Compaq StorageWorks

Command Console V2.1 for RAID Array 8000/ESA12000

User's Guide

AA–RFA2C-TE   387405-003

Compaq Computer Corporation

Warning!
This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Achtung!
Dieses ist ein Gerät der Funkstörgrenzwertklasse A. In Wohnbereichen können bei Betrieb dieses Gerätes Rundfunkstörungen auftreten, in welchen Fällen der Benutzer für entsprechende Gegenmaßnahmen verantwortlich ist.

Attention!
Ceci est un produit de Classe A. Dans un environnement domestique, ce produit risque de créer des interférences radioélectriques, il appartiendra alors à l'utilisateur de prendre les mesures spécifiques appropriées.

JAPAN

この装置は，情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

USA

This equipment generates, uses, and may emit radio frequency energy. The equipment has been type tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules, which are designed to provide reasonable protection against such radio frequency interference. Operation of this equipment in a residential area may cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference. Any modifications to this device - unless expressly approved by the manufacturer - can void the user's authority to operate this equipment under part 15 of the FCC rules.

Contents

3   Installing the StorageWorks Command Console for the Fibre Channel Switch

4   Command Console's Windows, Menus, and Tools

5   Maintenance and Troubleshooting

Figures (Continued)

Figures (Continued)

Figures (Continued)

Tables

Revision Record

This Revision Record provides a concise publication history of this guide. It lists the guide revision levels, release dates, and reasons for the revisions. It also describes how the changes to affected pages are marked in the guide.

The following revision history lists all revisions of this publication and their effective dates. The publication part number is included in the Revision Level column, with the last entry denoting the latest revision. This publication supports Command Console V2.1a for the RAID Array 8000/12000 storage subsystem, a Graphical User Interface (GUI) for the management of storage subsystems.

Revision Level	Date	Summary of Changes
AA–RFA2A–TE (387405-001)	November 1998	Original Release.
AA–RFA2B–TE (387405-002)	December 1998	Revised Chapter 2.
AA-RFA2C-TE (387405-003)	April 1999	Supports Version 2.1 with Fibre Channel Switch.

About This Guide

This User's Guide provides instructions on how to run and use Command Console V2.1 for the HSG80 Controller, a Graphical User Interface for the management of Compaq's RAID Array 8000/ESA12000 controller.

Audience

You should be familiar with 32-bit Windows operating systems (Windows 95 or Windows NT), RAID terminology and, if you are using Command Console over a network, TCP/IP networking fundamentals.

Document Organization

This guide contains the following:

Chapter 1: Introduction

The Introduction defines and provides an overview of Command Console.

Chapter 2: Running and Using Command Console

Chapter 2 describes how to choose a storage subsystem connection method and how to launch and run Command Console.

Chapter 3: Installing the StorageWorks Command Console for the Fibre Channel Switch

This chapter contains instructions for installing StorageWorks Fibre Channel Switch software (Client and Agent) required to support operation of the Fibre Channel switch in a Windows NT environment.

Chapter 4: Command Console's Windows, Menus and Tools

This covers Command Console's Control Panel. Menu Bar, Toolbar, and graphical elements.

Chapter 5: Maintenance and Troubleshooting

Chapter 5 provides data and configuration file backup recommendations, firmware upgrading procedures and troubleshooting.

Appendix A: Creating Storagesets

This appendix provides supplemental guidelines for determining and planning storage requirements.

Appendix B: Fibre Channel Switch Installation and Basic Set-up

This Appendix is designed to give you basic instructions for installing the Fibre Channel switch and for making the necessary connections.

Glossary

A glossary is provided which defines the terms related to RAID storage products.

Associated Documents

RA8000/ESA12000 HSG80 Solution Software V8.3 Installation Reference Manual for your specific host platform.

Quick Start Guide - RA8000/ESA12000 HSG80 Solution Software V8.3 for your specific host platform.

Quick Setup Guide StorageWorks Fibre Channel Switch, AA-RHCOA-TE

Fibre Channel Switch User's Guide, AA-RHBYA-TE

In addition to these manuals you should refer to all documentation included in your Software Solutions Kit.

Conventions

This guide uses the following documentation conventions:

Table 2 Style Conventions

Style	Meaning
boldface type	For emphasis and text-based user input
italic type	For emphasis, manual titles, Windows menus, screens, and filenames
`plain monospace type`	DOS, UNIX and text-based applications

NOTES, CAUTIONS, and WARNINGS

NOTE

A note calls attention to information that can be helpful in understanding the operation of the product.

CAUTION

Caution calls attention to an operating procedure or practice that could result in damage to the product if not correctly performed. Do not proceed beyond this symbol until you fully understand and meet the indicated conditions.

WARNING

Warning calls attention to a procedure or practice that could result in personnel injury if not correctly performed. Do not proceed beyond this symbol until you fully understand and meet the indicated conditions.

Getting Help

If you have a problem and have exhausted the information in this guide, you can get further information and other help in the following locations.

Compaq Web Site

The Compaq Web Site has information on this product as well as the latest drivers and Flash ROM images. To access the Compaq Web Site, log on to the Internet at:

http://www.compaq.com/products/storageworks

Telephone Numbers

For the name of your nearest Compaq Authorized Reseller:
In the United States, call 1-800-345-1518
In Canada, call 1-800-263-5868
For Compaq technical support:
In the United States and Canada, call 1-800-386-2172

1

Introduction

This chapter defines Command Console, describes Command Console's features and provides network and device support information.

1.1   What is Command Console?

Command Console is a graphical user interface (GUI) for managing StorageWorks RAID array products from a console running on Microsoft Windows NT® or Windows 95®. With Command Console you can configure virtual disks, receive notification on events, and monitor your storage systems.

Command Console provides storage array management over a TCP/IP network, serial port, or SCSI bus. The three connection methods are shown in Figures 1 –1, 1–2, and 1–3.

In a TCP/IP network environment, Command Console works by means of the client-agent process. For the purpose of this book, the network connection shown in Figure 1 –1 uses a Fibre Channel switch.

StorageWorks Command Console can consist of any number of Clients and Agents.

- The **Client** is the graphical user interface program designed for use on Windows NT® (Intel or Alpha), or Windows 95® operating systems. It provides an easy method of configuring, operating, monitoring, and troubleshooting your storage subsystem.

- An **Agent** is a companion program that is installed on one of the supported operating systems to which a controller is connected. The Agent lets the Client communicate with your storage system over a network.

Figure 1–1 Network Connection
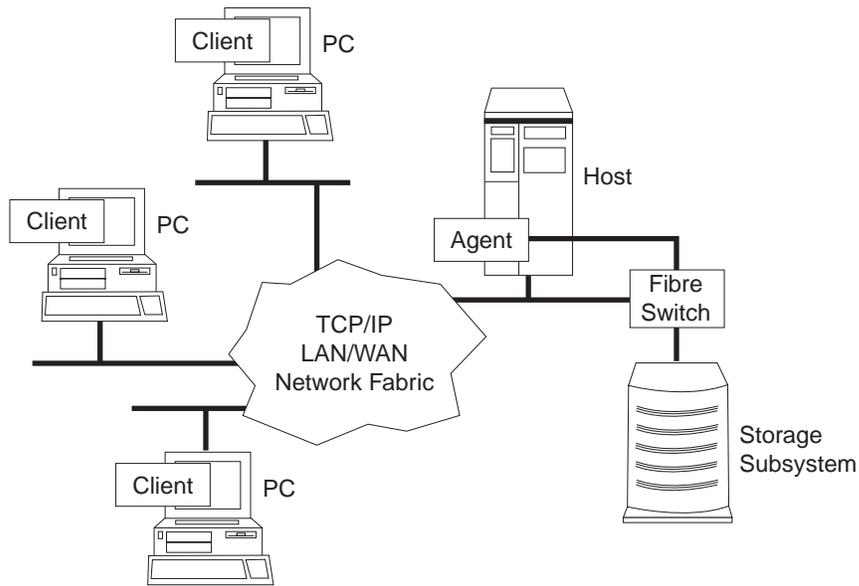
SHR-1250

Figure 1–2 Local Serial Port Connection
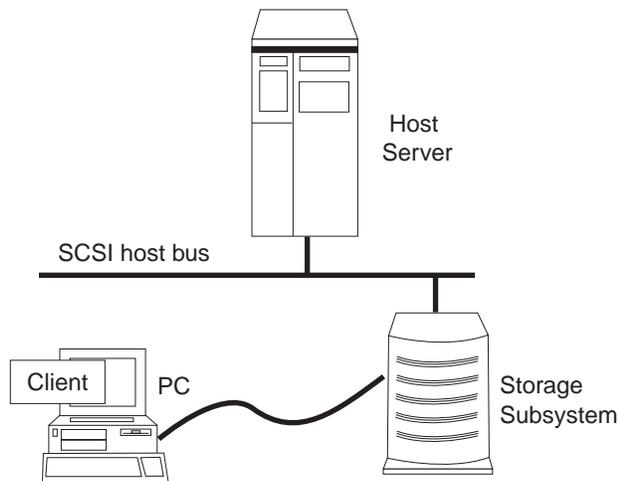
shr-1184a

Figure 1–3 SCSI Bus Connection
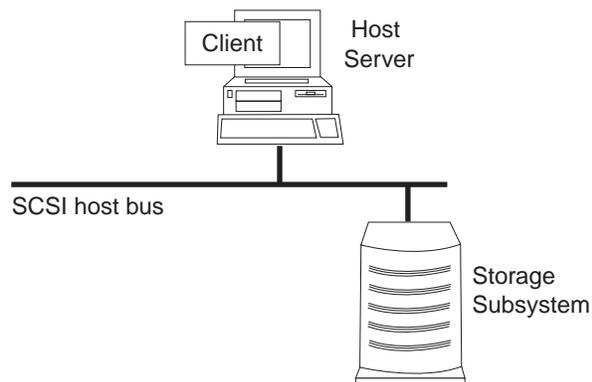
shr-1184b

1.2   Command Console Features

- Manages StorageWorks RAID Arrays over a serial port, SCSI bus, or TCP/IP network.
- Configures RAID arrays using menus, icons, and system views.
- Features an explorer-like navigation window and right click on mouse for properties.
- Configures arrays from files. Array configurations can be saved as files and used as backup or to automatically configure other arrays.
- Upgrades subsystem software.
- Monitors storage subsystems. Subsystem monitoring includes power, cooling, temperature, device, controller, and cache monitoring.
- Protects your data and configuration with password protection option.
- Interfaces with Windows NT's logging capability. Lists RAID array events in the Windows NT's event log.
- Notifies you of subsystem problems by pager.

1.3   Online Help

Command Console's online help is complete and fully indexed, providing instructions and information on all of Command Console functions and connection options. Online help explains RAID levels, has an excellent glossary, and covers the following major topics:

- An overview of RAID Array 8000/ESA 12000 Command Console applet manager
- Configuring Your Controller
- Configuring Your Storage Subsystem
- Maintenance

- How to…
- References

Command Console Help uses standard, WinHelp conventions along with the following:

Abbreviations---The term "Client" is used to refer to the Command Console Client program. The term "Agent" is used to refer to the Command Console Agent program.

F1---Displays context-sensitive help.

1.4 Command Console Client Hardware and Software Requirements

The minimum requirements to properly run the Command Console client are listed in Table 1–1.

Table 1–1 Minimum Hardware and Software Requirements

Item	Requirement
Management Station	Minimum of an Intel 486/66 Mhz with 16MB of RAM memory and 10 MB free hard disk space
Operating System	Windows NT 3.51, 4.0, or Windows 95
Monitor	VGA at 800x600 resolution, minimum
Modem	Hayes-compatible (for event notification)
Serial Port	RS-232 (for local connection)
Network Adapter	TCP/IP-compatible network card (for distributed systems only)
Controller Compatibility	The RAID Array 8000/ESA 12000 Controller HSG80
Host software	ASPI SCSI driver required for host port virtual terminal connection (NT only)

1.5   Operating System and Controller Support

Command Console Client runs on Windows 95 or NT PCs with Intel or Alpha CPUs running Windows NT. For specific version number information consult the *HSG80 Release Notes* for your host operating system.

1.6   Network Specifications

Command Console is a terabyte-class RAID Array manager. Fully scaleable, Command Console can manage a single, locally connected RAID array or RAID arrays in a multi-node, TCP/IP network environment. Command Console can support any number of clients and agents in a distributed environment.

1.7   Planning Your Distributed Storage System

Plan your distributed system. Locate the host platforms and storage subsystems. Determine the RAID level needed at each location. Selecting a RAID level is based on your performance, cost, and availability goals. Refer to *Appendix A, Creating Storagesets*, of this manual or online help for further information and guidelines for RAIDset selection.

2

Running and Using Command Console Client

This Chapter covers how to install, launch, and use Command Console Client.

2.1   Installing StorageWorks Command Console Client

Before you install a new version of Client, remove the old version by selecting *Settings* from the Start menu; then *Control Panel*. Double-click on the *Add/Remove Programs* icon. Select the *Command Console V2.1* applet, click the *Add/Remove* button and follow the prompts. Do the same for the other installed Command Console applets.

Figure 2–1 Add/Remove Programs Properties Sheet

Install New Version of Command Console Client

Command Console Client installs from a CD -ROM disk using a standard Windows installation routine on a Windows 95 or Windows NT platform. The program is self-extracting and stores Command Console Client into the directory `C:\Program Files\SWCC` by default. During setup, you have the option to change the disk or directory location.

To install Command Console:

The Command Console Client software is installed on your host system from the StorageWorks Command Console CD-ROM. After you insert the CD-ROM, the program should start automatically.

NOTE

If the program does not start automatically after inserting the CD-ROM, go to the CD-ROM drive directory and click on install.bat).

The installation program provides two options (reference Figure 2 –2):

- Solution Software for ACS v8.4 with FC Switch

- Solution Software for ACS v8.3 with FC Hub

Click on the Solution Software option appropriate to your ACS version.

