Enabled or Disabled	Disabled

Execution Throttle

Specifies the maximum number of commands executing on any one port. When a port's execution throttle is reached, no new commands are executed until the current command finishes executing. The valid options for this setting are 1-256. The default (optimum) is 256.

Fast Command Posting

Decreases command execution time by minimizing the number of interrupts. The default is Enabled.

>4 GByte Addressing.
Enable this option if the system has more than 4 GB of memory available. The default is Disabled.

LUNs per Target
Specifies the number of LUNs per target. Multiple LUN support is typically for redundant array of independent disks (RAID) boxes that use LUNs to map drives. The default is 0.

Enable LIP Reset
Determines the type of loop initialization process (LIP) reset that is used when the operating system initiates a bus reset routine. When this setting is yes, the driver initiates a global LIP reset to clear the target device reservations. When this setting is No, the driver initiates a global LIP reset with full login. The default is No.

Enable LIP Full Login
Instructs the ISP chip to log in to all ports after any LIP. The default is Yes.

Enable Target Reset
Enables the drivers to issue a Target Reset command to all devices on the loop when a SCSI Bus Reset command is issued. The default is Yes.

Login Retry Count
Specifies the number of times that the software tries to log in to a device. The default is 30 retries.

Port Down Retry Count
Specifies the number of times the software retries a command to a port returning port down status. The default is 30 retries.

Drivers Load RISC Code
When this setting is enabled, the host adapter uses the RISC firmware that is embedded in the software driver. When this setting is disabled, the software driver loads the RISC firmware found on the system. The default is Enabled.

Note: The driver being loaded must support this setting. If the driver does not support this setting, the result is the same as disabled regardless of the setting. Leaving this option enabled guaranties a certified combination of software driver and RISC firmware.

Enable Database Updates
When enabled, it allows the software to save the loop configuration information in flash memory when the system powers down. The default is No.

Note: This option usually applies to Windows NT and Windows 2000 operating environments.

Disable Database Load
When enabled, the device database is read from the Registry during driver initialization. When disabled, the device database is created dynamically during driver initialization. The default is No.

IOCB Allocation
Specifies the maximum number of buffers from the firmware's buffer pool that are allocated to any one port. The default is 256 buffers.

Extended Error Logging

Provides additional error and debug information to the operating system. When enabled, events are logged into the Windows NT or Windows 2000 Event Viewer. The default is Disabled.

Extended Firmware Settings

You can access this option from the Configuration Settings menu. The default settings for the FAStT Host Adapter are listed in Table 38 and described immediately following the table.

Table 38. Extended firmware settings

Setting	Options	Default
Extended Control Block	Enabled or Disabled	Enabled
RIO Operation Mode	0, 1, 2, 3, 4	0
Connection Options	0, 1, 2, 3	3
Class 2 Service	Enabled or Disabled	Disabled
ACK0	Enabled or Disabled	Disabled
Fibre Channel Tape Support	Enabled or Disabled	Disabled
Fibre Channel Confirm	Enabled or Disabled	Disabled
Command Reference Number	Enabled or Disabled	Disabled
Read Transfer Ready	Enabled or Disabled	Disabled
Response Timer	0-255	0
Interrupt Delay Timer	0-255	0

Extended Control Block

Enables all other extended firmware settings. The default is Enabled.

RIO Operation Mode

Specifies the reduced interrupt operation (RIO) modes, if supported by the software driver. RIO modes allow posting multiple command completions in a single interrupt (see Table 39). The default is 0.

Table 39. RIO operation modes

Option	Operation mode
0	No multiple responses
1	Multiple responses, 16 bit handles, interrupt host
2	Multiple responses, 32 bit handles, interrupt host
3	Multiple responses, 16 bit handles, delay host interrupt
4	Multiple responses, 32 bit handles, delay host interrupt

Connection Options

Defines the type of connection (loop or point to point) or connection preference (see Table 40). The default is 3.

Table 40. Connection options

Option	Type of connection
0	Loop only
1	Point-to-point only

Table 40. Connection options (continued)

2	Loop preferred, otherwise point to point
3	Point-to-point, otherwise loop

Class 2 Service
Enables Class 2 service parameters to be provided during all automatic logins (loop ports). The default is Disabled.

ACK0 Determines the type of acknowledge (ACK) used. When this setting is enabled, sequence ACK is used. When this setting is disabled, frame ACK is used. The default is Disabled.

Note: The Class 2 Service setting must be enabled to use the ACK0 setting.

Fibre Channel Tape Support
This setting is reserved for Fibre Channel tape support. The default is Disabled.

Fibre Channel Confirm
This setting is reserved for Fibre Channel tape support. The default is Disabled.

Command Reference Number
This setting is reserved for Fibre Channel tape support. The default is Disabled.

Read Transfer Ready
This setting is reserved. The default is Disabled.

Response Timer
Contains the value (in 100-microsecond increments) used by a timer to limit the time waiting accumulating multiple responses. For example, if this field is 8, the time limit is 800 microseconds. The default is 0.

Interrupt Delay Timer
Contains the value (in 100-microsecond increments) used by a timer to set the wait time between accessing (DMA) a set of handles and generating an interrupt. The default is 0.

Scan Fibre Channel Devices

Use this option to scan the FC loop and list all the connected devices by loop ID. Information about each device is listed, for example, vendor name, product name, and revision. This information is useful when configuring your FAStT Host Adapter and attached devices.

Fibre Disk Utility

This option scans the FC loop bus and lists all the connected devices by loop ID. You can select a disk device and perform a low-level format or verify the disk media.

CAUTION:
Performing a low-level format destroys all data on the disk.

Loopback Data Test

This option performs a data test using an FC loop or a loopback connector.

