

IBM TotalStorage™ Network Attached Storage 300G
(Model G25)



Release Notes

Network Attached Storage 300G

4/23/01

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED ON AN "AS IS" BASIS. In no event will IBM be liable for damages arising directly or indirectly from any use of the information contained in this document.

First Edition (April 23, 2001)

This edition applies to the IBM 5196 TotalStorage Network Attached Storage 300G (Model G25) (product number 5196-G25).

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 may address your comments to:

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

You can also submit comments on the Web at www.ibm.com/networking/support/feedback.nsf/docsoverall.

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.

© **Copyright International Business Machines Corporation 2001. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Release notes	1
Functional notes	1
Failover for redundant Ethernet	1
ASM PCI adapter - onboard planar communication on a single machine	1
Administration notes	2
Determining who is using the network-attached storage	2
Recovery procedure for the system hard disk drive	2
AntiVirus protection.	3
Memory notes.	3
Clustering installation and initial configuration	4
Cluster installation requirements	4
Cluster installation	7
Cluster configuration	13
Checking the FastT host-bus adapter settings and fibre-channel connectivity	16

Release notes

This document contains the release notes for this storage product.

Refer to www.storage.ibm.com/support/nas for possible updates to these release notes and for copies of the base documents.

Functional notes

The following section(s) contain information on 300G function.

Failover for redundant Ethernet

The 10/100 Ethernet adapter that is built-in (on the planar) of the engine is not redundant-capable. However, the IBM 10/100 Ethernet Server Adapters (NIC adapter) that you install in the PCI card slots are redundant-capable. If you install two NIC adapters and connect them to the same logical segment, you can configure the server to support an Ethernet-adapter failover function.

In failover mode, if the primary Ethernet adapter detects a link failure, all Ethernet traffic associated with it is switched to the redundant (secondary) adapter. This switching occurs without your intervention. When the primary link is restored to an operational state, the Ethernet traffic switches back to the primary Ethernet adapter. The switch back to the primary Ethernet adapter can be automatically or manually controlled, depending on the setup.

ASM PCI adapter - onboard planar communication on a single machine

You will need to follow one of two methods to enable communication between the ASM PCI adapter and the onboard planar on a single machine:

1. Using a single ASM interconnect cable (with dual RJ-11 plugs)
 - a. Connect one end of the internal ASM interconnect cable to the J-40 connector on the system board.
 - b. Connect the other end (the RJ-11 socket) of the internal ASM interconnect cable to the knockout slot on the back panel of the machine until it locks into place.
 - c. Connect one connector on the ASM interconnect cable to the RJ-11 socket you just installed on the back panel (in step 1b).
 - d. Connect the other connector to the RJ-11 socket on the ASM PCI adapter.
2. Using two ASM interconnect cables (each with a single RJ-11 plug)
 - a. Connect one end of the internal ASM interconnect cable to the J-40 connector on the system board.
 - b. Connect the other end (with the RJ-11 socket) of the internal ASM interconnect cable to the knockout slot on the back panel of the machine until it locks into place.
 - c. Connect the first ASM interconnect cable to the RJ-11 socket you just installed on the back panel (in step 2b).
 - d. Connect the second ASM interconnect cable to the RJ-11 socket on the ASM PCI adapter.
 - e. Connect the two ASM interconnect cables with a single Category 5 Ethernet cable (by plugging one end of the Ethernet cable into the “black box” on the first ASM interconnect cable, and the other end into the “black box” on the second ASM interconnect cable).

Administration notes

The following section(s) contain information on administering the 300G.

Determining who is using the network-attached storage

Occasionally, the administrator may want to know who is using the network-attached storage. The administrator can determine this information as follows:

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

Recovery procedure for the system hard disk drive

Chapter 7 of the *User's Reference* documents the Recovery CD-ROM procedure. This section documents additional steps and considerations not included in the *User's Reference*. Please use the following procedure to recover a ("failed") 300G node.

1. On the other ("working") node of the 300G, 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. Underneath it, locate the name of the failed node, right-click on 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 working node of the 300G.
3. Use the Recovery CD-ROM procedure described in Chapter 7 of the *User's Reference* on the failed node.