Figure 2–2 Solution Software (ACS Version) Selection Menu
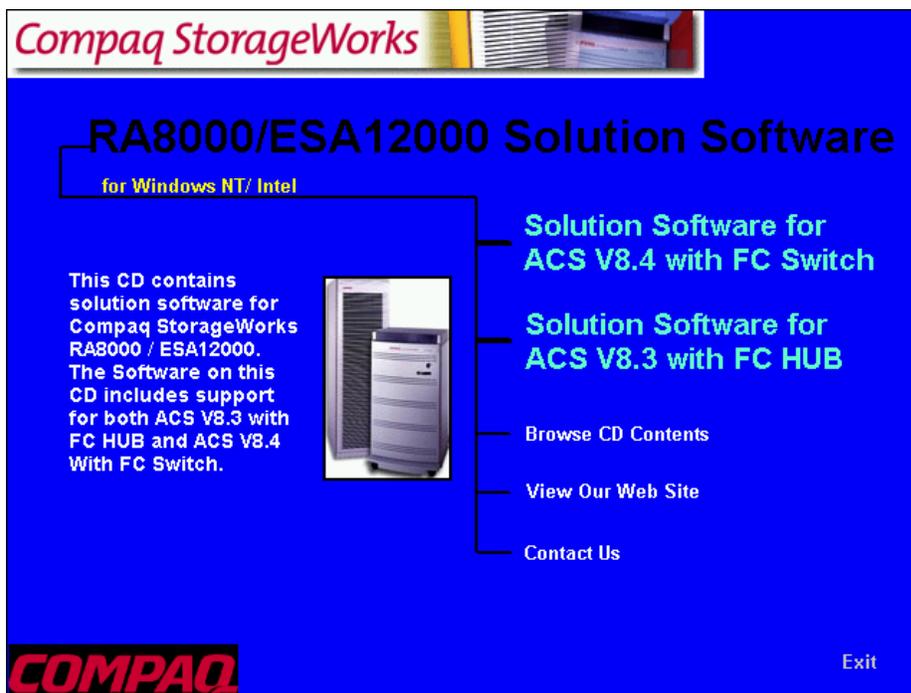
From the Fibre Channel Solution Software menu that appears, select *Install StorageWorks Command Console* option. The StorageWorks Command Console software will be automatically installed.

Follow the instructions in the setup program to complete the installation.

Command Console Client installs the Program Group *Command Console* and places server icons within the group. Client also inserts server selections on the Start menu (Figure 2–3).

Figure 2–3 Command Console Client's Start Menu

The three choices Client provides you with are:

- **HSG80 Storage Window:** Displays the Storage Window. This choice lets you monitor and configure one storage subsystem using Client's graphical user interface. All connection choices are provided: serial line, SCSI bus and network (TCP/IP). Refer to *Command Console V2.1(HSG80) RAID Array 8000/ESA12000 User's Guide* for complete information.

- **CLI Window:** Displays the CLI Window. The CLI window lets you monitor and configure storage subsystems using text commands. All connection choices are provided: serial line, SCSI bus and network (TCP/IP).

- **StorageWorks Command Console:** Displays the Navigation Window. The Navigation Window is a network navigation tool used to manage and monitor storage subsystems over a TCP/IP network. This choice lets you monitor and configure one or many storage subsystems over a network using Client's graphical tools.

2.2   Starting Command Console Client

Command Console Client provides you with three applet choices: *HSG80 Storage Window, CLI Window,* and *StorageWorks Command Console*. Choosing *HSG80 Storage Window* applet allows you to connect to only one storage subsystem. The other applet choice, StorageWorks Command Console, is an applet manager which lets you connect to, monitor and configure many storage subsystems over a TCP/IP network. To start one of the applets, single‑click on the applet's icon in the *Command Console* program group or, from the Start menu, select the applet's name.

If you select the *HSG80 Storage Window* Client displays the *Connection Selection* dialog box (Figure 2–4). This selection lets you connect locally to your storage subsystem over a serial line or the SCSI bus or, if you choose this applet, you can connect to one subsystem over a TCP/IP network.

Figure 2–4 Connection Selection for Connecting to a Single Storage Subsystem
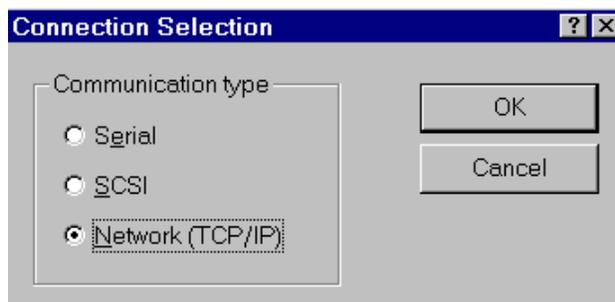
If you select the StorageWorks Command Console, the first window you will see is the
Navigation Window (Figure 2–5). The Navigation Window is an applet manager and is
used to monitor and configure many storage subsystem over a TCP/IP network.

Figure 2–5 Navigation Window
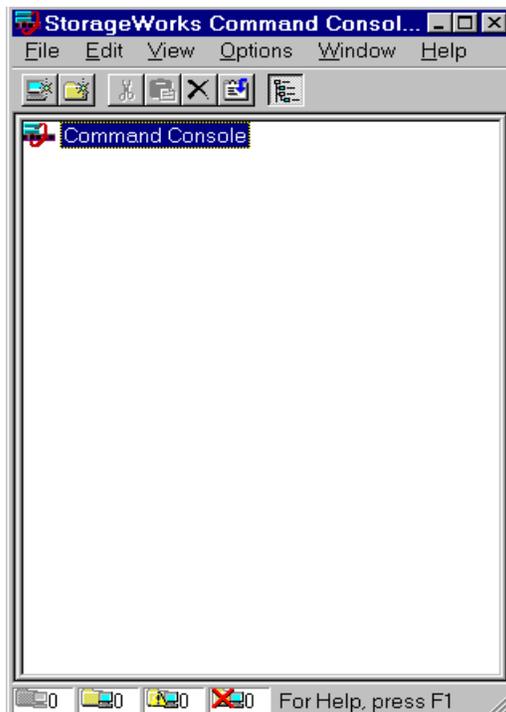
2.3   Choosing a Connection Method

Client offers three ways to connect to your storage subsystem: over the controller's serial
port, the SCSI bus or over a TCP/IP network. The serial port and SCSI bus connections
are local connections and allow you to connect to only one storage subsystem. You can
connect to one or many storage subsystems using a TCP/IP network connection.

You cannot make a network or SCSI bus connection to a storage subsystem until at least
one virtual disk has been created on it or the Command Console LUN (CCL) is enabled.
If your subsystem is not configured this way, you must establish a connection to the
controller by means of a serial connection and create a virtual disk or enable the CCL.
See *Establishing a Serial Connection* in this section and also 2.7, *Creating a Virtual Disk*.

Establishing a Serial Connection

The simplest connection to a storage subsystem is a direct, cable connection from the Client's host system to one of the storage subsystem's serial ports.

1. Select the *HSG80 Storage Window* applet.
2. When the connection selected dialog box displays, click the *Serial* radio button, then click *OK* to display the *HSG80 Management: Serial Connection* dialog box (Figure 2–6).
3. On the *Connect Serial* dialog box, from the drop-down menu, select the PC COM port your RAID Array 8000/ESA12000 controller is connected to. Select a *baud rate*, a *Subsystem Physical View* and a *Subsystem Grid View*; then click the *Connect* button to display the Storage Window. When the Storage Window displays, you are connected to your storage subsystem.

Figure 2–6 Serial Connection Dialog Box for Storage Window
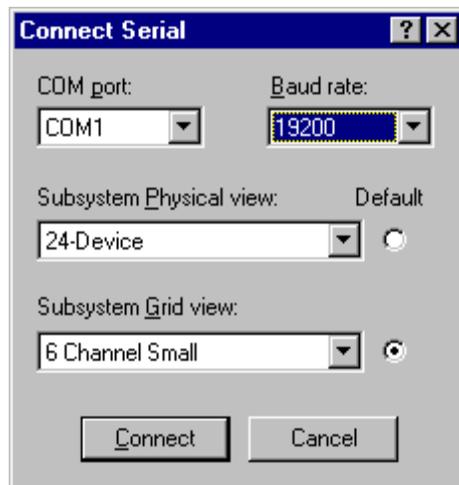
Establishing a SCSI Connection

If Client's host operating system is Windows NT, you can connect to, monitor and manage one storage subsystem over the SCSI bus.

Before you can connect to your storage subsystem over the SCSI bus, at least one virtual disk must be configured for the controller. To create the first virtual disk you must use a serial connection to the subsystem. The procedure to create the first virtual disk is summarized as follows:

- Establish a serial connection to your RAID subsystem using the *Establishing a Serial Connection* from Section 2.3*, Choosing a Connection Method*.

- Create a virtual disk. Refer to Section 2.7, *Creating a Virtual Disk*.

- Reboot your system.

- Start the Disk Administrator, create the partition and assign the drive letter for the virtual disk.

To connect over the SCSI bus proceed as follows:

1. Select *Command Console* from the Start menu, then *HSG80 StorageWindow*.
2. When the connection dialog box displays, click the *SCSI* radio button.
3. When the *HSG80 Management: Local SCSI Connection* dialog box appears (Figure 2–7), select a drive that belongs to the storag e subsystem; then click *OK*. After a brief period of time, Command Console will connect to the RAID Array 8000/ESA12000 and display the Storage Window.

Figure 2–7 Local SCSI Connection Dialog Box

Establishing a Network Connection

To communicate with your storage subsystem over a network, two software applications are required: Command Console Client and Command Console Agent. Client is the graphical user interface which installs and runs on a Windows NT or Windows 95 host system. Agent installs and runs unseen in the background on the storage subsystem's host computer. Agent provides the software interface between the host storage subsystem and any number of Client sessions, allowing you to configure and monitor your storage subsystem from many locations.

First you install Client and then Agent. Together the two provide the software link required for communication over a network. The installation of Agent is covered in the *RA8000/ESA12000 Solution Software V8.3 Installation Reference Manual* for your specific host platform, refer also to online help.

The basic steps for connecting and communicating over a network:

- Ensure that you have at least one virtual disk created on your storage subsystem or the CCL is enabled. If necessary, create one using a serial connection to the RAID Array 8000/ESA12000 or enable the command console LUN.

- Install and configure Agent on the storage subsystem's host computer.

- Select either *HSG80 Storage Window* (for a single storage subsystem) or, if you want to configure and monitor multiple storage subsystems over a network, select the *StorageWorks Command Console* applet manager. Choosing *StorageWorks Command Console* displays the Navigation Window, a graphical tool for managing RAID Array 8000/ESA12000 subsystems over a network.

Connecting To One Raid Array 8000/ESA12000 Over the Network

1. Select *HSG80 Storage Window* from the Start Menu.
2. When Client displays the *Connection Selection* (Figure 2–4) dialog box, click the TCP/IP radio button. Client displays the *Connect Network (TCP/IP)* dialog box (Figure 2–8).

Figure 2–8 Connect Network (TCP/IP)
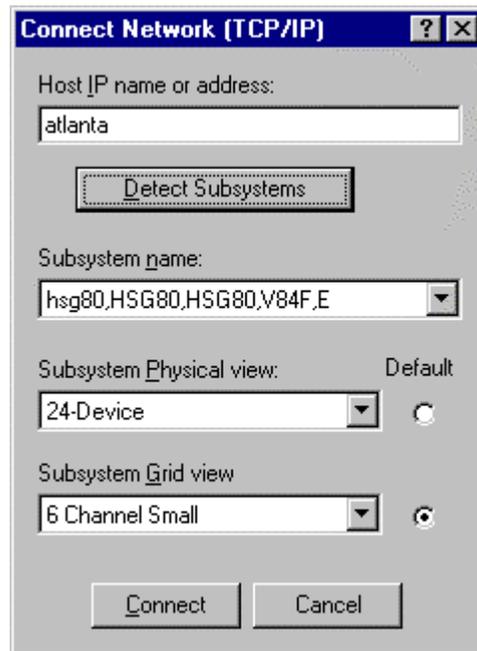
3.   Enter the host IP name *(atlanta* is used in this example) or address in the text box;
     then click *Detect Subsystems.* Client will search and find the storage subsystem
     connected to *sunday* and will display it in the *Navigation* window. Select the storage
     subsystem, *hsg80, HSG80, HSG80, V84F,E* in this example; then click *Connect.*
     After a short period of time, Client displays the Storage Window. (The name of the
     storage subsystem is defined during Agent installation).

Creating and Managing a Network of RAID Array 8000's and ESA 12000's

If you select *StorageWorks Command Console* from the Start menu, Client displays the Navigation Window, shown in Figure 2–9. The Navigation Window is a navigation tool that is used to manage and monitor many storage subsystems over a network. Refer to Section 2.4, *Using the Navigation Window* in this guide.

Figure 2–9 Navigation Window

2.4   Using the Navigation Window

The Navigation Window is a network management tool that you can use to create, monitor, and configure a network of storage subsystems. The Navigation Window has its own Menu Bar. The window can be moved and minimized. It can be sized by dragging on its corners and borders. You start building a network of RAID Array 8000/ESA12000 storage subsystems by adding systems to the Navigation Window. An example of a single-host network is shown in (Figure 2–10). The host system is named a *tlanta* and the folder *for atlanta* is shown expanded.

Figure 2–10 Navigation Window

Adding a System to the Network

1. From the File menu, click *Add System* (Figure 2–11).

2. Enter a Domain Name Service (DNS) name or the IP address in the *Host name or TCP/IP address:* text box (shown in Figure 2–12) and click *Apply*. After you click the *Apply* button, Client adds an icon for *sunday* in the Navigation Window (Figure 2–13).

3. When the second *Add System* dialog box appears, click the *Close* button.

Figure 2–11 Adding a System

Select *Add System* from the File menu.

Figure 2–12 Add System Dialogue Box
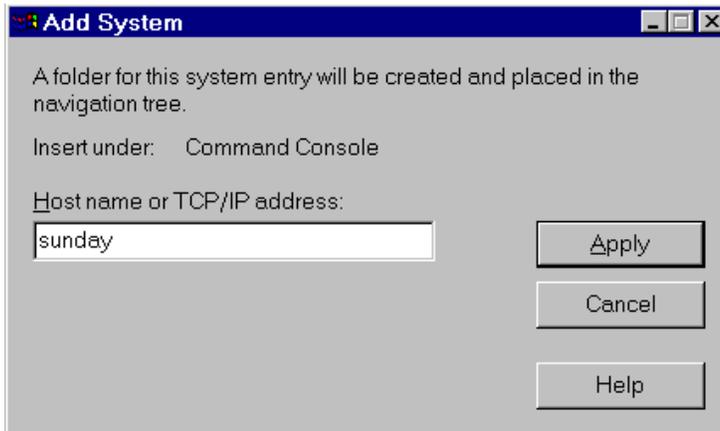
Figure 2–13 Sunday Added to the Navigation Window
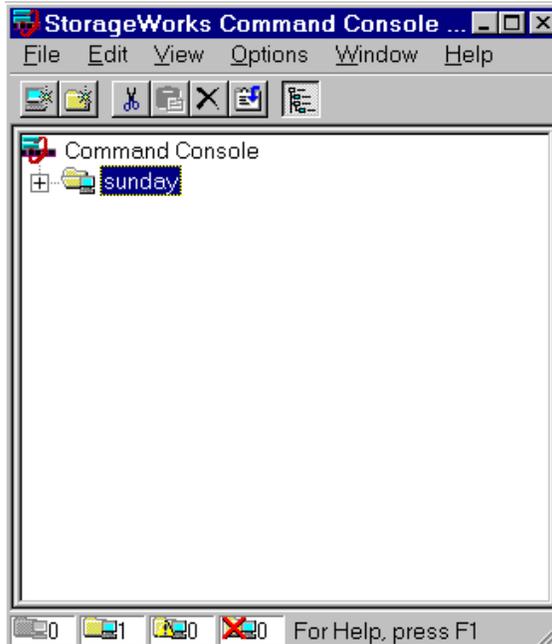
Deleting a System

1. In the Navigation Window, click on and highlight the system to you want to delete.
2. From the Edit menu, select *Delete*.