CAUTION:
Performing this test will disrupt data if tested in a FC loop.

Select Host Adapter

Use this setting to select a specific FAStT Host Adapter if you have multiple FAStT Host Adapters in your system.

Appendix F. Communication adapters

This appendix describes the PCI adapters and their correct placement. For more detailed information on the adapters, refer to the hardware installation guide.

The NAS Gateway 300 has a Fibre Channel adapter (Fibre Channel) as standard feature. The NAS Gateway 300 also has the following optional adapters:

- Alacritech 100x4 Quad-Port Server Accelerated Adapter
- IBM Gigabit Ethernet SX Server Adapter
- IBM PCI Ultra160 SCSI adapter (LVD/SE)
- Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter
- PRO/1000 XT Server Adapter by Intel
- Qlogic 2340 1-port Fibre Channel adapter
- Qlogic 2342 2-port Fibre Channel adapter
- Remote Supervisor Adapter

PCI adapter placement

These sections explain how to determine the correct positions for PCI adapters.

Adapter placement rules

Adapter placement rules guide which adapter to install in a PCI slot. These rules consist of *priority* and *slot location*. Priority is the sequence in which you install the adapters. Determining the slot location requires eliminating already filled slots and placing the adapter into the first available slot according to its possible slot locations. The priority and slot locations are shown in Table 41. To determine the location of the slots:

1. Arrange the adapters in order, one having the highest priority, using the Priority column in Table 41.
2. Place the highest priority adapter into the first available slot location listed in the slot location column of Table 41.
3. Repeat step 2 until all adapters are installed.

Example 1: If you are installing a single IBM Gigabit Ethernet SX Server Adapter, that adapter should be installed in slot 2.

Example 2: If you are installing five adapters—one Qlogic 2340 1-port Fibre Channel adapter, two IBM Gigabit Ethernet SX Server Adapters, one IBM PCI Ultra160 SCSI adapter (LVD/SE), and one Alacritech 100x4 Quad-Port Server Accelerated Adapter— they should be installed in the following slots:

- Slot 1 — Empty
- Slot 2 — IBM Gigabit Ethernet SX Server Adapter
- Slot 3 — IBM Gigabit Ethernet SX Server Adapter
- Slot 4 — Alacritech 100x4 Quad-Port Server Accelerated Adapter
- Slot 5 — IBM PCI Ultra160 SCSI adapter (LVD/SE)
- Slot 6 — Qlogic 2340 1-port Fibre Channel adapter

Table 41. Adapter installation rules

Priority	Adapter	Slot location	Maximum quantity
1	Qlogic 2342 2-port Fibre Channel adapter	6	1

Table 41. Adapter installation rules (continued)

Priority	Adapter	Slot location	Maximum quantity
2	Qlogic 2340 1-port Fibre Channel adapter	6, 5	2
3	Remote Supervisor Adapter	1	1
4	IBM PCI Ultra160 SCSI adapter (LVD/SE)	5	1
5	Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter	2, 3, 4 5	2
6	IBM Gigabit Ethernet SX Server Adapter	2, 3, 4, 5	2
7	PRO/1000 XT Server Adapter by Intel	2, 3, 4, 5	2
8	Alacritech 100x4 Quad-Port Server Accelerated Adapter	2, 3, 4, 5	2

Adapter placement charts

The following tables show you where to install the PCI adapters.

Notes:

1. In the following charts, FC in slot 6 denotes a the SAN adapter, which must be either a Qlogic 2342 2-port Fibre Channel adapter or a Qlogic 2340 1-port Fibre Channel adapter. The SAN adapter is always in slot 6.
2. The Remote Supervisor Adapter is the only adapter that is supported in slot 1. For this reason, the Remote Supervisor Adapter can be added to any of the combinations in these charts that do not already list a Remote Supervisor Adapter .
3. An engine can be ordered with a maximum of one adapter for tape backup (either a Qlogic 2340 1-port Fibre Channel adapter or a IBM PCI Ultra160 SCSI adapter (LVD/SE)). In these charts, *Tape* refers to the tape backup adapter. If there is a tape backup adapter, it is always installed in slot 5.
4. The minimum number of Ethernet adapters in each engine is one, while the maximum number of Ethernet adapters in each engine is four.
5. The onboard Ethernet controllers on each engine of the NAS Gateway 300 are used to connect the engines. For this reason, there must be an Ethernet adapter installed in one of the engines of the NAS Gateway 300 for network connectivity. The configurations that meet this requirement and are valid for the NAS Gateway 300 appear in boldface type in the tables.

Adapter abbreviations

FC Qlogic 2340 1-port Fibre Channel adapter or Qlogic 2342 2-port Fibre Channel adapter

GB IBM Gigabit Ethernet SX Server Adapter

CEN PRO/1000 XT Server Adapter by Intel

EN4 Alacritech 100x4 Quad-Port Server Accelerated Adapter

CENA Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter

RSA Remote Supervisor Adapter

Tape Tape includes one and only one of the following:

- IBM PCI Ultra160 SCSI adapter (LVD/SE)
- Qlogic 2340 1-port Fibre Channel adapter