Note: If you have installed more processor memory on the failed node, you will need to attach a keyboard and display to the node to step through the menu-driven BIOS configuration program which appears during two of the node restarts that take place during the recovery process (as a result of the additional memory having been installed).

The BIOS configuration program appears for the first time when you restart the failed node from the Recovery CD-ROM (after you use the IBM TotalStorage Recovery Enablement Diskette to allow the node to restart from CD-ROM). The BIOS configuration program appears the second time when you restart the node without the Recovery CD-ROM (after the Recovery CD-ROM has finished copying its preloaded software image onto the node's internal hard disk drive).

To step through the BIOS configuration program, click **Continue** on the first screen, then click **Continue** again, then click **Exit Setup**, and finally, click **Yes, save and exit Setup**. Once you have done this for the second time, you may detach the keyboard and display from the failed node and allow the recovery process to complete automatically.

4. Reinstall all software updates you had installed on the failed node since you installed the 300G out of the box. Or, if the Recovery CD-ROM you used in step 3 on page 2 is a newer version than the one you received with the 300G, reinstall only those software updates that are newer than those on the Recovery CD-ROM.
5. 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 6.
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, then click the **IBM NAS Admin icon**, and then from the IBM NAS Administration console that appears, select **Computer Management**, then **Disk Management**.
 - b. Start a Windows 2000 for NAS user interface session to the node, then select **Disks and Volumes**, then select **Disks and Volumes** again, and then provide your administrator user name and password when prompted.

Once Disk Management has started, do the following:

- a. In the Disk Management window, right-click on 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 and change the volume label to *Maintenance*.
- e. Click **Finish** to close the wizard.

The partition will then be formatted. Once 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*, size (approximately) *5.9 GB*.

6. On the failed (now “recovered”) node, follow the procedures outlined in “Clustering installation and initial configuration” on page 4 for configuring a joining node. The recovered node will rejoin the cluster which already contains the other (working) node. You will also need to reconfigure any load balancing you had set up prior to recovery, such 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. See “Load balancing” on page 13 for more details on configuring load balancing.

AntiVirus protection

You can perform AntiVirus scanning of 300G storage from clients having the appropriate access permissions. Also, you can install Norton AntiVirus Version 7.1 or later on the 300G engine using normal Windows 2000 software installation procedures.

Depending on configuration options, AntiVirus scanning may use substantial CPU or disk resources. Therefore, you should carefully select scanning options and scan schedule.

Memory notes

The following section(s) contain information on adding memory.

Adding more engine memory to increase performance

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

1. Purchase one or two 512M memory SIMMs (part number 33L3127) from Options by IBM at www.pc.ibm.com/us/options/
2. Follow the instructions in Chapter 3, section *Replacing memory modules*, of the *Installation Guide*.
3. Before rebooting the box, 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 the *User's Reference*), you will have to attach a keyboard and display and answer questions about the additional memory you have installed.

Clustering installation and initial configuration

The Model G25 uses Microsoft's Cluster Server (MSCS) software to provide clustering technology for your SAN. Clustering ensures availability of your SAN, 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 Model G25 detects the failure and begins a failover process in less than 10 seconds, and completes the process within 60 seconds. Failover/Failback will include Active/Active support for the CIFS protocol and Active/Active support for the NFS protocol.

This chapter gives the details for installing and initially configuring MSCS on the Model G25. Administrative concepts and procedures are provided in the online help and at the following Web sites:

<http://www.microsoft.com/windows2000/library/technologies/cluster/default.asp>

http://www.microsoft.com/ntserver/support/faqs/clustering_faq.asp

Note: Please use the Release Notes for the Model G25 in conjunction with the information presented in this chapter.

Cluster installation requirements

Before you configure the Model G25 nodes for clustering function, ensure that the following requirements are met:

Network Requirements

- A unique NetBIOS cluster name
- 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 6 shows a summary of the networking information you will need 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.x.

Notes:

1. If you are not the system administrator, contact that person for the applicable IP addresses.
2. One additional static IP address will be required for creating each clustered group.
3. 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.
4. 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

Although not required, the use of a fault-tolerant RAID configuration is recommended for the shared storage disk arrays. Clustering provides fault-tolerant support for the resources you set up to be used by the Cluster Server, and building in additional support will increase the cluster's availability.