Creating and Using New General Folders

Use folders to help organize your storage. Folders can be dragged and dropped within the Navigation Window. Folders can only be placed under the Command Console root or under another general folder in the Navigation Window. Folders can be used to group systems and other general folders, but they cannot be used to group controllers. In this example we have created the general folder *My Folder* and have placed one host, *atlanta,* in the folder (Figure 2–14). To create a new folder:

1. From the File menu, choose *New Folder*.
2. Name the new folder in the Navigation Window by double-clicking its label and entering a new name.

Figure 2–14 Adding a General Folder
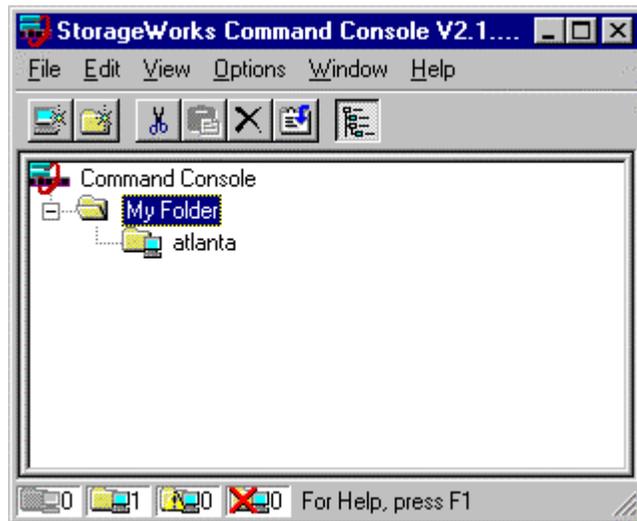
Renaming a Folder

1. Single-click on the folder name in the Navigation Window.
2. Enter the new name. Only general folders that you create can be renamed.

Viewing the Hierarchy in the Navigation Window

1. Click on a folder on the Navigation Window to display its contents.
2. Click the plus signs (+) to display more folders, systems, and controllers. To open a folder, double-click on it.

Viewing and Modifying System Folder Properties

1. Right-click on a system icon in the Navigation Window.
2. Click *Properties* from the shortcut menu to view the system folder properties.

Figure 2–15 System Properties for Sunday
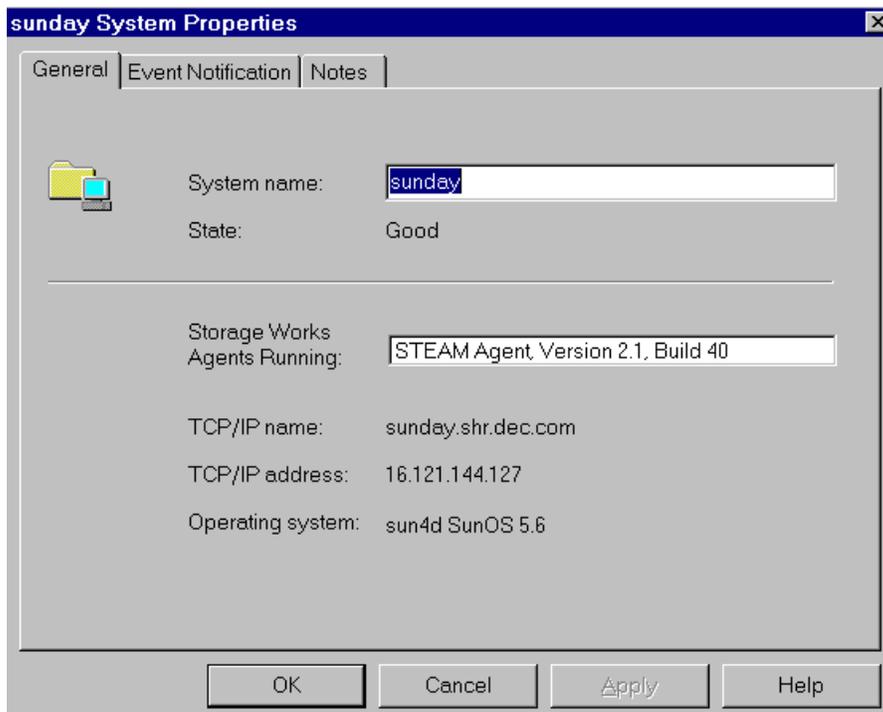
Opening a Storage Window

1. Connect to a storage subsystem.

     Storage Window

2. In the Navigation Window, double-click on a system folder.
3. Double-click the Storage Window icon to open a Storage Window (Figure 2–16).

Figure 2–16 Storage Window
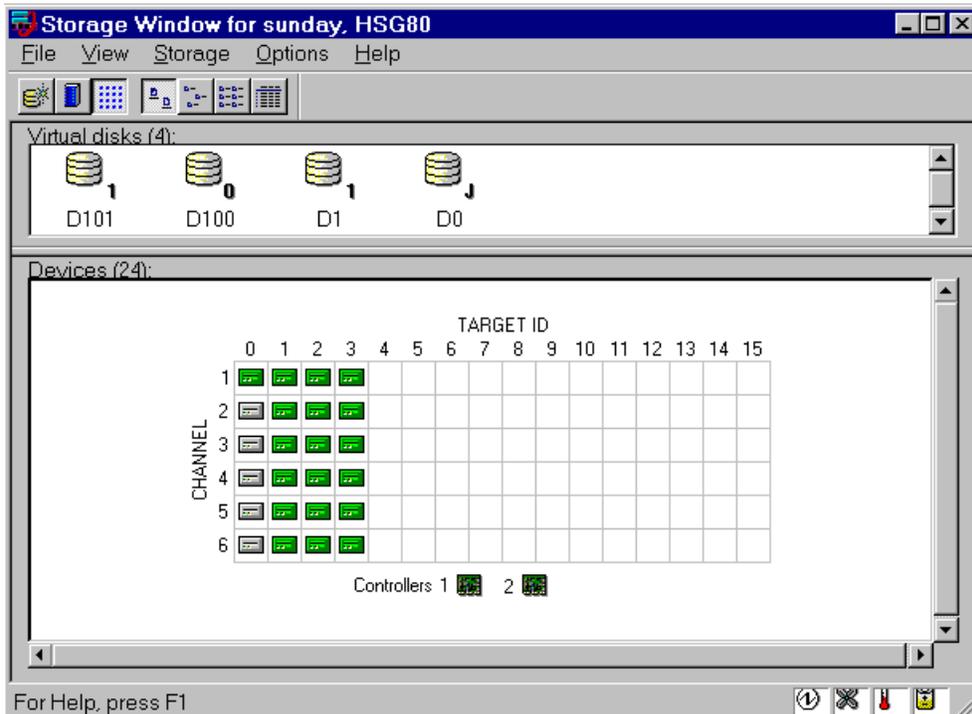
2.5 Configuring a Controller

Your controller's operating parameters are stored in property sheets. Controller property sheets are accessed by double-clicking on a controller icon in the Storage Window or right clicking on the icon and selecting *Properties*. Property sheets are tabbed. To access a sheet click on its tab.

Changes to controller settings require a controller restart to place the changes in effect. The program prompts you for confirmation before it restarts your controller.

Access the controller's property sheets by double-clicking the controller's icon in the Storage Window to display its property sheets. When you double-click on a controller's icon, the *General* controller properties sheet displays (Figure 2 –17). The controller has six property sheets. Each sheet is accessed by clicking on a tab.

Figure 2–17 Controller Icon

Controller Icon

Double-click on the controller icon in the Storage Window to configure the controller.
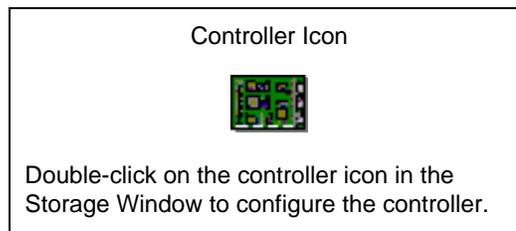
The RAID Array 8000/ESA 12000 Controller has six property sheets: General, Host Ports, Cache, Communications LUN, Connections, and Battery.

1. The *General* tab (Figure 2–18) displays the properties of the top and bottom controllers, such as type, serial number, SCSI address, firmware revision, hardware revision, and common parameters.

Figure 2–18 General Controller Properties Tab
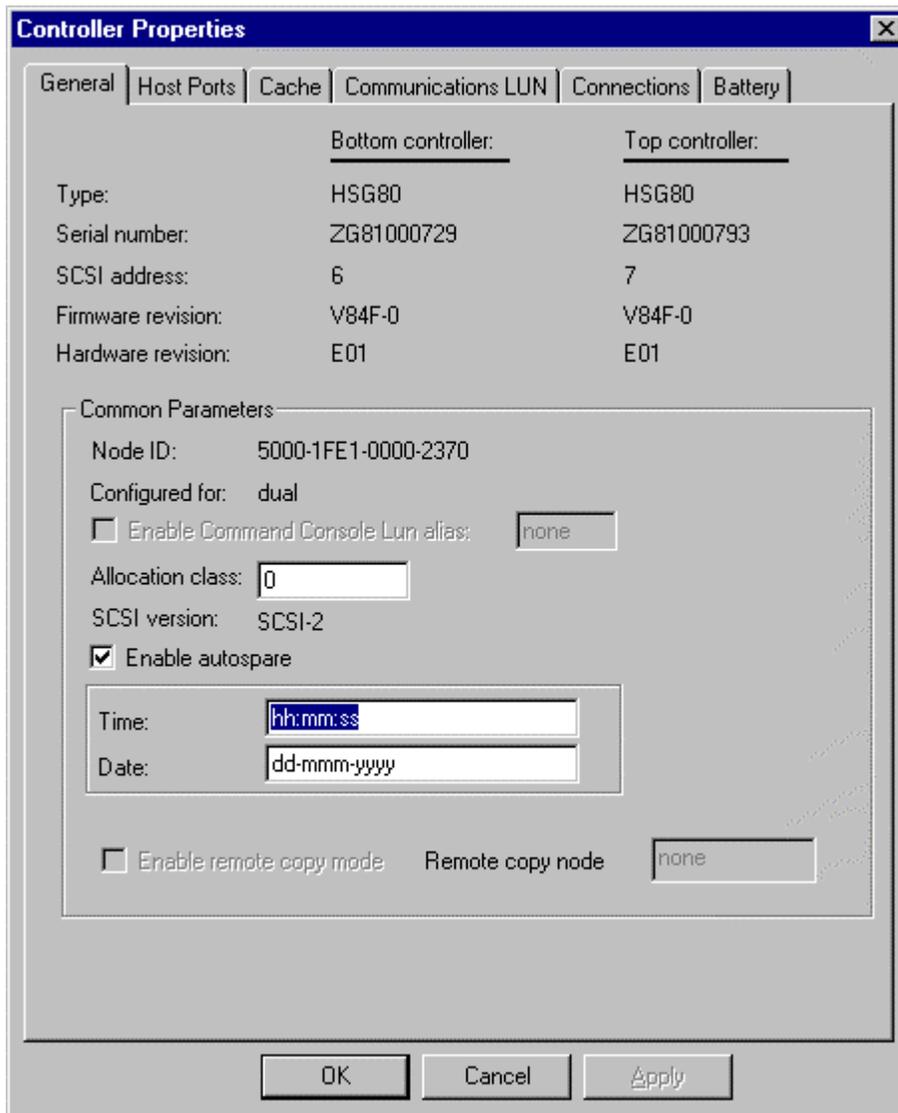
2. Click the *Host Ports* tab (Figure 2–19) to display the host port operating parameters for Host Ports 1 and 2. For platform specific settings, refer to the *Quick Start Guide - RA8000/ESA12000 HSG80 Solution Software V8.3* for your specific host platform packaged with your RAID Array 8000/ESA12000.

Figure 2–19 Host Ports Controller Properties Tab
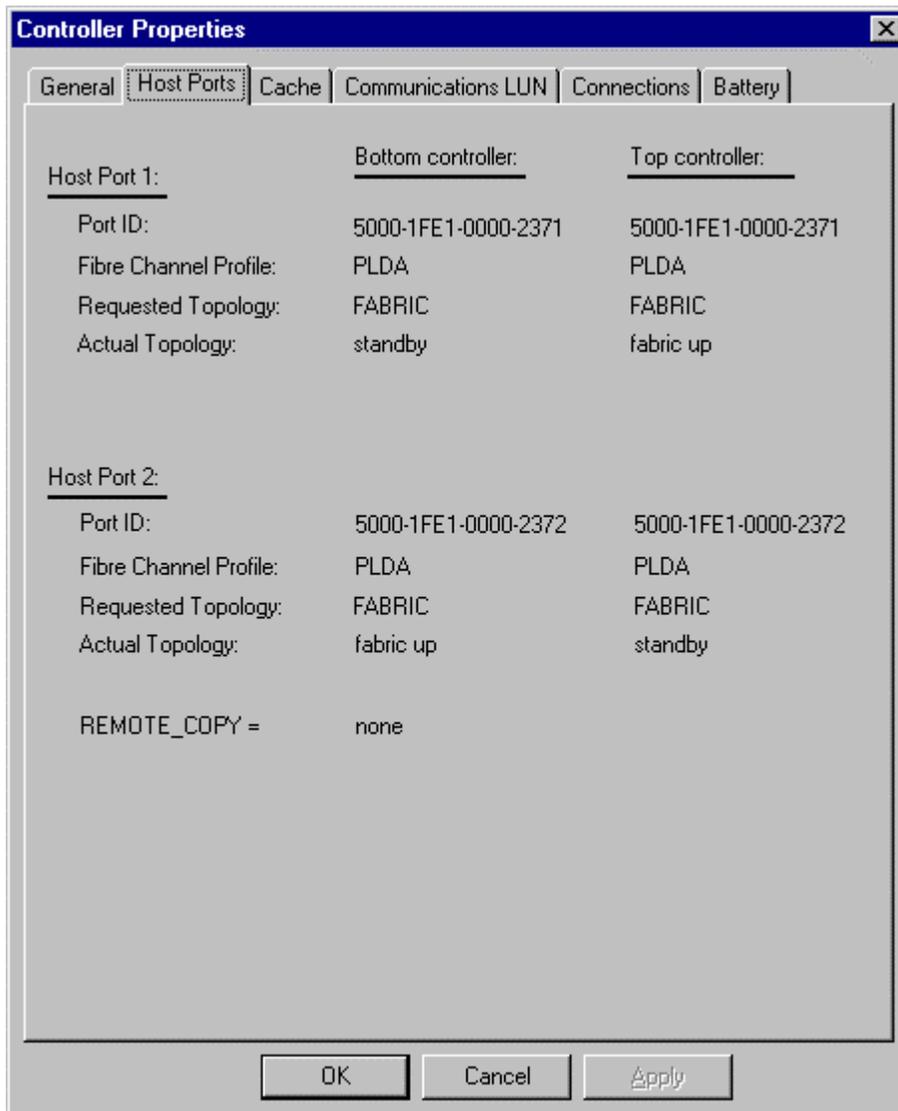
3. Click the *Cache* tab to check the cache size (Figure 2 –20). Confirm the following:

- Cache flush_time: (seconds): is 10.
- Respond to internal cache battery condition is selected.