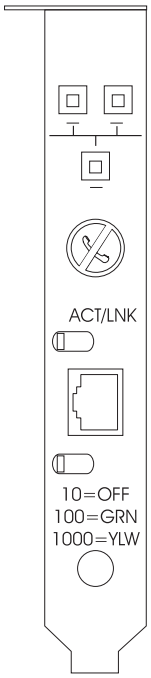

Figure 15. PRO/1000 XT Server Adapter by Intel

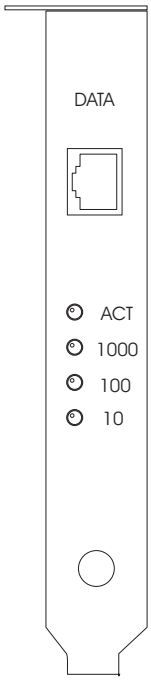

Figure 16. Alacritech 1000x1 Single-Port Server and Storage Accelerated adapter

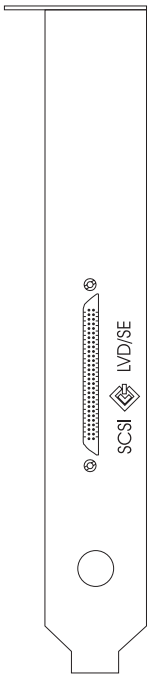

Figure 17. IBM PCI Ultra160 SCSI adapter (LVD/SE)

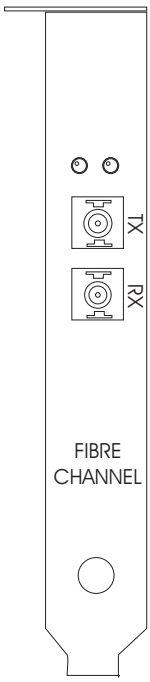

Figure 18. Qlogic 2340 1-port Fibre Channel adapter

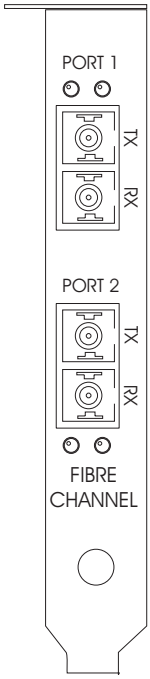

Figure 19. Qlogic 2342 2-port Fibre Channel adapter

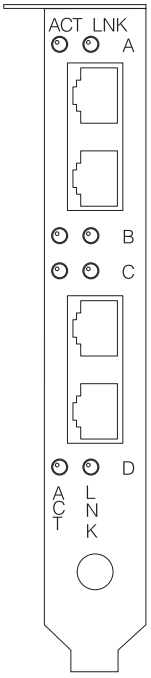

Figure 20. Alacritech 100x4 Quad-Port Server Accelerated Adapter

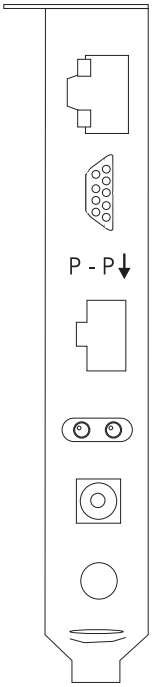

Figure 21. Remote Supervisor Adapter

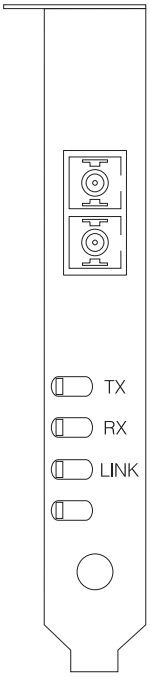

Figure 22. IBM Gigabit Ethernet SX Server Adapter

No options

Table 42 shows the adapter placement with no options.

Table 42. No options

Configuration	PCI Slot-1 (32-bit)	PCI Slot-2 (64-bit)	PCI Slot-3 (64-bit)	PCI Slot-4 (64-bit)	PCI Slot-5 (64-bit)	PCI Slot-6 (64-bit)
						FC

RSA only options

Table 43 shows the adapter placement with only an RSA option.

Table 43. RSA only options

Configuration	PCI Slot-1 (32-bit)	PCI Slot-2 (64-bit)	PCI Slot-3 (64-bit)	PCI Slot-4 (64-bit)	PCI Slot-5 (64-bit)	PCI Slot-6 (64-bit)
RSA	RSA					FC

Tape only options

Table 44 shows the adapter placement with only a tape backup option.

Table 44. Tape only options

Configuration	PCI Slot-1 (32-bit)	PCI Slot-2 (64-bit)	PCI Slot-3 (64-bit)	PCI Slot-4 (64-bit)	PCI Slot-5 (64-bit)	PCI Slot-6 (64-bit)
LVD	RSA				LVD	FC
FC1	RSA				FC1	FC

Network only options

Table 45 shows Ethernet network options without tape backup.

Table 45. Ethernet network options

Configuration	PCI Slot-1 (32-bit)	PCI Slot-2 (64-bit)	PCI Slot-3 (64-bit)	PCI Slot-4 (64-bit)	PCI Slot-5 (64-bit)	PCI Slot-6 (64-bit)
CENA	RSA	CENA				FC
GB	RSA	GB				FC
CEN	RSA	CEN				FC
EN4	RSA	EN4				FC
CENA, CENA	RSA	CENA	CENA			FC
CENA, GB	RSA	CENA	GB			FC
GB, GB	RSA	GB	GB			FC
CEN, CEN	RSA	CEN	CEN			FC
CEN, CENA	RSA	CENA	CEN			FC
CEN, GB	RSA	GB	CEN			FC
EN4, EN4	RSA	EN4	EN4			FC
EN4, CEN	RSA	CEN	EN4			FC
EN4, CENA	RSA	CENA	EN4			FC
EN4, GB	RSA	GB	EN4			FC
EN4, EN4, CEN	RSA	CEN	EN4	EN4		FC
EN4, EN4, CENA	RSA	CENA	EN4	EN4		FC
EN4, EN4, GB	RSA	GB	EN4	EN4		FC
EN4, CEN, CEN	RSA	CEN	CEN	EN4		FC
EN4, CEN, CENA	RSA	CENA	CEN	EN4		FC
EN4, CEN, GB	RSA	GB	CEN	EN4		FC
EN4, CENA, CENA	RSA	CENA	CENA	EN4		FC
EN4, CENA, GB	RSA	CENA	GB	EN4		FC
EN4, GB, GB	RSA	GB	GB	EN4		FC
CEN, CEN, CENA	RSA	CENA	CEN	CEN		FC
CEN, CEN, GB	RSA	GB	CEN	CEN		FC
CEN, CENA, CENA	RSA	CEN	CENA	CEN		FC
CEN, CENA, GB	RSA	CENA	GB	CEN		FC
CEN, GB, GB	RSA	GB	GB	CEN		FC
CENA, CENA, GB	RSA	CENA	CENA	GB		FC
CENA, GB, GB	RSA	CENA	GB	GB		FC
EN4, EN4, CEN, CENA	RSA	CENA	CEN	EN4	EN4	FC
EN4, EN4, CEN, GB	RSA	GB	CEN	EN4	EN4	FC
EN4, EN4, CEN, CEN	RSA	CEN	CEN	EN4	EN4	FC
EN4, EN4, CENA, CENA	RSA	CENA	CENA	EN4	EN4	FC
EN4, EN4, CENA, GB	RSA	CENA	GB	EN4	EN4	FC