Table 1. Networking information worksheet for the public connection

Cluster component	Information needed
Cluster	<p>Cluster name:</p> <p>IP address:</p> <p>Subnet:</p>
First node	<p>Computer name: (For example: IBM5196-23H1234)</p> <p>IP address:</p> <p>Subnet:</p> <p>Gateway:</p> <p>Preferred DNS:</p>
Second node	<p>Computer name:</p> <p>IP address:</p> <p>Subnet:</p> <p>Gateway:</p> <p>Preferred DNS:</p>
Domain to join	<p>Domain name:</p>

Cluster installation

There are four basic steps to follow when you install cluster services on the first and second nodes and create the cluster.

After you have completed the steps for the first node (by referring to the appropriate section), complete the steps for the second node.

First node creation steps:

- | | |
|---------------------------------|---|
| 1. Networking setup | See “Networking setup (step 1 of 4)” on page 8. Note the cautionary statement at the beginning of this section. |
| 2. Shared storage setup | See “Shared storage setup (step 2 of 4)” on page 10. |
| 3. Domain setup | See “Joining a domain (step 3 of 4)” on page 11. |
| 4. Cluster configuration | See “Cluster configuration (step 4 of 4)” on page 12. |

Second node creation steps:

- | | |
|---------------------------------|---|
| 1. Networking setup | See “Networking setup (step 1 of 4)” on page 8. |
| 2. Shared storage setup | This step is not required when creating the second node. However, if you are having trouble with the Fibre Channel connection, you can use the steps in “Checking the FastT host-bus adapter settings and fibre-channel connectivity” on page 16 to diagnose the problem. |
| 3. Domain setup | See “Joining a domain (step 3 of 4)” on page 11. |
| 4. Cluster configuration | See “Cluster configuration (step 4 of 4)” on page 12. |

Note: You can carry out the following steps with a keyboard, monitor, and mouse, or you can use the IBM Advanced Appliance Configuration Utility (that is, in “headless” mode) to detect the node for initial network configuration. For details on the utility, see the section on the IBM Advanced Appliance Configuration Utility, in Chapter 3 of the *User’s Reference*.

Networking setup (step 1 of 4)

CAUTION:

Before you begin installing cluster services on the first node, power off the second node. This is required to prevent corruption of data on the shared storage devices. Corruption can occur if both nodes were to try to simultaneously write to the same shared disk that is not yet protected by the clustering software.

Configuring the interconnect (Private) network adapter: Do the following steps on both nodes to configure the interconnect (Private) network adapter. 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 IBM 10/100 Netfinity Fault Tolerant Adapter.
3. Right-click the adapter icon and click **Properties**.
4. Click **Configure**, select the **Advanced** tab, and verify that the following characteristics are set:

External PHY	100 Mbps Full Duplex
Full Duplex	UTP - Full Duplex
TP Mode	On (Default)

5. Click **OK**.
6. If prompted to restart the node, select **NO**.

After you return to the Properties panel for **My Network Places**, do the following:

1. Rename the connection to **Private**.
2. Go to the Properties panel for the connection, select **Internet Protocol (TCP/IP)** from the components section, and click **Properties**.
3. Set the IP address to 10.1.1.1 for the first node, or 10.1.1.2 for the second node, depending on which node you are installing.
4. Ensure a Subnet Mask of 255.0.0.0.
5. Click on **Advanced**, and select the **WINS** tab.
6. Select the **Disable NetBIOS over TCP/IP** radio button.
7. Click **OK**.
8. Select **YES** when you are asked if you want to continue using an empty Primary WINS address.
9. Click **OK** on the Internet Protocol (TCP/IP) Properties panel.
10. Click **OK** on the Local Area Connection Properties (Private) panel.

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.

Do the following steps on both nodes to configure the public local area connections.

1. Right-click **My Network Places**, then click **Properties**.

2. Right-click the active **Local Area Connection** (the one that is **not** the IBM 10/100 Netfinity Fault Tolerant Adapter).

When you do this step, the connection that uses the IBM 10/100 Netfinity Fault Tolerant Adapter 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. Click **Rename**, and then type **Public**, and press **Enter**.