Figure 2–20 Cache Controller Properties Tab
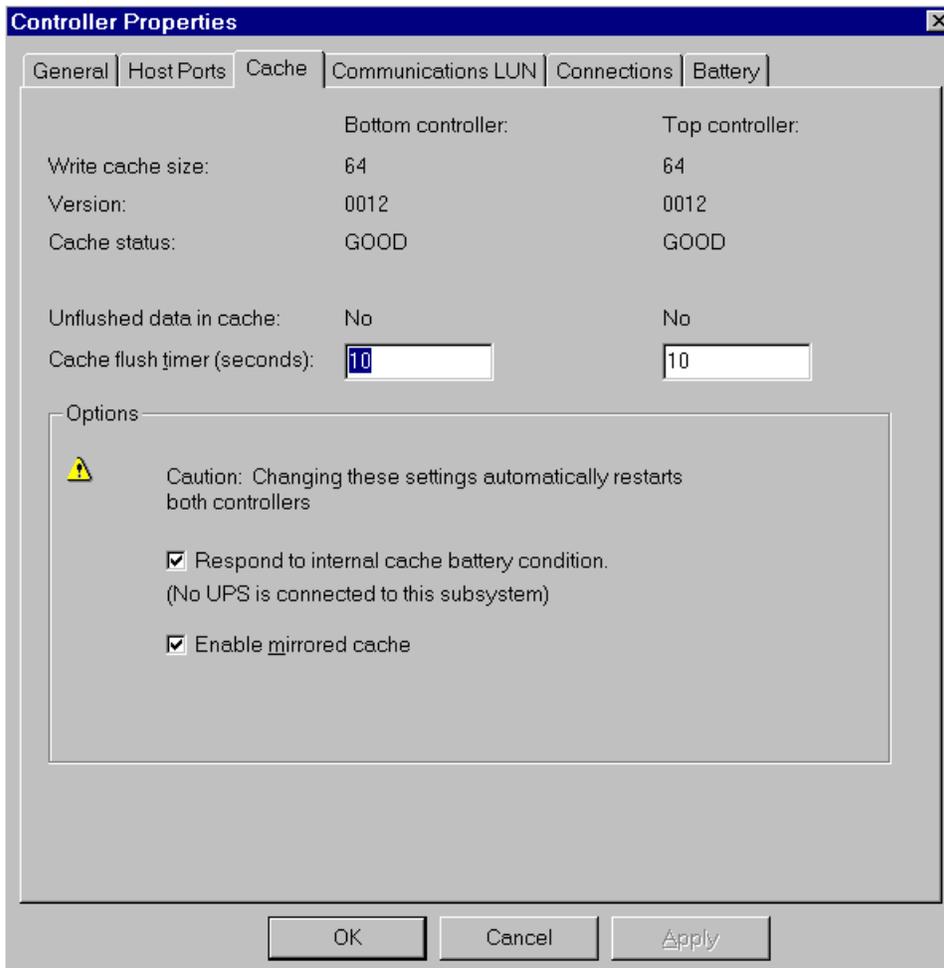
4. Click the *Communications LUN* tab (Figure 2–21) to display and set the communications LUN properties.

Figure 2–21 Communications LUN Controller Properties Tab
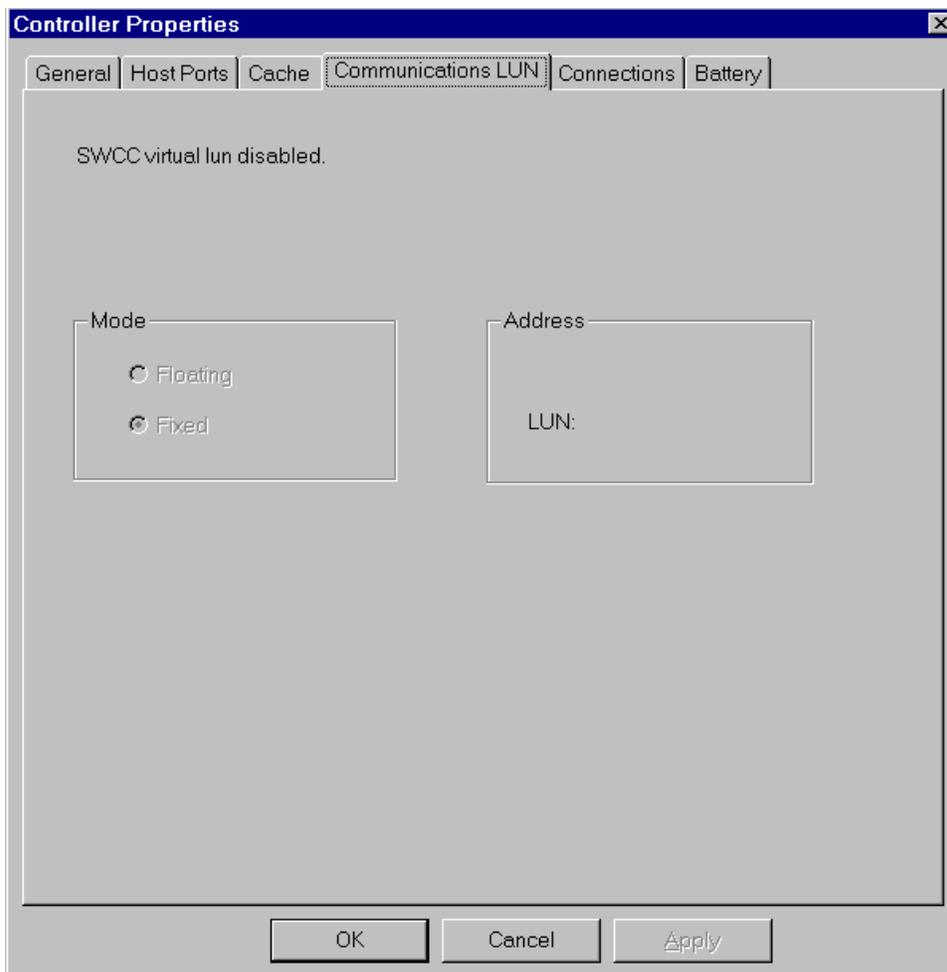
5.   Click the *Connections* tab (Figure 2–22) to display and set the connections properties.

Figure 2–22 Connections Properties Tab
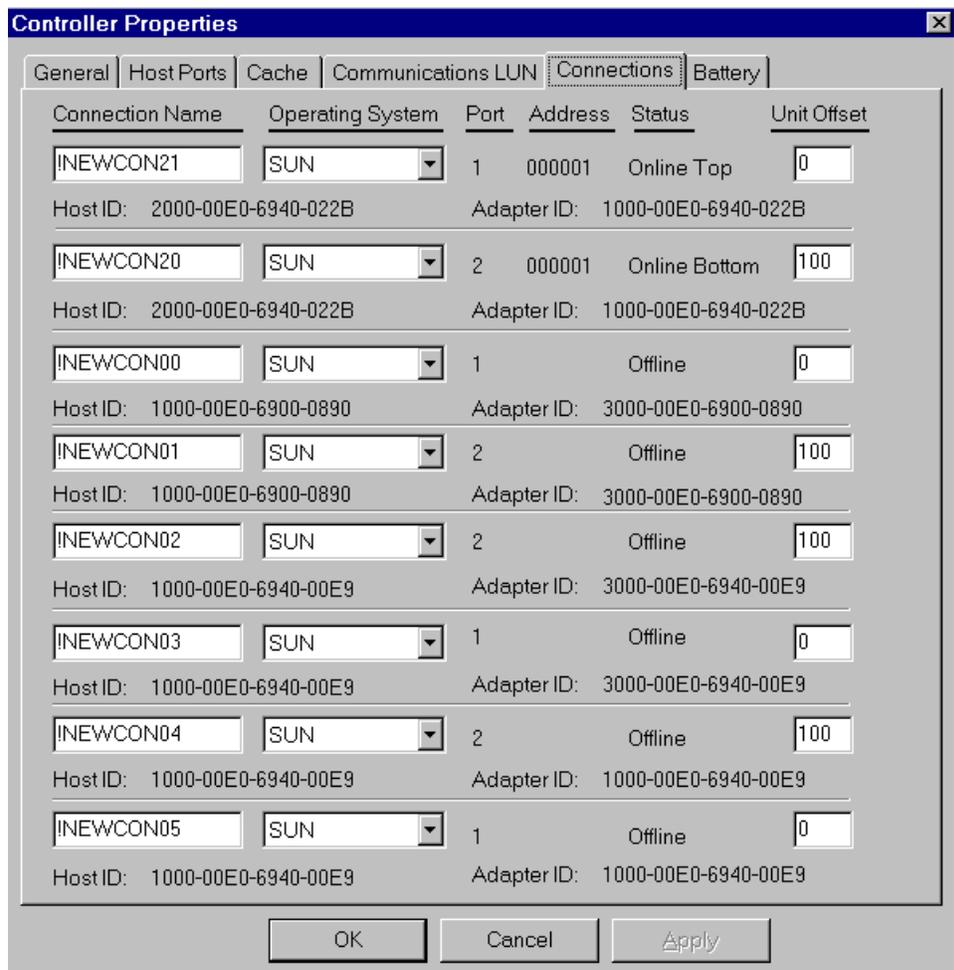
6. Click the *Battery* tab (Figure 2–23) to display the battery properties.

Figure 2–23 Battery Controller Properties Tab
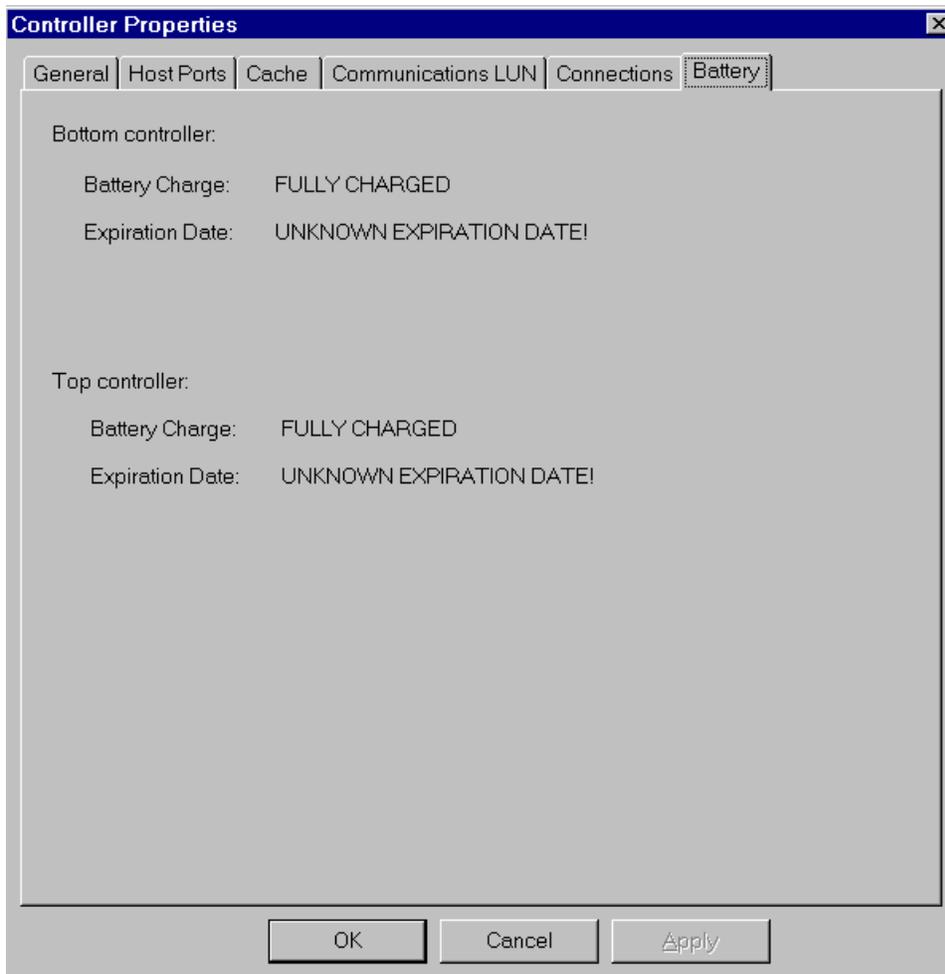
2.6 Adding Devices

You must make your physical devices known to your controller before you can create virtual disks from them. You can also add devices online without restarting the controller. To add a device to the subsystem configuration, from the Navigation window open a Storage window within a desired subsystem. Select a slot for the device and then from the *Storage* menu select *add*. Enter the SWCC password in the Security Check dialog box (Figure 2–24) that appears and click the *OK* button.

Figure 2–24 Add Devices

2.7 Creating a Virtual Disk

Command Console Client can create a number of different types of logical storage units called virtual disks on your RAID Array 8000/ESA12000. You can create:

- Single-device virtual disks (JBODs)
- Striped virtual disks (RAID 0)
- Mirrored virtual disks (RAID 1)
- Striped mirrored virtual disks (RAID 0+1)
- Striped parity virtual disks with floating parity disk (RAID 5)

Virtual disks are created using Command Console Client's Add Virtual Disk Wizard.

NOTE

Many SWCC operations require an authorization password, a security check box will appear and the user will be prompted to enter a password (Figure 2–25).

Figure 2–25 Security Check Box
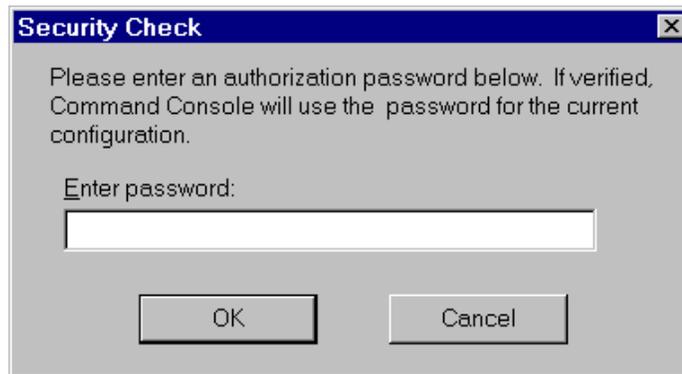
Start the Wizard:

Access the subsystem's Storage Window. Select *Add Virtual Disk* from the Storage menu
to start the Virtual Disk Wizard. To start the Virtual Disk Wizard, Access the subsystem's
Storage Window, and select *Add Virtual Disk* from the Storage menu.

1. Select the RAID Level

The first step is to select the type of virtual disk you want to create. Click the radio button of the RAID level you want (Figure 2–26), then click *Next*. In this example we will create a RAID 5 virtual disk.