Table 45. Ethernet network options (continued)

Configuration	PCI Slot-1 (32-bit)	PCI Slot-2 (64-bit)	PCI Slot-3 (64-bit)	PCI Slot-4 (64-bit)	PCI Slot-5 (64-bit)	PCI Slot-6 (64-bit)
EN4, EN4, GB, GB	RSA	GB	GB	EN4	EN4	FC
EN4, CEN, CEN, CENA	RSA	CENA	CEN	CEN	EN4	FC
EN4, CEN, CEN, GB	RSA	GB	CEN	CEN	EN4	FC
EN4, CEN, CENA, CENA	RSA	CENA	CENA	CEN	EN4	FC
EN4, CEN, CENA, GB	RSA	CENA	GB	CEN	EN4	FC
EN4, CEN, GB, GB	RSA	GB	GB	CEN	EN4	FC
EN4, CENA, CENA, GB	RSA	CENA	CENA	GB	EN4	FC
EN4, CENA, GB, GB	RSA	CENA	GB	GB	EN4	FC
CEN, CEN, CENA, CENA	RSA	CENA	CENA	CEN	CEN	FC
CEN, CEN, CENA, GB	RSA	CENA	GB	CEN	CEN	FC
CEN, CEN, GB, GB	RSA	GB	GB	CEN	CEN	FC
CEN, CENA, CENA, GB	RSA	CENA	CENA	GB	CEN	FC
CEN, CENA, GB, GB	RSA	CENA	GB	GB	CEN	FC
CENA, CENA, GB, GB	RSA	CENA	CENA	GB	GB	FC

Tape and network options

Table 46 shows Ethernet network options with tape backup.

Table 46. Tape backup with Ethernet network option

Configuration	PCI Slot-1 (32-bit)	PCI Slot-2 (64-bit)	PCI Slot-3 (64-bit)	PCI Slot-4 (64-bit)	PCI Slot-5 (64-bit)	PCI Slot-6 (64-bit)
TAPE, CENA	RSA	CENA			Tape	FC
TAPE, GB	RSA	GB			Tape	FC
TAPE, CEN	RSA	CEN			Tape	FC
TAPE, EN4	RSA	EN4			Tape	FC
TAPE, CENA, CENA	RSA	CENA	CENA		Tape	FC
TAPE, CENA, GB	RSA	CENA	GB		Tape	FC
TAPE, CEN, CEN	RSA	CEN	CEN		Tape	FC
TAPE, CEN, CENA	RSA	CENA	CEN		Tape	FC
TAPE, CEN, GB	RSA	GB	CEN		Tape	FC
TAPE, GB, GB	RSA	GB	GB		Tape	FC
TAPE, EN4, EN4	RSA	EN4	EN4		Tape	FC
TAPE, EN4, CEN	RSA	CEN	EN4		Tape	FC
TAPE, EN4, CENA	RSA	CENA	EN4		Tape	FC
TAPE, EN4, GB	RSA	GB	EN4		Tape	FC
TAPE, EN4, EN4, CEN	RSA	CEN	EN4	EN4	Tape	FC
TAPE, EN4, EN4, CENA	RSA	CENA	EN4	EN4	Tape	FC
TAPE, EN4, EN4, GB	RSA	GB	EN4	EN4	Tape	FC
TAPE, EN4, CEN, CEN	RSA	CEN	CEN	EN4	Tape	FC
TAPE, EN4, CEN, CENA	RSA	CENA	CEN	EN4	Tape	FC

Table 46. Tape backup with Ethernet network option (continued)

Configuration	PCI Slot-1 (32-bit)	PCI Slot-2 (64-bit)	PCI Slot-3 (64-bit)	PCI Slot-4 (64-bit)	PCI Slot-5 (64-bit)	PCI Slot-6 (64-bit)
TAPE, EN4, CEN, GB	RSA	GB	CEN	EN4	Tape	FC
TAPE, EN4, CENA, CENA	RSA	CENA	CENA	EN4	Tape	FC
TAPE, EN4, CENA, GB	RSA	CENA	GB	EN4	Tape	FC
TAPE, EN4, GB, GB	RSA	GB	GB	EN4	Tape	FC
TAPE, CEN, CEN, CENA	RSA	CENA	CEN	CEN	Tape	FC
TAPE, CEN, CEN, GB	RSA	GB	CEN	CEN	Tape	FC
TAPE, CEN, CENA, CENA	RSA	CENA	CENA	CEN	Tape	FC
TAPE, CEN, CENA, GB	RSA	CENA	GB	CEN	Tape	FC
TAPE, CEN, GB, GB	RSA	GB	GB	CEN	Tape	FC
TAPE, CENA, CENA, GB	RSA	CENA	CENA	GB	Tape	FC
TAPE, CENA, GB, GB	RSA	CENA	GB	GB	Tape	FC

Appendix G. Fibre Channel adapter event logs

You can view event logging for the Fibre Channel adapter for troubleshooting problems using the event viewer. The detailed event code is displayed at offset 34 (hex). Table 47 gives a list of detailed event codes for the adapter. For some of the event codes, additional data will be recorded in the least significant 16 bits of the longword. Additional data might also be recorded in the longword at offset 10 (hex).