When you do these renaming steps for the second node, ensure that the local area connection name for each physically connected network is identical on each server. See Table 2 for an example.

4. Use the networking information in Table 1 on page 6 to enter the networking addresses:
 - Right-click **My Network Places**
 - Click **Properties**
 - Right-click the **Public** icon, then click **Properties**.
 - Select **Internet Protocol (TCP/IP)**
 - Click **Properties**, select **Use the following IP address:**, and enter the addresses for the IP, subnetmask, default gateway, and preferred DNS server.
5. If needed, configure the WINS, HOSTS, or whichever method you will be using for names resolution. You can view this information by clicking the **Advanced** button on the Properties window.
6. Click **OK** on each panel to back up to the Properties window, then continue with “Configuring the interconnect (Private) network adapter” on page 8.

If you are using multiple public networks, do not place additional paired adapters on the same IP network, unless you are going to use adapter teaming or adapter load balancing. Without adapter teaming or adapter load balancing, the additional paired adapters will not show up in the Cluster Server Administrator that you will set up later.

Verify Network Connectivity and Names Resolution: Verify network connectivity and names resolution after you have installed clustering on the second node.

To verify that the private and public networks are communicating properly, open an MS-DOS command prompt and 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 2. Example of local area connection names and network adapter IP addresses

Node	Local area connection name	Network adapter IP address
1	Private Cluster Connection	10.1.1.1
1	Public Cluster Connection 1	192.168.1.12
1	Public Cluster Connection 2	192.168.2.12
2	Private Cluster Connection	10.1.1.2
2	Public Cluster Connection 1	192.168.1.13
2	Public Cluster Connection 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 second node. You can do this from any machine that is physically connected to each node's networks.

In case you need to view the addresses, retrieve them using the **ipconfig** command on each node:

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

To check or change the Network Binding order

1. Right-click **My Network Places** and then select **Properties**
2. Select **Advanced, Advanced Settings**.

The top-most connection is first in the binding order. Typically, this will be the most used network adapter. You can reorder the position of the adapters by selecting them, then pressing the **up** or **down** arrow keys, then pressing **OK**.

Shared storage setup (step 2 of 4)

The following steps need to be followed if your shared storage is not configured, and only for the first node. If your shared storage is configured, the only steps you need to do are to create the Quorum disk and ensure that the disks are labeled the same on each node.

Configure the fibre-attached storage: You can configure the fibre-attached storage with the following steps or contact your disk administrator to have them done.

1. Double-click the **IBM NAS Admin** icon, then click **Storage**, then **Disk Management**.
2. Create a logical drive for the Quorum disk (RAID 5 is recommended for performance and redundancy). The drive should be at least 500 MB in size, but no larger than 1 GB. Also, the Quorum disk requires its own LUN.
3. Do not create any arrays that are RAID 0, as this is not supported.
4. If needed, see "Checking the FastT host-bus adapter settings and fibre-channel connectivity" on page 16.
5. Rescan the disks. The **Rescan** function is located under the **Action** menu.

Note: You cannot dynamically add storage to a defined volume when that volume is in a cluster.

Format the logical drives: Follow this procedure to format logical drives.

Note: 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.

1. After rebooting, open **IBM NAS Admin** and select **Disk Management (Local)**, in the **Storage** folder.
2. You will see a Write Signature and Upgrade Disk Wizard pop up on the screen. Click **Cancel**.
3. Right-click on the bottom right, where it says **Disk 1**, and select **Write Signature**.

4. Select **Write Signature** for all disks that will be accessed by the NOS (All disks in view).

On the Quorum drive, right-click and select **Format**. Enter the following parameters:

- File System of *NTFS*
- Volume Label of *Quorum Disk*
- Drive Letter of *Q* (suggested)

Do **not** enable disk compression, and select **Finish**.

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. If you do not plan to use the NAS Backup Assistant for PSM Persistent Image backups, assign a drive letter of *F* for the first drive, *G* for the second drive, and so on.