Figure 2–26 Add Virtual Disk Wizard Step 1
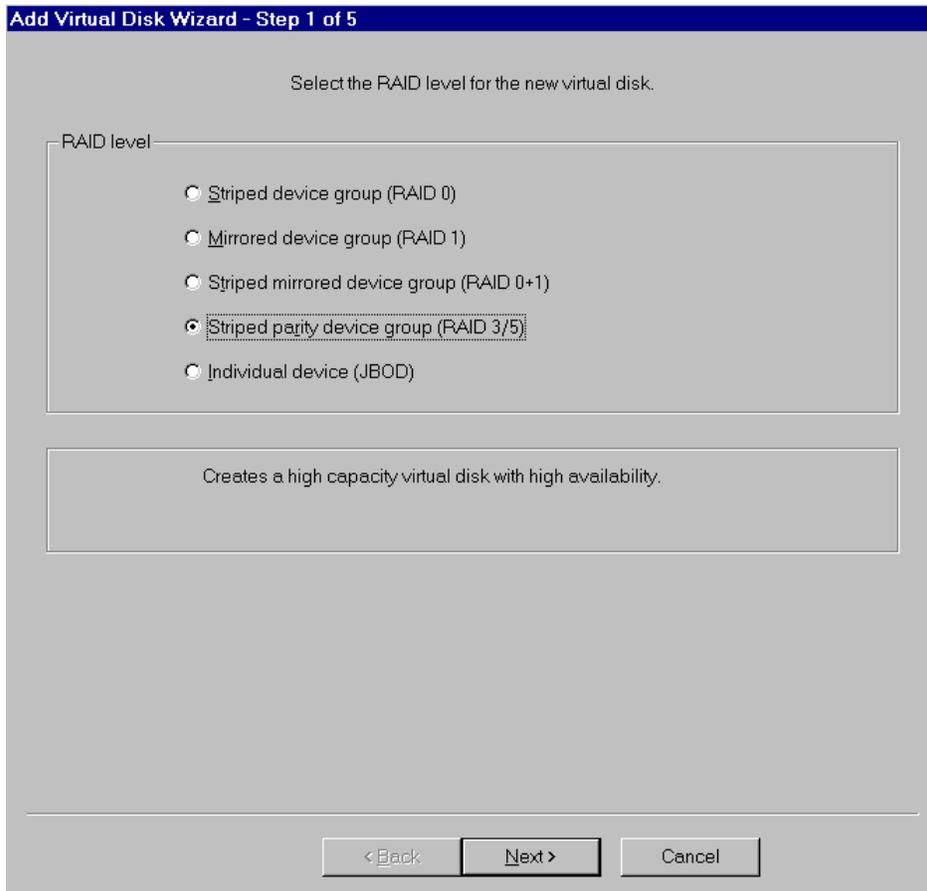
2. Select the Devices for Your Virtual Disk

Select the devices you want to include in the virtual disk by clicking on them in the *Available storage*: window pane (Figure 2–27). As you click them, they are listed in the *Selected devices:* window pane. Each wizard step provides your with the information you need to successfully create a virtual disk. In this example the wizard tells you that you need to select *at least* three disks to create a RAID 5 virtual disk.

Figure 2–27 Add Virtual Disk Wizard Step 2
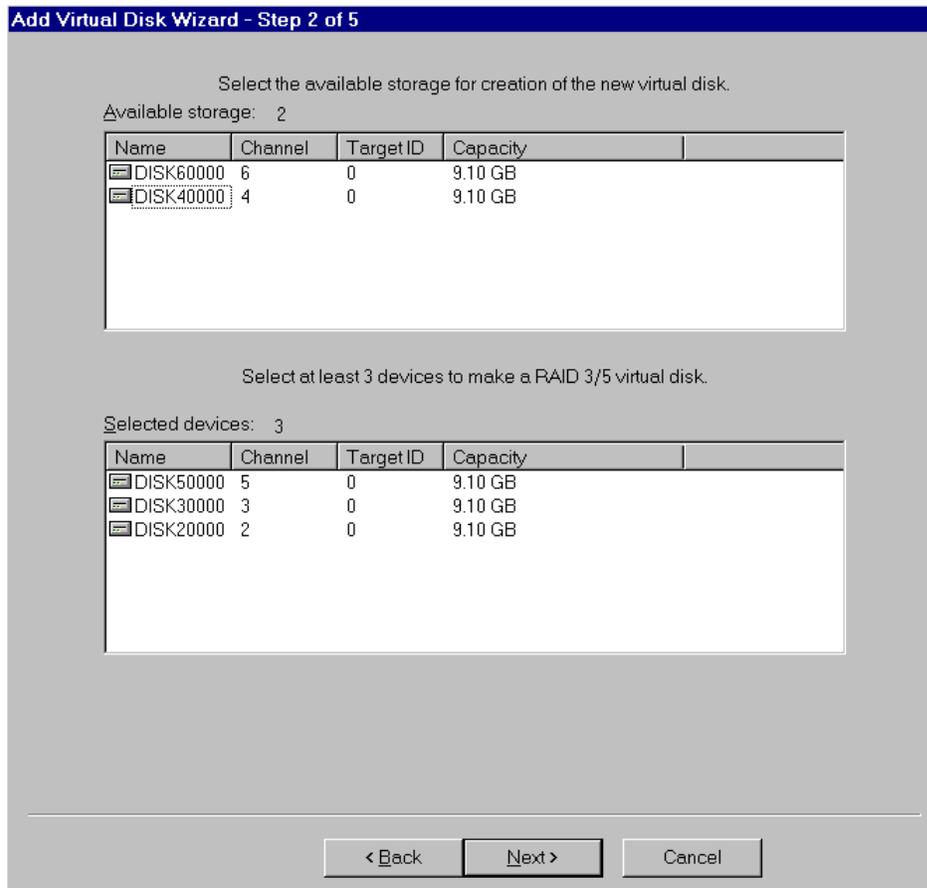
3. Set the Capacity of the Virtual Disk

Set the capacity of the virtual disk in the *Capacity for virtual disk:* window (Figure 2–28).
The applet will inform you of the minimum and maximum capacity available to you.
Click *Next to* continue.

Figure 2–28 Add Virtual Disk Wizard Step 3
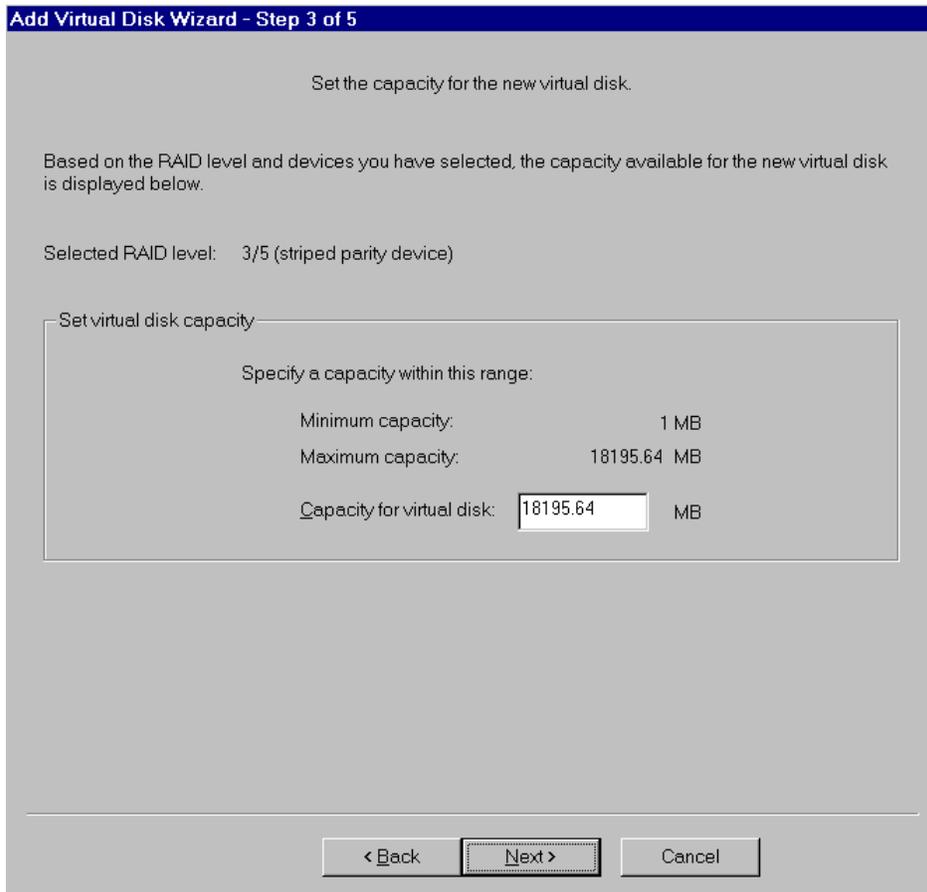
4. Set the ID, LUN, and Operating Parameters

Specify the unit number, operating parameters, and options for the virtual disk; then click *next*. If you wish to save the controller configuration during the process of the wizard, select the *Save controller configuration to virtual disk setting* at step 4 of 5 (see Figure 2–29). Making this selection changes the default setting.

Figure 2–29 Add Virtual Disk Wizard Step 4

5. Confirm and Create the Virtual Disks 1 and 2

The final Add Virtual Disk Wizard window Step 5 (Figure 2 –30), recaps the choices you have made in Steps 1 through 4. If you are satisfied with your choices, click *Finish* to start creating the virtual disks.

Figure 2–30 Add Virtual Disk Wizard Step 5
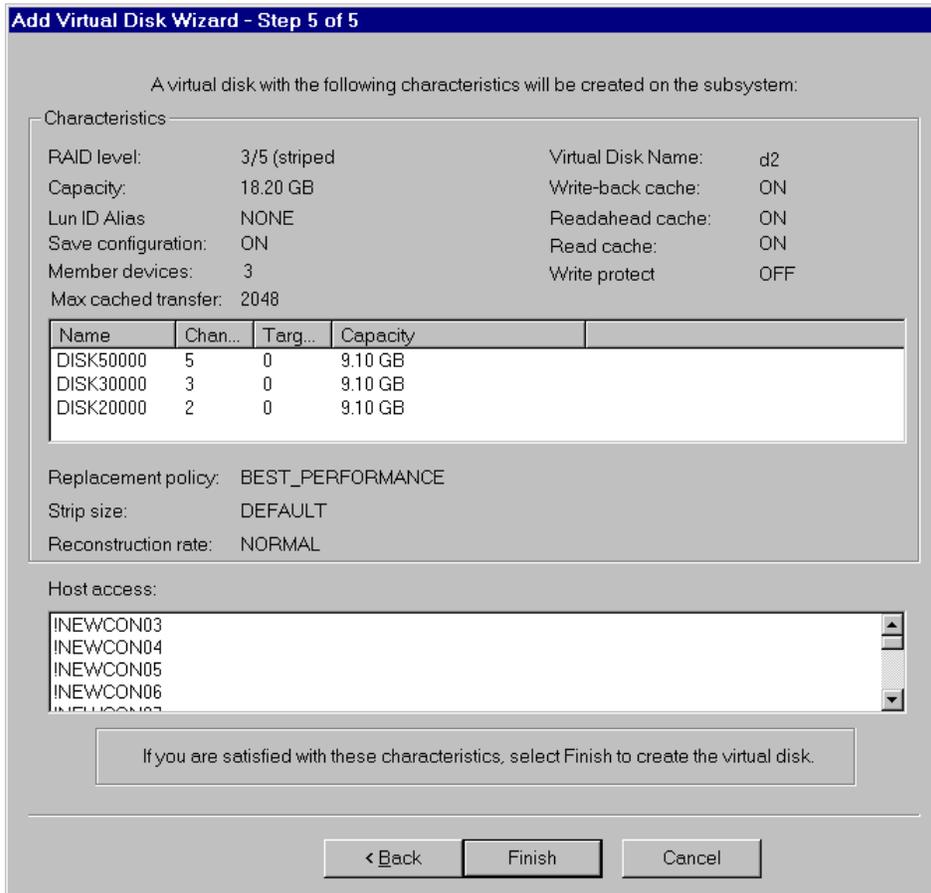
Figure 2–31 shows what the Storage Window looks like *during* the creation of the virtual disk. You should not use the virtual disks until they are finished being built.

Figure 2–31 Storage Window During Creation of Virtual Disks D2
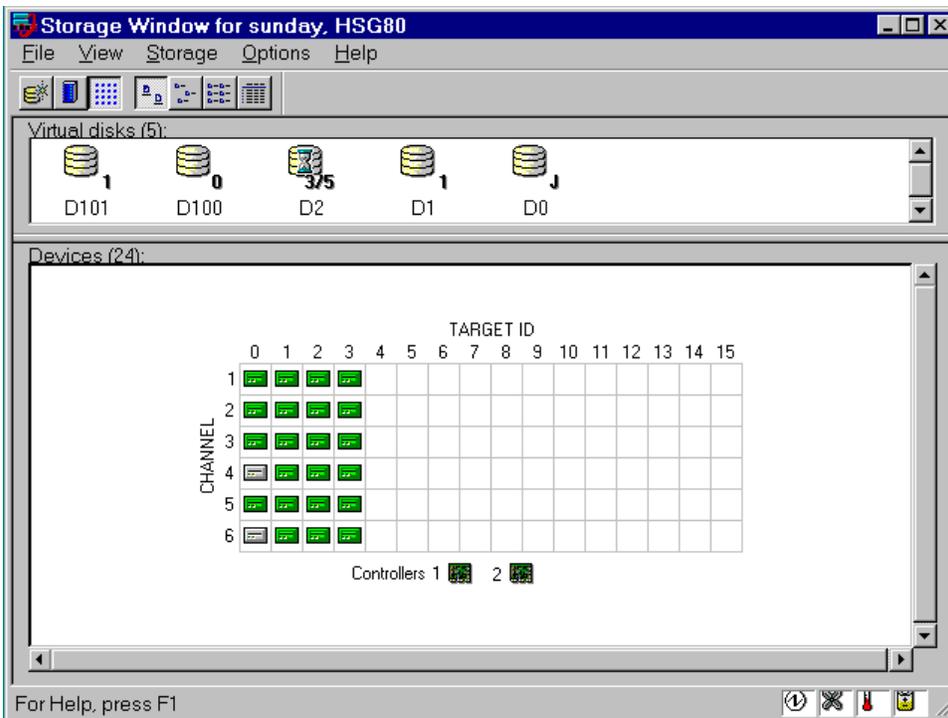
2.8   Deleting Virtual Disks

CAUTION

The virtual disks in your storage subsystem are logical units that contain your user data. Although they do not exist in a physical sense, it is their logical structure that ties together the physical pieces of your data spread across their members. Use extreme caution when deleting a virtual disk. You may be deleting valuable, user data.

1. Under Windows NT, use the Disk Administrator to first delete the partition. Under UNIX, ensure that any filesystem(s) that may have been mounted on the device have been unmounted and removed. If the drive was part of a logical volume the device should also be removed from the logical volume system.