Certain codes will be logged only if you set a Fast!UTIL parameter to enable additional event logging. These codes are indicated by an asterisk (*). By default, these events are not logged.

If an error occurs that is not listed in Table 47, contact the IBM Support Center (1 800 426-7378 in the U.S.). In all other countries, contact your IBM reseller or IBM marketing representative.

Table 47. Fibre Channel adapter error codes

Event code offset 34h	More data offset 10h	Description	Suggested action
4002xxxx	yyyy00zz	Host interface error: xxxx = mailbox1; yyyy = mailbox2; zz = command	Hardware DMA error: replace adapter
4005xxxx	Yyyy00zz	Mailbox command error: xxxx = mailbox1; yyyy = mailbox2; zz = command	Normally indicates loop down, check all cabling
4005xx6F	yyyyyyzz	Login Fabric port mailbox command error: xx = adapter; state yyyyyy = port id; zz = loop id	Normally indicates loop down, check all cabling
* 80010000	00000000	Reset detected	Not logged during normal operation
8003xxxx	yyyyzzzz	RISC request queue transfer error: xxxx = mailbox1; yyyy = mailbox2; zzzz = mailbox3	Hardware error: replace adapter
8004xxxx	yyyyzzzz	RISC response queue transfer error: xxxx = mailbox1; yyyy = mailbox2; zzzz = mailbox3	Hardware error: replace adapter
* 80100000	0000xxxx	LIP occurred: xxxx = mailbox1	Not logged during normal operation
* 80110000	xxxxyyzz	Link up 2200: xxxx = current ISP connection mode (0 = Loop, 1 = P2P); yy = ISP connection option 0 = Loop, 1 = P2P, 2 = Loop->P2P, 3 = P2P->Loop; zz = starting loop id for remote devices. 2100: xxxx = 0000; yyyy = 0000	Not logged during normal operation
80120000	00000000	Link down error	Not logged during normal operation
80130000	0000xxxx	LIP reset occurred: xxxx = mailbox1	Not logged during normal operation

Table 47. Fibre Channel adapter error codes (continued)

Event code offset 34h	More data offset 10h	Description	Suggested action
**80300000	xxxxyyzz	Link mode up: xxxx = current ISP connection mode 0 = Loop 1 = P2P; yy = ISP connection option (0 = Loop, 1 = P2P 2 = Loop->P2P 3 = P2P->Loop); zz = starting loop id for remote devices	Not logged during normal operation
**8036aabb	xxxxyyzz	Point-to-Point update configuration: xxxx = mailbox1; yy = current ISP connection mode (0 = Loop, 1 = P2P); zz = ISP connection option 0 = Loop, 1 = P2P, 2 = Loop->P2P, 3 = P2P->Loop; aa = starting loop id value for remote devices; bb = current retry count for the ISP initialization mode	Not logged during normal operation
* F0000000	00000000	Restarting RISC firmware	Initial driver load or loop down longer than 4 minutes
* F0030004	00xx00yy	Reset command completion error: xx = CD-ROMB opcode; yy = target loop ID	Not logged during normal operation
* F0030005	00xx00yy	Command aborted by OS: xx = CD-ROMB opcode; yy = target loop ID	Not logged during normal operation
F0030028	00xx00yy	Port unavailable, command completion error: xx = CD-ROMB; opcode yy = target loop ID	Check target device and cabling
F0030029	00xx00yy	Port logged out command completion error: xx = CD-ROMB; opcode yy = target loop ID	Check target device and cabling
F003001C	00xx00yy	Target device queue full (SCSI status 28 from target): xx = CD-ROMB opcode; yy = target loop ID	Check target device and cabling
* F00A0000	0000xxxx	RISC firmware state during adapter initialization: xxxx = firmware state	Not logged during normal operation
F00B0000	00000000	Reset ISP chip failed	
F00D0000	00000000	Fail to allocate non-cached memory	
F00E0000	00000000	Fail to map ISP registers	
F00F0000	00000000	Fail to load RISC code	
F0100000	0000xxxx	Fail to start RISC code: xxxx = mailbox0	
F0110000	0000xxxx	Fail to initialize firmware xxxx = mailbox0	
F0120000	0000xxxx	Fail to get firmware state: xxxx = mailbox0	
* F0130000	00000000	Port Update notification (RISC database changed)	

Table 47. Fibre Channel adapter error codes (continued)