Note: If you plan to use the NAS Backup Assistant to perform PSM Persistent Image backups (see Chapter 4 of the *User's Reference* for more information about the NAS Backup Assistant), you should reserve drive letter *F* for use by the NAS Backup Assistant. Do this because the NAS Backup Assistant, in performing a persistent image backup, will map the persistent image created for the backup to the next available drive letter in the system. Thus if drive letter *F* is assigned to a drive that is owned by the other node at the time that the persistent image for backup is mapped, and the drive fails over to the node on which the backup is being performed during the backup, the backup will fail. If you reserve drive letter *F* for the NAS Backup Assistant, then you would assign drive letter *G* for the first drive, *H* for the second drive, and so on.

At this point, you have set up shared storage (if this is the first node). You can now close IBM NAS Admin and continue with "Joining a domain (step 3 of 4)".

Joining a domain (step 3 of 4)

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.

Documentation suggests that you can configure each node in the cluster as a domain controller. However, in some instances, configuring each node as a domain controller might make them unusable. This could happen if a failure occurred on a PDC and the Backup Domain Controller (BDC) took over. Then, if the PDC were to come back up before the BDC can finish Active Directory updating, both systems would be in a state of conflicting controlling memberships.

To join a node to a domain:

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 do these steps to join a domain:
 - a. Select the **Domain** radio button
 - b. Type in the name of your domain and click **OK**.
 - c. When prompted, enter the Administrator user id and password and press **OK**.
4. Close the **System Properties**.
5. Restart the computer.

After the computer restarts, it is recommended that you do not login to the domain as this will change your desktop.

Cluster configuration (step 4 of 4)

At this point, you have completed the preparation steps for cluster installation and are ready to configure the cluster. Do the following steps:

1. Start the Cluster Configuration Wizard: open **IBM NAS Admin**, then the **Cluster Tools** folder, and click the **Cluster Setup** icon.
2. You will be asked to verify that you have completed the three steps that precede this cluster configuration step. If you have, click the **Continue** button.
3. Depending on whether you are configuring the first node or the second (joining) node, click the **First Node** or **Joining Node** button.
4. If you are configuring the first node, do these steps, otherwise go to step 5.
 - a. The Cluster Information panel appears. Enter the data for the following fields (some of this data comes from Table 1 on page 6):
 - Administrator's ID
 - Password
 - Domain name
 - Cluster name
 - Cluster IP address
 - Subnet mask
 - Quorum drive (select from the pulldown)
 - b. After you have entered the data, click **Continue**.
 - c. You will be asked to verify the information. If it is correct, click **Yes** to start the configuration. You will see a message indicating that configuration takes a few minutes.
 - d. If you are asked by the Cluster Service Configuration Wizard to select a user account, enter the user name and password for the domain account you want the cluster service to use. If you are not asked, cluster configuration completes.
 - e. If you are asked to select a disk on which to store cluster checkpoint and log files, do the following:
 - 1) Select the disk on which the Quorum is located (for instance, Q, if this is what you specified earlier) and click **Next**.
 - 2) The Cluster Information panel displays. Click **Finish**.If you are not asked, cluster configuration completes.
 - f. Cluster configuration completes for the first node, the Cluster Administrator opens, and you can now go through the steps for the joining node, starting with "Networking setup (step 1 of 4)" on page 8.
5. If you are configuring the joining node:
 - a. Click the **Joining Node** button.
 - b. The First Node Information panel appears. Enter the name of the **first** node.
 - c. You may be asked to confirm the Administrator name and password. Enter that information and click **Finish**.
 - d. You will see a message indicating that configuration takes a few minutes. When configuration completes, the Cluster Administration function starts.

Cluster configuration

This section contains procedures to assist you in configuring basic cluster functions. It is assumed that the cluster installation has completed without errors, and both cluster nodes are running.

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

Initial configuration

The following steps are needed to reset the size of the logfile and set the priority and purpose of the private network.

1. To start the initial configuration of the clustering function, 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. Underneath it are the Groups (there is one disk group for each disk that you created on the shared storage), Resources, Cluster Configuration folders, and the two nodes in the cluster. Under the nodes are the respective active groups, active resources, and network interfaces. These folders allow you to select cluster resources on either node to view or change their properties.
3. Click on the cluster name to see the status of the cluster nodes in the right hand pane. The state of both nodes should be "Up."
4. Right click on the cluster name and select **Properties**. The Properties window has four tabs:
 - **General**
 - **Quorum Disk**
 - **Network Priority**
 - **Security**
5. Select **Quorum Disk**, and change the *Reset quorum log at:* field from 64 KB to 4096 KB.
6. Select **Network Priority** to view all networks acknowledged by the cluster server, then select the private network connection and move it to the top for cluster communication priority by clicking the **Move Up** button.
Doing this provides internal communication to the private network before attempts are made to communicate over any public networks that are installed.
7. 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.
8. Click **Apply**, then **OK**.