2. Click on the icon of the virtual disk you want to delete.

3. On the Storage menu, select *Virtual Disk* , then *Delete*. Or press the *Delete* key on your key board.

4. If Client prompts you for your password, enter it.

5. When Client prompts to confirm the change, click the *Yes* button to continue. Client deletes your virtual disk from your configuration and refreshes the Storage Window.

2.9   Modifying Virtual Disks

You can modify the characteristics of your virtual disks in two ways: you can change its operating characteristics or remove members.

CAUTION

The virtual disks in your storage subsystem are logical containers that hold your user data. Use extreme caution when modifying the characteristics of a virtual disk. You may be putting valuable, user data at risk.

2.9.1 Changing Virtual Disk Operating Parameters

You can change the characteristics of any of your virtual disks by accessing their property sheets.

1. Double-click the icon, or right-click and choose *Properties* of any virtual disk to access its property sheets. The first sheet displayed is the *General* tab (Figure 2–32). This tab is informational and lists the characteristics of the virtual disk. The example given in Figure 2–32 shows the properties for the virtual disk, D6 which was created in the the previous section.

Figure 2–32 Virtual Disk Properties General Tab
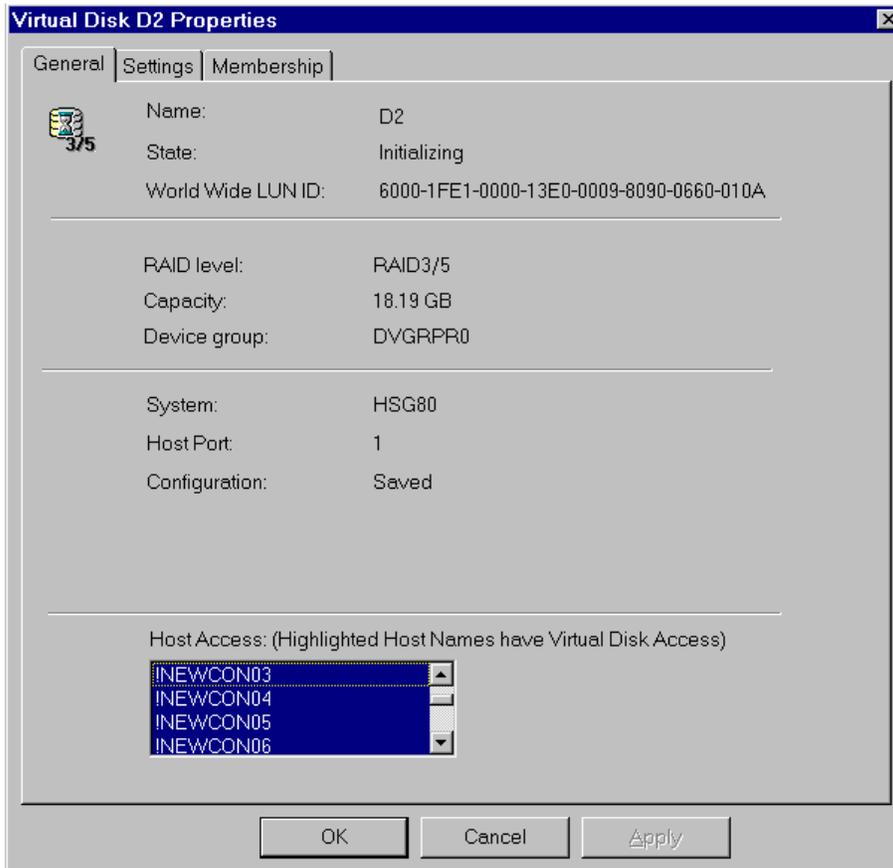
2. Click the *Settings* tab to access the settings property sheet (Figure 2 –33). You can change the chunk size, reconstruction rate, replacement policy, and maximum cached transfer blocks. You can also enable or disable writeback cache by clicking on the checkbox. Click the *OK* button to save the changes and exit.

Figure 2–33 Virtual Disk Properties Settings Tab
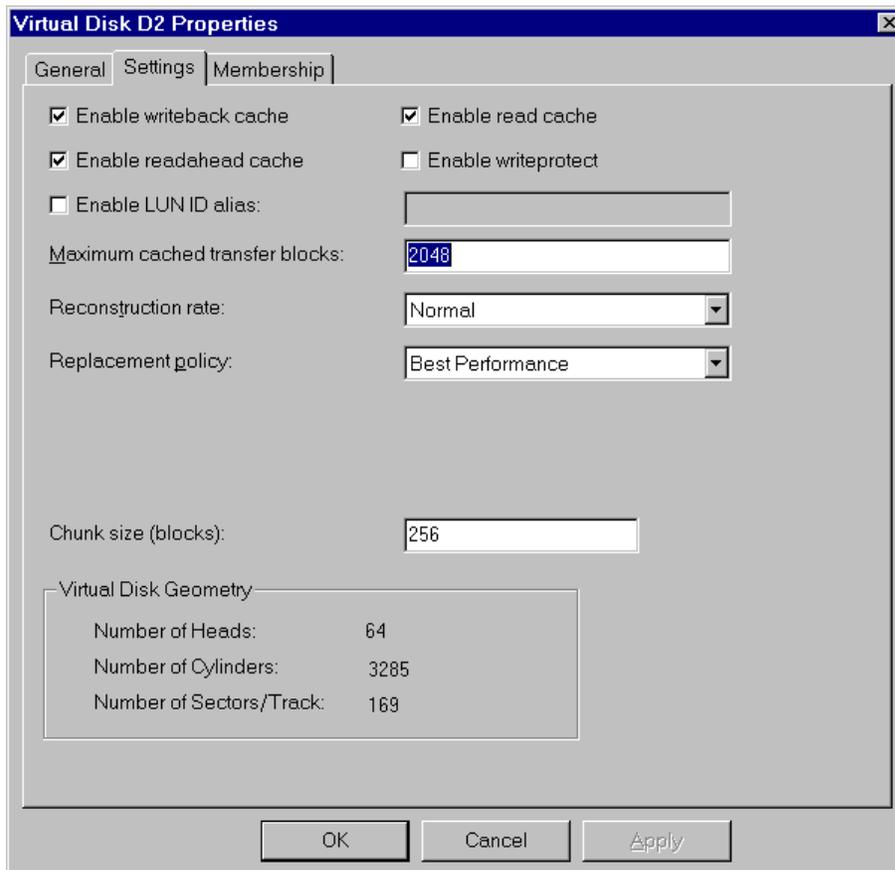
3. Click the *Membership* tab (Figure 2–34) to display member devices properties. In this window pane, you can view a device by name, and the device's current state, channel, target, and capacity.

Figure 2–34 Virtual Disk Properties Membership Tab
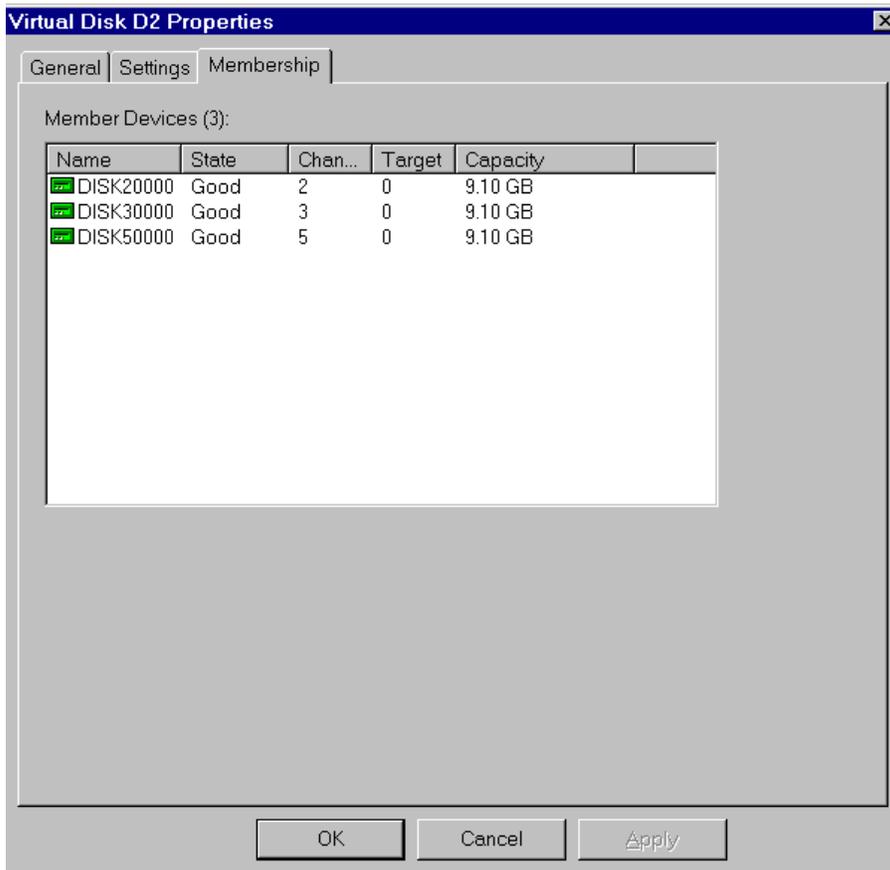
2.10 Setting Passwords and Security Options (Network Only)

Password security prevents unauthorized users from changing or removing storage configurations. In Command Console, any operation that involves changing controller, virtual disk, or device settings requires a password. You do not need a password to view Command Console storage configurations. A password is not required when Client is connected locally to the storage subsystem via a serial port or host SCSI bus connection. The two security options, selectable from the Storage Window's Options menu are: *View Only* and *Make changes*.

Figure 2–35 Security Options Menu

View Only Mode

In view only mode, users can view configurations and monitor storage status, but cannot make changes to controller, virtual disk, or device properties. To prevent others from making changes, from the Options menu, choose *View Only*.

Password Security

Password security prevents unauthorized users from changing or removing storage configurations. You created a password when you installed a Command Console Agent on a system. The password is encrypted and stored in the file where the Command Console Agent resides. To set or change a password, use the Agent Maintenance Menu.

Command Console allows a valid password to make configuration changes for one hour. If there is no user input for one hour, Command Console resets itself to view only. Any configuration change resets the timer and keeps the password active for one hour.

2.11 Managing and Creating Spare Devices

Making a spare device part of a virtual disk will protect the integrity of the RAID setup. This proves especially true for virtual disks that have RAID requirements of two or more devices. Should one device fail, the virtual disk will automatically activate the spare device, as an instant replacement for the failed one.

To create a spare:

1.   Click the device in the Storage Window that you want to make as a spare. (The device must be available, and it must have a capacity equal to or greater than the lowest capacity drive in the group for an automatic failed drive replacement.)

2.   From the Storage menu, choose *Device*, then choose *Make Spare*.

How a Spare Works

*   If a device fails in a RAID 1, RAID 0+1, or RAID 5 virtual disk, the spare automatically replaces the failed device and the controller reconstructs all virtual disks of which the failed device was a member.

*   Once the controller writes data to a spare, the spare becomes part of the same device group in which a device failed.

2.12   Clearing Failed Devices

You can easily identify a failed device in the Storage Window because its icon will appear with a red X covering it. To clear the failed device, from the Storage menu, select *Device*, and then select *Delete*.

2.13   Using Configuration Files

You can use a configuration file to save a particular subsystem configuration and to view or reconstruct it at a later time. You can revert to one of these saved configurations at any time. Client reads the configuration file you choose and sets up your subsystem accordingly.

Saving Configuration Settings to a File

1.   From the Storage menu select *Controller Configuration*, then select *Save* (Figure 2-36).

Figure 2–36 Configuration File Submenu
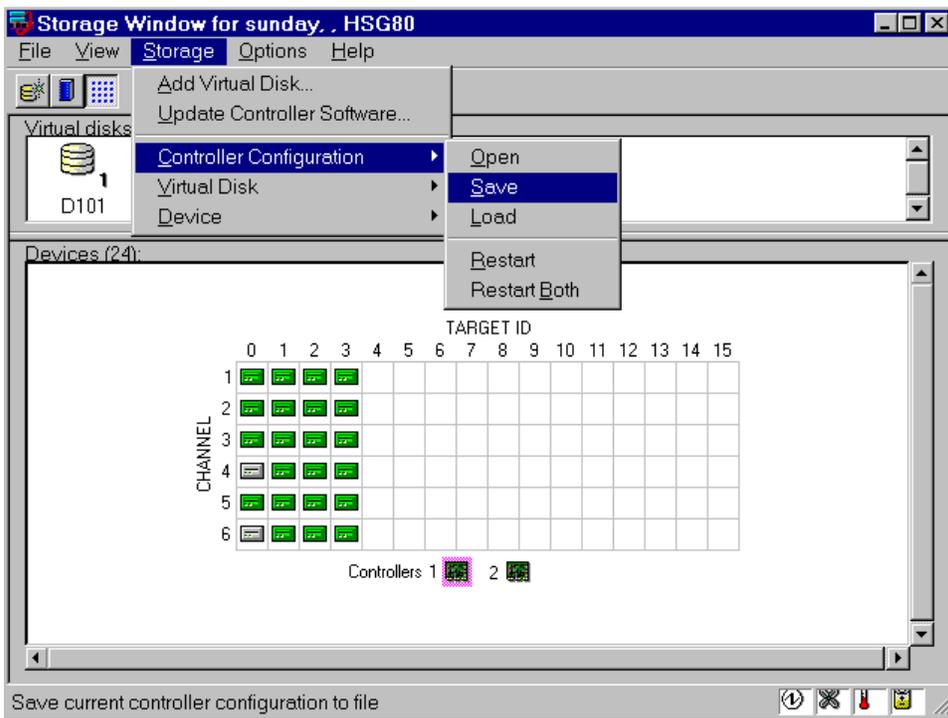
2. When the *Save As* dialog box displays, specify a file name in the *File name:* window.(Figure 2–37), then click the *Save* button. The current configuration settings, including caches, LUNS, host port, and stripe size will be saved to a file at location you specify.

Figure 2–37 Saving to a File
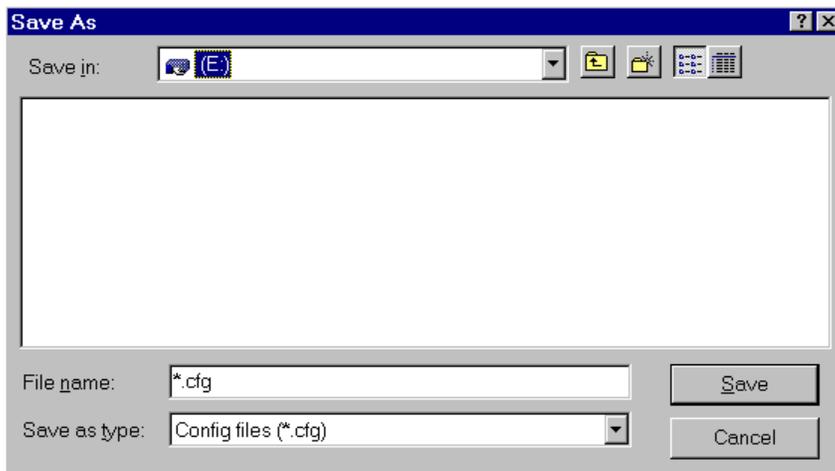
NOTE

In order to restore the current configuration settings, all controllers must be attached to a serial connection.