Event code offset 34h	More data offset 10h	Description	Suggested action
* F0140000	xxxxxxxx	RSCN notification (Name server change detected): xxxx = RSCN Information	
* F0150000	00xx00yy	Name server query rejected (v6 2100): xx = Reason Code; yy = Explanation Code (Valid if reason code is 0x09)	
* F0150000	xxxxyyzz	Name server query rejected (v7 2100/2200): xxxx = response status; yy = Reason Code; zz = Explanation Code (Valid if reason code is 0x09, for example, if zz = 0x09, yy = 07, this means no SCSI device found.)	
* F0160000	00000000	Driver reset called; command timed out	
* F0170000	00xxxxxx	Fabric port login (for information only): xxxxxx = Port Id	
F0180000	000000xx	Excessive link errors, loop down: xx = number of link errors per second	
* F0190000	00000000	Verify firmware checksum failure	
* F01B0000	000000xx	Device marked offline after being not-ready longer than port down retry count: xx = loop ID of device	
* F01C0000	000000xx	Bad type field in IOCB from RISC: xx = IOCB type	
* F01D0000	00000000	Error downloading post RISC code	
* F01Exxxx	Yyyyzzzz	Error running post RISC code: xxxx = mailbox0; Yyyy = mailbox1; zzzz = mailbox2	
* F01Fxxyy	Zzzzzzzz	DMA 64 bit (PAE) configuration (for information only): xx = Dma64BitAddressess flag set by W2K; yy = Dma64BitAddressess flag set by driver; zzzzzzzz = driver adapter flags	
F0200000	Xxxxyyyy	Error ISP not accessible: xxxx = ISP host command and control; yyyy = ISP interrupt status	
* F0210000	xxyy00zz	ISP connection option/topology (for information only): xx = ISP connection option from NVRAM; yy = previous ISP topology; zz = current ISP topology code: 0000 = Loop, 0001 = FL_Port, 0002 = N_Port to N_Port, 0003 = F_Port	
* F0220000	0000xxxx	External RISC ram parity error (for 2200G only): xxxx = number of parity errors detected	

Table 47. Fibre Channel adapter error codes (continued)

Event code offset 34h	More data offset 10h	Description	Suggested action
* F0230000	Xxxxyyyy	Subvendor Id not match (for information only): xxxx = actual subvendor id; yyyy = expected subvendor id	

Glossary of terms and abbreviations

This glossary includes terms and definitions from:

- *The American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies can be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. Definitions are identified by the symbol (A) after the definition.
- The *ANSI/EIA Standard - 440A: Fiber Optic Terminology*, copyright 1989 by the Electronics Industries Association (EIA). Copies can be purchased from the Electronics Industries Association, 2001 Pennsylvania Avenue N.W., Washington, D.C. 20006. Definitions are identified by the symbol (E) after the definition.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.

A

adapter load balancing. The ability of several adapters in a team to be active simultaneously, with the outbound-traffic load balanced across all the adapters in the team; spreading tasks among adapters improves performance by preventing uneven distribution of workload. If one adapter in the team fails, the outbound traffic is redistributed across the remaining active adapters in the team. See also *teaming*.

assigned disk. A disk that is mapped to a logical drive.

attachment. A port or a pair of ports, optionally including an associated optical bypass, that are managed as a functional unit. A dual attachment includes two ports: a port A and a port B. A single attachment consists of one port: port S.

B

Basic Input/Output System (BIOS). The personal computer code that controls basic hardware operations, such as interactions with diskette drives, hard disk drives, and the keyboard.

BIOS. See *Basic Input/Output System*.

bits per second (bps). The rate at which bits are transmitted per second. Contrast with *baud*.

bps. See *bits per second.*

buffer. See *buffer storage*.

buffer storage. (1) A special-purpose storage or storage area allowing, through temporary storage, the data transfer between two functional units having different transfer characteristics. A buffer storage is used between non-synchronized devices, a serial and a parallel device, or between devices having different transfer rates. (2) In word processing, a temporary storage in which text is held for processing or communication (T).

bus. See *data bus*.

C

CIFS. See *Common Internet File System.*

client. A computer system or process that requests access to the data, services, or resources of a server (another computer system or process). Multiple clients may share access to a common server.

client-server model. A common way to describe network services and the model user processes (programs) of those services.

cluster. (1) A station that consists of a control unit (a *cluster controller*) and the terminals attached to it. (2) A group of APPN nodes that have the same network ID and the same topology database. A cluster is a subset of a network identifier (NETID) subnetwork. See also *high-availability cluster multiprocessing (HACMP)* and *network identifier (NETID)*.

collision avoidance. In carrier sense multiple access with collision avoidance (CSMA/CA), the process of sending a jam signal and waiting for a variable time before transmitting data. The process is designed to avoid two or more simultaneous transmissions.

Common Internet File System (CIFS). A protocol that enables collaboration on the Internet by defining a remote file-access protocol that is compatible with the way applications already share data on local disks and network file servers.

connect. In a LAN, to physically join a cable from a station to an access unit or network connection point. Contrast with *attach*.

CRU. See *customer-replaceable unit*.

customer-replaceable unit (CRU). An assembly or part that a customer can replace in its entirety when any of its components fail. Contrast with *field-replaceable unit* .

D

DASD. See *direct access storage device*.

data bus. A bus used to communicate data internally and externally to and from a processing unit, storage, and peripheral devices (A).

device parity protection. A function that protects data stored on a disk-unit subsystem from being lost because of the failure of a single disk unit in the disk-unit subsystem. When a disk-unit subsystem has device parity protection and one of the disk units in the subsystem fails, the subsystem continues to run. The disk-unit subsystem reconstructs the data after the disk unit in the subsystem is repaired or replaced. See also *RAID*.

DHCP. See *Dynamic Host Configuration Protocol*.

DIMM. See *dual inline memory module*.

direct access storage device (DASD). A mass-storage medium on which a computer stores data. Contrast with *random access memory (RAM)*.

DNS. See *Domain Name System*.

Domain Name System (DNS). In the Internet suite of protocols, the distributed database system used to map domain names to IP addresses.

drive bay. A receptacle in an appliance into which you insert a hard-disk-drive module. The bays are in storage units that can be located in a different rack from the appliance.

dual inline memory module (DIMM). A small circuit board with memory-integrated circuits containing signal and power pins on both sides of the board.

Dynamic Host Configuration Protocol (DHCP). A protocol defined by the Internet Engineering Task Force (IETF) that is used for dynamically assigning IP addresses to computers in a network.