Load balancing

When you configure cluster resources, you should set up load balancing 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 load balancing:

1. Select a disk group and bring up its Properties panel by right-clicking it.
2. Click on 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 **arrow** button to move the node to the Preferred Owners pane.

Each disk group should have a preferred owner so that, when both nodes are running, all resources contained within each disk group will 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.

If you restart a cluster node, those resources that are preferentially owned by this restarted node failback to the node once the cluster service detects that the node is operational. For the failback to happen, you need to have defined the appropriate failover policy. If you have not defined the restarted node as the preferred owner for the resources, then they do not failback to the node.

Failover setup

The failover of resources under a disk group on a node allows 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 fail, then the whole group will failover 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 your Threshold for Disk Group Failure.
For example, if a network name fails, clustering services attempts to failover the group 10 times within 6 hours, but if the resource fails an eleventh time, the resource will remain 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 failback 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 usage.

User creation

The creation of users is performed through normal procedures. Users do not need to be created exclusively for use on the cluster resources. You will need to 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.

File share creation

The creation of file shares involves dependencies on a physical disk, a static IP address, a network name, and a file share name. The various 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.

Physical disk

Is the base resource in which to store user data. It is not dependent on any other resources except for 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 will not change if the resource is moved to the other node.

Static IP address

Is a virtual address that will bind 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, rather 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 through the Cluster Administrator on the same physical network.

Note: Note that the cluster IP Address is not to 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

Provides an alternate computer name for an existing named computer. It is physically dependent on an IP address on 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 to a specific node and can be failed over to another node in the cluster. Users access the groups using this virtual server.

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

An example of how to create a clustered file share follows. Assume you are creating a file share in Disk Group 2:

1. Select **Disk Group 2**, right-click it, then select **New**, and then **Resource** to bring up the New Resource panel. Here you name and describe the resource, and select its resource type. You can also specify a different group other than the default group that appears when the New Resource panel appears. After the main New Resource panel, you can define possible owners and dependencies on succeeding panels. On the dependencies panel, you define the resources that must be brought online by the cluster service before the new resource can be brought online.
Using the New Resource panel, you can then create the following resources.
2. Create an IP address resource using your first public network, and give it an address of 192.168.1.200 and a name of ipaddress1. Your first node's public network is 192.168.1.101 and your second node's is 192.168.1.102. Bring the resource online.
3. Create a network name resource, NN1, and make it dependent on ipaddress1 and bring the resource online.
4. Create a file share resource, fs1, dependent on Drive G: and NN1. You then set up the location, comment, maximum number of users, and share permissions. Bring the resource online.

A Clustered NFS Share is created using a separate resource type under the Cluster Administrator.

Checking the FastT host-bus adapter settings and fibre-channel connectivity

FastT Check is provided for diagnostic purposes. You can use it to determine if your physical fibre channel connections are in place.

1. Go to **IBM NAS Admin**, expand **Storage** and select **NAS Utilities**. Then click on **FastT Check** on the left hand side.
2. Connect using the *hostname* of the system you are using with Terminal Services.
3. Once connected, select the **QLA2200 Adapter**, and verify that you see all Fibre Controllers that you are physically connected to.

If you see a red "X" on the QLA2200 Adapter icon, check the fibre connection. (A green icon means connections are in place.) If, after checking the connection, the QLA2200 adapter still does not connect, go through the hardware diagnostics.

4. Click the **Utilities** button on the bottom right-hand side.
5. Change from **Link Status** to **NVRAM Settings**.
6. Under **Advanced Adapter Settings** (about half-way down), look at **Luns per Target**. It should be set to 0. If not, change the value to 0. Do this for any other installed adapters.



Part Number: 38P7712



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

(1P) P/N: 38P7712