Restoring Configuration Settings from a File

You can automatically restore/configure your entire storage subsystem from a previously saved configuration file. Client can configure your subsystem from a file only if the file corresponds to your subsystem's hardware configuration. Make sure that the file you select corresponds to your subsystem's hardware configuration.

CAUTION

Make sure you are completely aware of the configuration information contained in the configuration file you choose. If you choose a file that is not compatible with the current configuration of the data on your devices, you put your data at risk.

---

1. From the Storage menu select Controller, then select Configuration, and then, select Load.

2. Click one of the load option radio buttons:

   If you need to reconfigure a failed controller with existing virtual disks, select *Load configuration only,*. (Figure 2–38). Client recreates your virtual disks, but doesn't initialize them. Your data is not overwritten.

Figure 2–38 Load Configuration
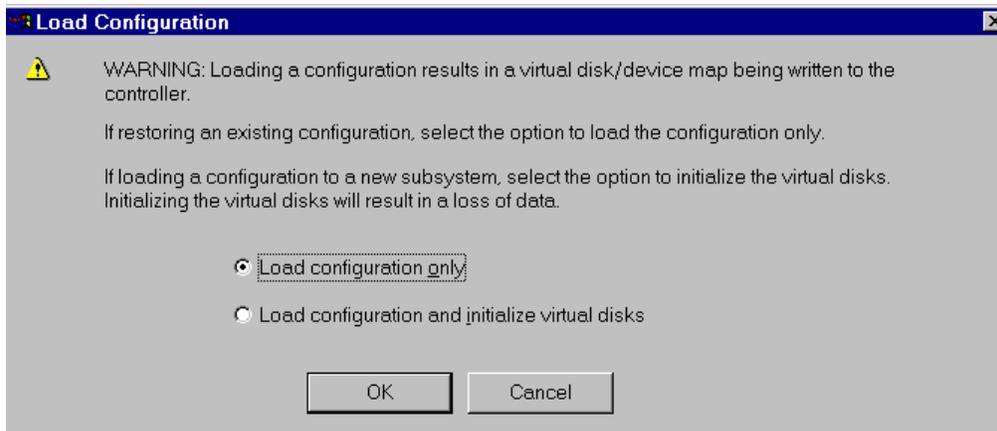
   If you are configuring a new system, select *Load configuration and initialize virtual disks*. **Caution:** Do not use this option on an existing system. Client creates your virtual disks and initializes them. Your data will be lost.

Click, *OK*. When the *Open* dialog box displays (Figure 2–39), specify the location (path) and file name in the text box, then click the *Open* button. The configuration settings as defined in the file, including caches, LUNS, host port, and stripe size will be restored.

Figure 2–39 Open Dialog Box
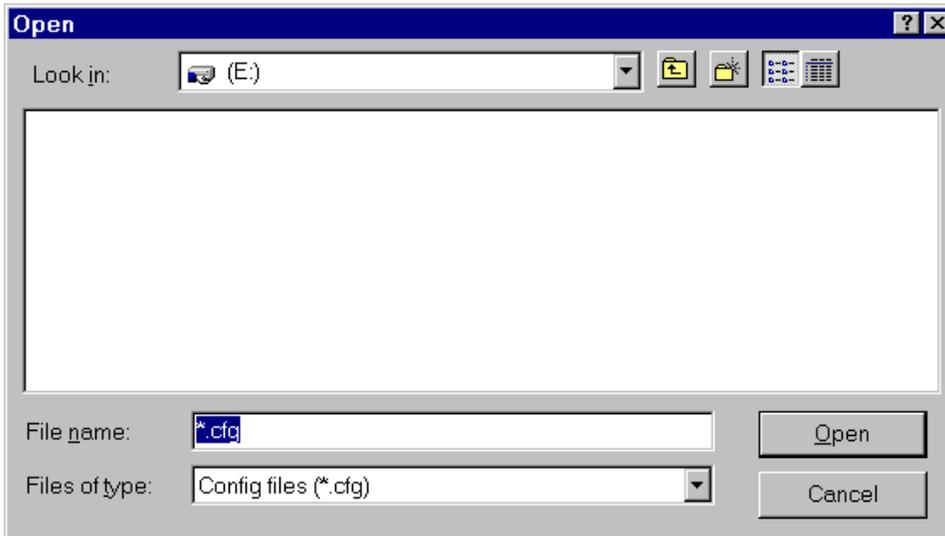
CAUTION

You will be unable to restore the original configuration settings if the original hardware and connections have been replaced, or if the original configuration settings were not saved before implementing the new changes.

2.14   Monitoring the Condition of Your RAID Array 8000/ESA12000

Command Console provides you with two ways to monitor your storage subsystem; visually on the PC screen or remotely using pagers. It may be helpful to familiarize yourself with the Legend of Icons (Figure 2–40) for Command Console as they are used to show the status of Virtual disks, Devices, Controllers, Cache Battery, and the Environmental Monitoring Unit (EMU).

Figure 2–40 Legend of Icons for Command Console
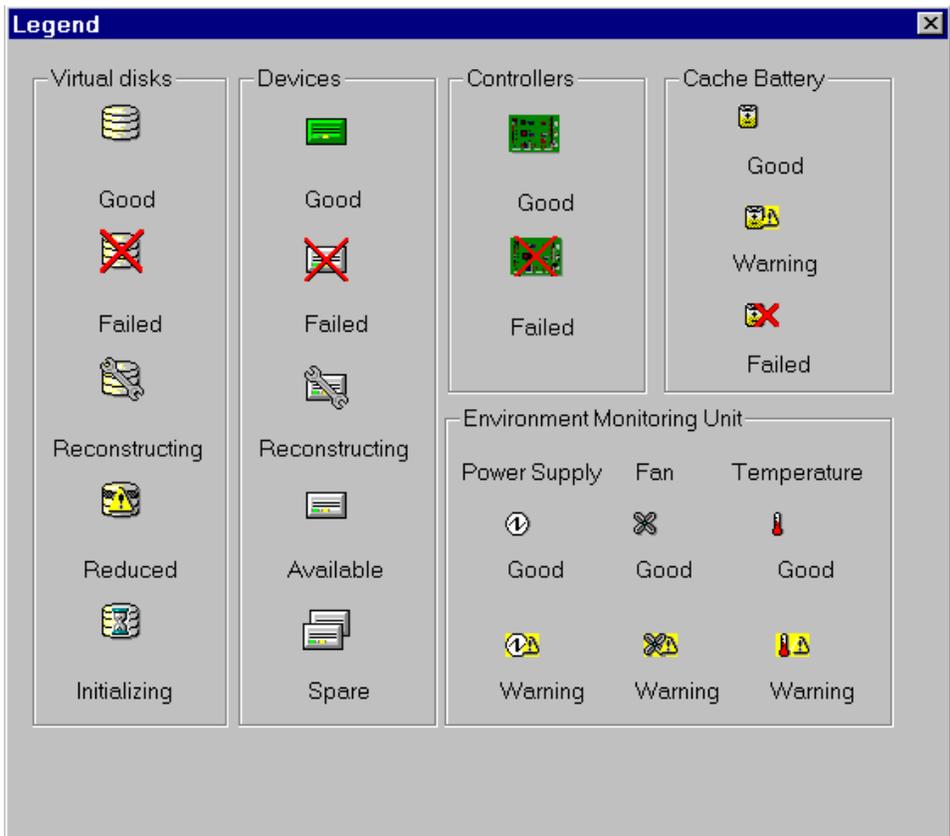
2.14.1 Monitoring Visually

Icons in the Navigation Window and Storage Window show the status of host systems, storage subsystems, devices and controllers. In addition, an EMU indicator in the lower-right corner summarizes the condition of the power supply, UPS, and fan.

If you are connected locally to your RAID Array 8000/ESA12000, you must manually refresh the screen to check system status. To refresh the screen from the Storage Window's View menu select *Refresh*.

Navigation Window Icons Show System Status

The icons in the Navigation Window change to show system status. For example, if a device fails in a storage system connected to a controller on a system, the failure will be indicated by the system folder. A yellow exclamation mark indicates a warning condition and a red X failure. The Navigation Window Status Bar (Figure 2 –41), located in the bottom right of the window, summarizes host system connection status.

Figure 2–41 Navigation Window Status Bar
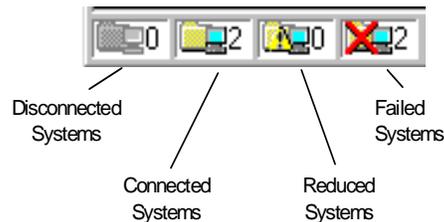
Storage Window Icons Show Device, Controller, and Virtual Disk Status

The Storage Window provides detailed information in both physical and logical views of an individual storage subsystem in the Storage Window. In each view, device and virtual disk states are represented by icons. These icons change to indicate status. A red X through a device indicates device failure. The Storage Window Status Bar (Figure 2 –42) indicates power, fan, temperature, and battery states.

Figure 2–42 Storage Window Status Bar
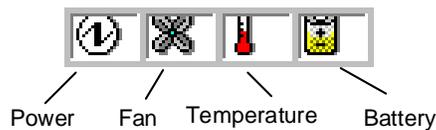
2.14.2 Monitoring the Subsystem Using Your Pagers

Command Console will monitor events on your storage subsystem and notify you of critical events by pager. You can use either an alphanumeric or numeric pager. If you use an alphanumeric pager text messages will be sent to you. The message will include the host system name, storage subsystem name and a text event message. For numeric pagers you must assign numbers to your host systems and storage subsystems. The event message will also be a four-number code, indicating a critical event.

Command Console can send messages that notify either you or a group of people you designate at different shift times.

The steps to setting up critical event notification by pager are:

- Set pager preferences. This includes setting your modem's COM port, pager service baud rate and polling interval.

- Define a user profile. This includes entering your pager number, identification number (alphanumeric pagers only), and the time period in which you want to be notified.

- For numeric pagers only, assign pager codes. This includes assigning three-digit numbers to the host systems and storage subsystems you want to monitor.

- As an option, create a group of people to notify from the user profiles you have defined.

Setting Pager Preferences

1. From the Options menu, click *Preferences*.

Figure 2–43 Event Notification Preferences Settings
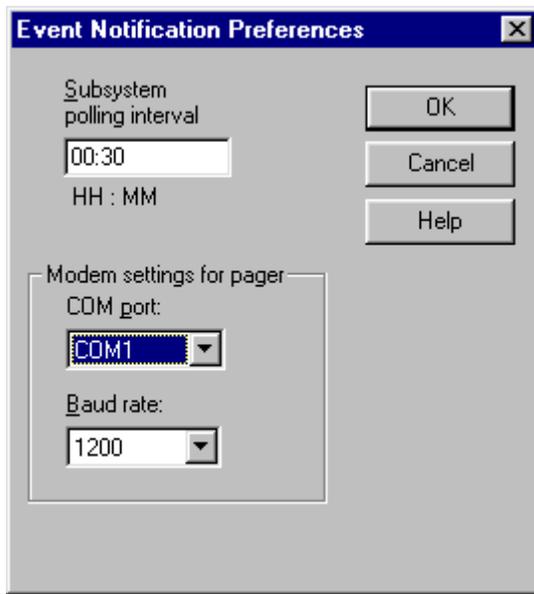
2. Enter the polling interval in the subsystem polling interval box. The standard, default polling interval is 30 minutes.
3. Select the COM port that your modem uses on your computer.
4. For alphanumeric pagers select the baud rate, usually 1200, of your paging service provider's TAP modem.

Defining a User Profile

1. From the Options menu, click *Event Notification*, then User Profiles.
2. Click the *New* button and enter a name on the *General* tab. A description is optional. Use *Edit* to make changes to an existing profile.
3. Select the *Pager* tab (Figure 2–44).

Figure 2–44 Pager Tab Configured for an Alphanumber Pager
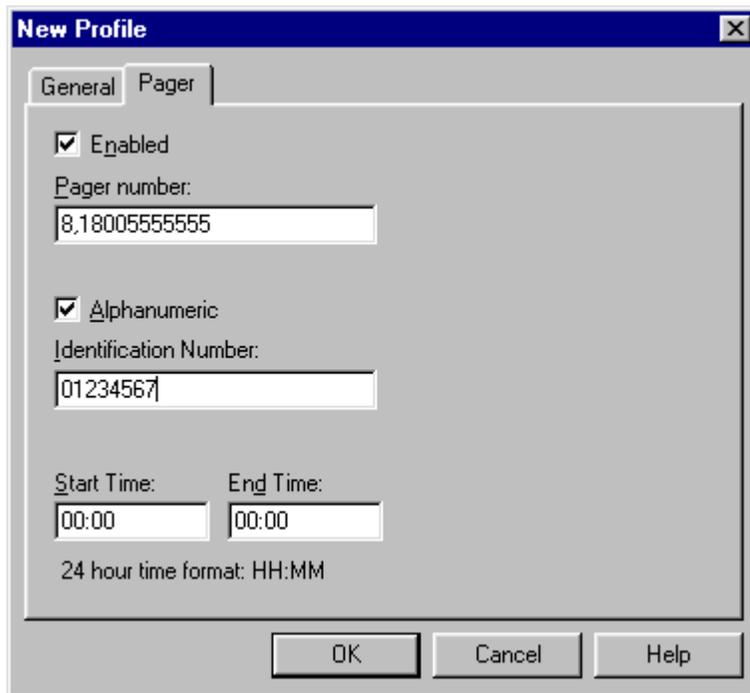
4. Click the *Enabled* checkbox.

5. If you have a numeric pager, enter its phone number in the *Pager number:* box.

6. If you have an alphanumeric pager, enter service provider's TAP phone number in the *Pager number:* box. Click the *Alphanumeric* checkbox and enter the pager identification number in the *Identification Number:* box.

7. Enter the start and end times in HH:MM, 24-hour format. For example: 08:00 as start, and 17:00 as end. Leave both fields set to 00:00 if you want to be notified 24 hours a day.

Tips and Notes for Setting Up Pagers

- If you need to dial a code to access an outside number, it must precede the phone number in the pager number field.

- Do not use commas with alphanumeric pagers, except if required to access an outside line.

- If you are configuring an alphanumeric pager you must supply the paging service provider's TAP service modem number in the pager number text box.

- When you enter a numeric pager number, you can optionally use one or more commas to specify a two-second wait interval. Each comma causes the pager dialer to wait two seconds before sending the pager identification number to the paging company. You may need to add additional commas using the trial and error method so that Command Console will pause long enough for the pager company greetings to complete.

- Not every paging company requires a pager identification number. If your pager is assigned a unique phone number to call the paging company, do not enter a pager identification number.

- Command Console supports the TAP paging protocol.