E

EISA. See *Extended Industry Standard Architecture*.

electromagnetic compatibility (EMC). The design and test of products to meet legal and corporate specifications dealing with the emissions and susceptibility to frequencies in the radio spectrum. Electromagnetic compatibility is the ability of various electronic equipment to operate correctly in the intended electromagnetic environment.

electrostatic discharge (ESD). An undesirable discharge of static electricity that can damage equipment and degrade electrical circuitry.

EMC. See *electromagnetic compatibility*.

engine. The unit that contains the processors that respond to requests for data from clients. The operating software for the IBM TotalStorage appliance resides in the engine.

equivalent paths. A collection of paths to the storage device. The paths have no switchover time penalty when changing from one path group to another while accessing the storage device.

error. A discrepancy between a computed, observed, or measured value or condition and the true, specified, or theoretically correct value or condition (A) (I). Contrast with *failure*.

ESD. See *electrostatic discharge*.

Ethernet. A standard protocol for a 10-Mbps baseband local area network (LAN) that allows multiple access and manages contention by using carrier sense multiple access with collision detection (CSMA/CD) as the access method.

Ethernet network. A baseband LAN with a bus topology in which messages are broadcast on a coaxial cable using a carrier sense multiple access/collision detection (CSMA/CD) transmission method.

expansion slot. In personal-computer systems, one of several receptacles in the rear panel of the system unit into which a user can install an adapter.

Extended Industry Standard Architecture (EISA). The PC bus standard that extends the AT bus (ISA bus) to 32 bits and provides support for bus master. It was announced in 1988 as a 32-bit alternative to the Micro Channel that would preserve investment in existing boards. PC and AT cards (ISA cards) can plug into an EISA bus.

F

fabric. A complex network using hubs, switches and gateways. For example, Fibre Channel uses a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices.

failback. The restoration of an appliance to its initial configuration after detection and repair of a failed network or component.

failover. (1) The automatic recovery of resources in the event of a network outage, or failure of the hardware or software. (2) A cluster event in which the primary database server or application server switches to a backup system due to the failure of the primary server.

Fast Etherchannel (FEC). A proprietary technology developed by Cisco that creates a team of two to four 10/100 Ethernet adapters or ports to increase transmission and reception throughput. Adapter fault tolerance is also supported by this technology.

feature code. A code used by IBM to process hardware and software orders.

FEC. See *Fast Etherchannel*.

fiber optic cable. See *optical cable*.

field-replaceable unit (FRU). An assembly that is replaced in its entirety when any one of its components fails. In some cases, a FRU may contain other FRUs. Contrast with *customer-replaceable unit*.

flash memory. A type of non-volatile storage device that must be erased in fixed blocks rather than single bytes.

FRU. See *field-replaceable unit*.

G

gateway. A device that acts as a router to transfer packets between networks, but occurs at the transport layer. See also *router*.

GEC. See *Gigabit Etherchannel*.

Gigabit Etherchannel (GEC). A proprietary technology developed by Cisco that creates a team of two Gigabit Ethernet adapters to increase transmission and reception throughput. Adapter fault tolerance is also supported by this technology.

H

hertz (Hz). A unit of frequency equal to one cycle per second.

Note: In the United States, line frequency is 60 Hz or a change in voltage polarity 120 times per second; in Europe, line frequency is 50 Hz or a change in voltage polarity 100 times per second.

host. (1) In TCP/IP, any system that has at least one Internet address associated with it. A host with multiple network interfaces may have multiple internet addresses associated with it. The host can be a client, a server, or both. (2) In Fibre Channel technology, any system that has at least one worldwide name associated with it. A host with multiple network interfaces may have multiple worldwide names associated with it.

I

IETF. See *Internet Engineering Task Force*.

iLUN. See *iSCSI client logical-unit number.*

Internet Engineering Task Force (IETF). The task force of the Internet Architecture Board (IAB) that is responsible for solving the short-term engineering needs of the Internet. The IETF consists of numerous working groups, each focused on a particular problem. Internet standards are typically developed or reviewed by individual working groups before they can become standards.

Internet Protocol (IP). A protocol that routes data through a network or interconnected networks. IP acts as an intermediary between the higher protocol layers and the physical network.

interrupt request (IRQ). An input found on a processor that causes it to suspend normal instruction execution temporarily and to start executing an interrupt handler routine.

IP. See *Internet Protocol*.

IRQ. See *interrupt request*.

iSCSI client logical-unit number (iLUN). A unique number that is assigned to each virtual logical unit number (VLUN). The iLUN for a single client starts at zero and increments sequentially.

J

jumper. A connector between two pins on a network adapter that enables or disables an adapter option, feature, or parameter value.

L

LAN. See *local area network*.

local area network (LAN). A network in which a set of devices is connected to one another for communication and that can be connected to a larger network.

logical drive. A unit of virtual storage that is made available to the network through virtual logical unit numbers (VLUNs) and iSCSI client logical-unit number (iLUNs). It consists of one or more physical disks that are combined using RAID 0, 1, 1E, 5, or 5E technology.

logical unit. A type of network-accessible unit that enables users to gain access to network resources and communicate with each other.

logical unit number (LUN). An identifier used on a SCSI bus to distinguish among up to eight devices (logical units) with the same SCSI ID.

loop. A closed unidirectional signal path connecting input/output devices to a system.

LUN. See *logical unit number*.

M

megahertz (MHz). A unit of measure of frequency. One megahertz equals 1 000 000 hertz.

MHz. See *megahertz*.

multicast address. A type of IP address, which identifies a group of interfaces and permits all of the systems that are in that group to receive the same packet of information.

N

N. See *newton*.

NAS. See *network-attached storage*.

NetBIOS. A standard interface to networks, IBM personal computers (PCs), and other compatible PCs. It is used on LANs to provide message, print-server, and file-server functions. Application programs that use NetBIOS need not manage the details of LAN data-link-control protocols.

network-attached storage (NAS). A task-optimized storage device directly attached to a network that operates independently of the general-purpose file servers.

Network File System (NFS). A protocol, developed by Sun Microsystems, Incorporated, that allows any host in a network to mount another host's file directories. After a file directory is mounted, it appears to reside on the local host.

network information services (NIS). A set of UNIX network services (for example, a distributed service for retrieving information about the users, groups, network addresses, and gateways in a network) that resolve naming and addressing differences among computers in a network.

newton (N). The unit of force required to impart an acceleration of one meter per second per second to a mass of one kilogram (1 m/s^2).

NFS. See *Network File System*.

NIS. See *network information services*.

O

optical cable. A fiber, multiple fibers, or a fiber bundle in a structure built to meet optical, mechanical, and environmental specifications (E).

P

path. In a network, a route between two nodes.

path group. A collection of equivalent paths. Storage devices may have one - *n* path groups.

PCI. See *Peripheral Component Interconnect*.

Peripheral Component Interconnect (PCI). A local bus for PCs from Intel that provides a high-speed data path between the CPU and up to 10 peripherals (video, disk, network, and so on). The PCI bus coexists in the PC with the industry standard architecture (ISA) or extended industry standard architecture (EISA) bus. ISA and EISA boards plug into an ISA or EISA slot, while high-speed PCI controllers plug into a PCI slot.

Persistent Storage Manager (PSM). Columbia Data Products software that creates multiple, point-in-time, persistent, TrueImage data views of any or all system and data volumes residing on network-attached storage. All persistent images survive system power loss, or a planned or unplanned reboot. Each instance of PSM seamlessly handles 250 concurrent images of up to 255 independent volumes for a total of 63 750 independent data images.

port. See *socket*.

PSM. See *Persistent Storage Manager*.

R

RAID. See *redundant array of independent disks*.

RAM. See *random access memory*.

random access memory (RAM). A temporary storage location in which the central processing unit (CPU) stores and executes its processes. Contrast with *direct access storage device (DASD)*.

redundant array of independent disks (RAID). A method of protecting data loss due to disk failure based on the Redundant Array of Independent Disks specification published by the University of California in 1987. See also *device parity protection*.

S

SAN. See *storage area network*.

SCSI. See *small computer system interface*.

server. In a network, a node that provides facilities to other stations; examples of servers are a file server, a printer server, and a mail server.

shielded twisted pair (STP). A cable medium consisting of a telephone wire wrapped in a metal sheath to eliminate external interference.

Simple Network Management Protocol (SNMP). In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application-layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB).

small computer system interface (SCSI). A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

SNMP. See *Simple Network Management Protocol*.

storage area network (SAN). A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services.

storage client network. A classic, interconnected, Fibre Channel fabric with a single, Fibre Channel, fabric name.

storage controller. A device (such as a RAID controller) that creates and manages other storage devices.

storage network. An arrangement that provides shared access to a set of logical unit numbers (LUNs) across one - *n* storage client networks.

storage port. An engine's connection point to a storage client network. A storage port is a member of a single fabric. See also *engine*.

storage unit. Hardware that contains one or more drive bays, power supplies, and a network interface. Some storage units contain RAID controllers; their storage unit is accessed by the appliance.

STP. See *shielded twisted pair*.

T

target. A collection of logical units that are directly addressable on the network. The target corresponds to the server in a client-server model.

TCP. See *Transmission Control Protocol*.

TCP/IP. See *Transmission Control Protocol/Internet Protocol*.

teaming. The grouping of two to four ports or adapters to increase transmission and reception throughput. Teaming creates a single, high-speed, fault-tolerant link that provides load balancing for both outbound and inbound traffic.

Tivoli Storage Manager (TSM). A client/server product that provides storage management and data access services in a heterogeneous environment.

Transmission Control Protocol (TCP). In TCP/IP, a host-to-host protocol that provides transmission in an internet environment. TCP assumes Internet Protocol (IP) is the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP). The Transmission Control Protocol and the Internet Protocol, which together provide reliable end-to-end connections between applications over interconnected networks of different types.

True Image data view. A data view that allows the file to be restored in the event of accidental deletion. It consists of point-in-time images that provide a near-instant virtual copy of an entire storage volume.

TSM. See *Tivoli Storage Manager*.

U

universal serial bus (USB). A serial-interface standard for telephony and multimedia connections to personal computers.

unshielded twisted pair (UTP). A cable medium with one or more pairs of twisted insulated copper conductors bound in a single plastic sheath.

USB. See *universal serial bus*.

V

virtual local area network (VLAN). A logical association of switch ports based upon a set of rules or criteria such as MAC addresses, protocols, network address, or multicast address. This concept permits resegmentation of the LAN without requiring physical rearrangement.

virtual logical unit number (VLUN). A subset of a logical drive.

VLAN. See *virtual local area network*.

VLUN. See *virtual logical unit number*.

volume. (1) A unit of storage on disk, tape, or other data-recording media. (2) A logical disk visible to an appliance over a storage network. A volume is a member of a single storage network of 1 - *n* fabrics. It can have 1 - *n* path groups of 1 - *n* equivalent paths.

W

Windows Internet Naming Service (WINS). A Microsoft program that provides a distributed database for registering and querying dynamic NetBIOS names to IP address mapping in a routed network environment.

WINS. See *Windows Internet Naming Service*.

Windows networking. A networking file-system protocol for the Windows operating system.

Index

H

I

K

L

M

N

O

P

Q

R

S

T

U

W

Printed in U.S.A.

GA27-4321-00