- See your pager documentation to determine which special characters are supported.

Setting Up Pager Codes for Numeric Pagers

Numeric pager support lets you receive notification about critical events by means of numerical coded messages. A numeric pager usually has three fields separated by hyphens. The first field signifies the area code, the second field signifies the exchange, and the third field signifies the four-digit extension. When you receive a page from Command Console, the first field signifies a three-digit number for the host system, the second field is a three-digit number for the subsystem controller, and the last field is a four-digit code for the critical event.

Figure 2–45 Understanding a Numeric Pager Readout
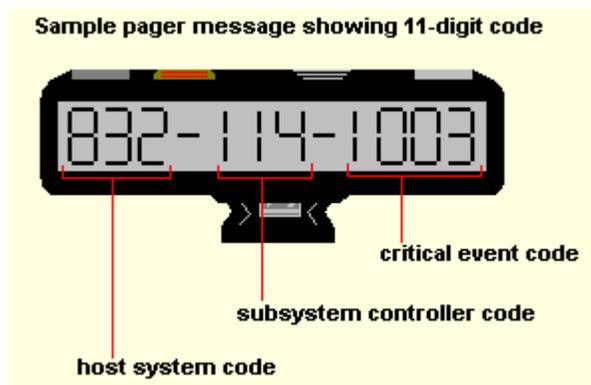
First, set up pager codes by assigning a three-digit code for each host system. Next assign another three-digit code for each subsystem controller connected to that host. The last four digits are pre-assigned event codes. If you choose, the event codes can be modified. System codes, subsystem controller codes, and event codes can only be numeric. Alphanumeric codes are not supported.

Assigning a Pager Code to a Host System

1. Click a host system in the Navigation Window.
2. From the File menu, click *Properties*.
3. Click the *Event Notification* tab.
4. Enter a three-digit number in the *Pager identification number:* text box.

Assigning a Pager Code to a Subsystem Controller

1. Click a subsystem controller in the Navigation Window.
2. From the File menu, click *Properties.*
3. Click the *Event Notification* tab.
4. Enter a three-digit number in the *Pager identification number* text box.

Viewing and Changing a Pager Critical Event Code

1. From the Options menu, click *Event Notification Pager Codes*.
2. Click the code you want to edit, then select *Edit*.
3. Type the new four-digit pager code in the text box. The default codes are:

Table 2–1 Pager Critical Event Codes

Critical Event	Code	Description
Disk Fault	1000	Physical device failure
Power Supply Fault	1001	
Fan Fault	1002	
Battery Fault	1003	
Temperature Fault	1004	
Controller Fault	1005	The controller has detected an internal fault
Communication Fault	1006	Agent cannot communicate with the subsystem
Other Controller Fault	1007	One controller is reporting a fault on the other
External Fault	1008	
LUN Fault	1009	Virtual disk failure

Setting Up Notification Lists

1. From the Options menu, click *Event Notification*, select *Notification List*.
2. Click the *New* button and enter a name for the new notification list in the *Name:* text box (shown in Figure 2–46).

Figure 2–46 Defining a Notification List
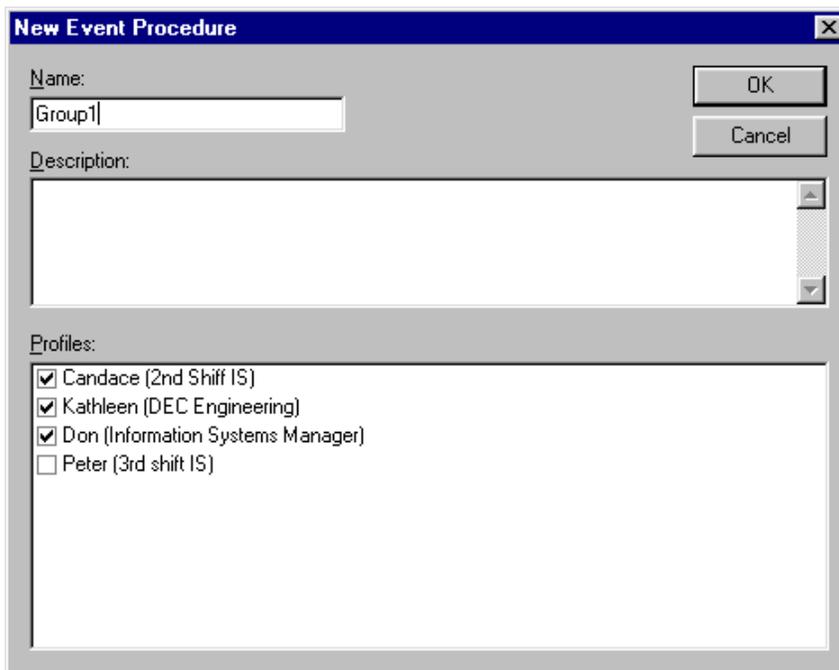
3. Write a brief description of the list (optional) in the *Description:* box.
4. Add members to the list by clicking the checkbox next to each person's name. When you are finished click *OK*.
5. Next, associate a host system or controller to the list you have just created.

Associating a Host System or Controller with a Notification List

1. Click a host system or a controller subsystem in the Navigation Window.
2. From the File menu, click *Properties*.
3. Click the *Event Notification* tab.
4. Select a list from the *Notification list:* drop-down list; then click *OK*.

Figure 2–47 Selecting a Group for Notification
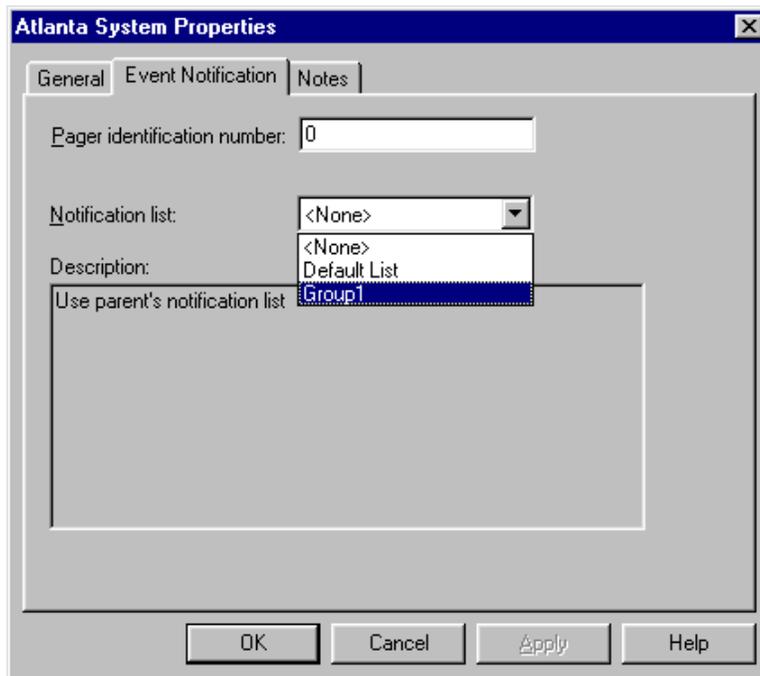
Tips and Notes for Creating Notification Lists

- Use notification lists to notify groups of people of critical events.

- You cannot add a user to a list until a user profile has been created for the user. Be sure you have created the user profiles first before creating a notification list.

- Use the property pages to associate notification lists with systems and subsystem critical events.

- You must first create a notification list before you can associate a system or controller with a notification list.

2.15 Using the Command Line Interpreter (CLI) Window

The Command Line Interpreter (CLI) is a text-based interface utility for monitoring and configuring your storage subsystem. Because the configuration of the storage subsystem can be changed using CLI commands, we recommend that you allow CLI access only to users who are familiar with CLI operation.

The CLI Window (Figure 2–48) has a command-entry area and a command response area below it. You can enter commands in the command-entry area, and the controller responds with the results of the entry in the response area.

Accessing the CLI Window:

There are two ways to access the CLI Window from the Start menu or by double-clicking on the CLI Window icon in the program group.

From the Start menu:

Select *Command Console* then CLI Window applet.

From the Navigation Window:

1. In the Navigation Window, click the subsystem's host system icon to expand the tree and display the CLI Window icon
2. Click the icon for the subsystem you wish to use the CLI Window with. The Navigation Tree expands the display to show the windows and pages available with this host system.
3. Click the CLI Window icon for your subsystem. If the program prompts you for your password, enter it and the CLI Window appears. Enter a CLI command in the command-entry area.

Figure 2–48 CLI Window
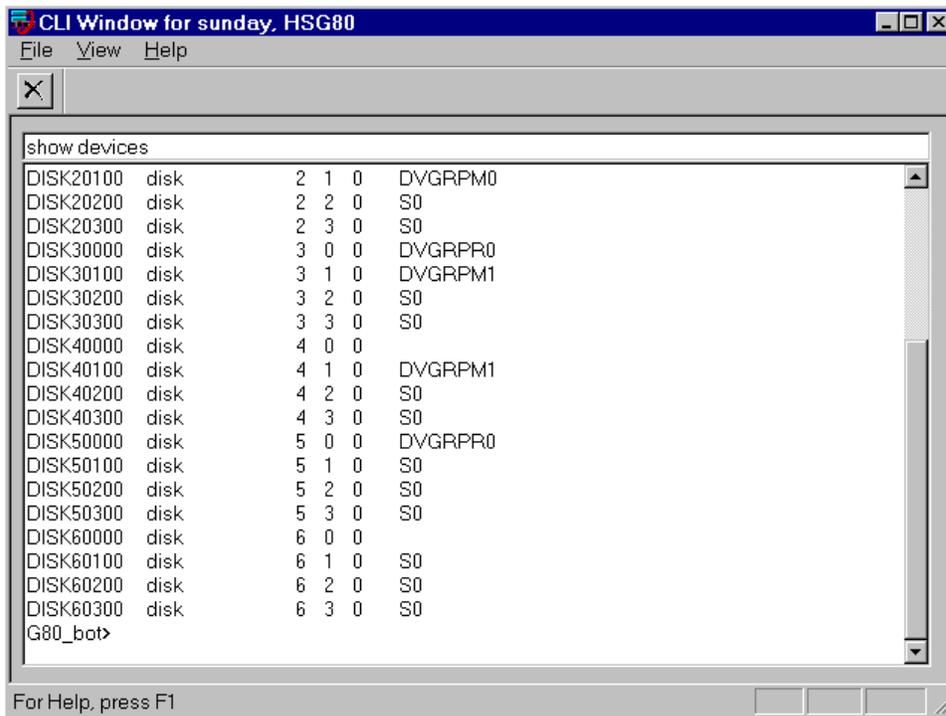
3

Installing the StorageWorks Command Console for the Fibre Channel Switch

This chapter contains instructions for installing StorageWorks Fibre Channel Switch software (Client and Agent) required to support operation of the Fibre Channel switch in a Windows NT environment. (This chapter is not applicable to fibre channel hub installation).

3.1   Initial Fibre Channel Switch Configuration

Prior to installing the StorageWorks Client and Agent support software for the Fibre Channel switch, you must properly configure the Fibre Channel switch regarding the following elements:

- Ethernet IP Address
- Ethernet Subnetmask
- Fibre Channel IP Address
- Fibre Channel Subnetmask
- Gateway Address
- Domain

Switch configuration is accomplished through the switch's Configuration Menu, accessed through its front panel (or Telnet connection). Please reference the Quick Setup Guide for your Fibre Channel switch for procedures to establish the proper configuration settings before proceeding.

NOTE

After configuring the switch from the front panel or Telnet connection, you must reboot the switch to implement changes.

3.2 Installing the Fibre Channel Switch Client

This chapter contains instructions on how to install the Fibre Channel Switch Client on Windows 95, and Windows NT version 4.0 (Alpha and Intel) for network access to the Fibre Channel Switch Agent. The Fibre Channel Switch Client, that provides the Fibre Channel Fabric window, renders a graphical interface for your Fibre Channel Switch Agent. The Fibre Channel Switch Client provides the Fabric window, which displays the status of your switches. A fabric is a collection of switches, bridges, and hubs. For this software release, the Fibre Channel Switch Client will only provide information about your switches.

CAUTION

Before you install the Fibre Channel Switch Client, do the following:

- Verify that you have installed Command Console Client version 2.1. The Client software packages must be on the same Client system. You can determine a Client's version by selecting *Help|About* in its window.

- Read the `fcswitch.txt` file for more information on the installation.

- Exit the Command Console before you install the Fibre Channel Switch Client. This will allow the Command Console Client to recognize the new Client.

- Verify that you have installed one of the following browsers: Netscape Navigator version 4.0 (or later), Internet Explorer version 4.0 (or later), or HotJava version 1.1 (or later). Your Fibre Channel Switch Client will use one of these browsers to display the status of your switch.

3.2.1 Beginning the Fibre Channel Switch Client Installation

To install the Fibre Channel Switch Client, perform the following tasks:

1. Place the *Compaq StorageWorks Fibre Channel Switch* CD, (part no. AG-RHC1A-TE, packaged with your Fibre Channel Switch), into the CD-ROM drive.

2. Run *File Manager* or *Windows Explorer.*

3. Navigate to the folder: *drive_letter:\SWCC\NTIntel\client*

4. Double-click on *Setup*.

3.2.2 Completing the Fibre Channel Switch Client Installation

1. Click <u>Next</u> to continue. The Command Console License Agreement window appears.

2. Click <u>Next</u> to continue. The software may display a message saying that it has found Command Console and that it will now install the Fabric window into the Command Console directory structure.

3. Click *OK.* The Fibre Channel Switch Client is installed into the directory in which the Command Console Client is located (`C:\Program Files\swcc`). In most cases, the Fibre Channel Switch Client will be installed into the following directory:

   ```
   C:\Program Files\swcc\FibreWindow
   ```

4. Start the asynchronous event service (AES).

AES must be running for the Fabric window to function properly. AES allows users to see traps and refreshes to the Fabric window (for example, to show that a switch has been added). AES starts when your system is booted.

Windows NT:

To stop or restart AES, click the *Services* icon located under the Control Panel. "*AsyncEventSvc*" is the entry for AES in the Services window. AES is one of several services that you can start, stop, pause, and continue under the *Services* icon

Windows 95:

To stop or restart AES, click the *Async Event Service* icon located under the Control Panel. To disable the automatic start of AES when your system boots, deselect the *Automatic Start Upon Boot* option, and then click *Apply*.

3.2.3 Removing the Fibre Channel Switch Client from Windows NT

When you remove the Fibre Channel Switch Client, you will remove the Fabric window. The Command Console Client will still be on your computer. The following procedure tells you how to remove the Fibre Channel Switch Client on Windows 95, and Windows NT.

1. Click *Settings* under the *Start* menu, then click *Control Panel*.

2. Double-click the *Add/Remove Programs* icon in the Control Panel. The Add/Remove Program Properties window appears.

3. Highlight *StorageWorksFibreApplet* located in the window, and then click *Add/Remove*. The computer removes the Fibre Channel Switch Client.

3.3   Installing the Fibre Channel Switch Agent

This chapter contains instructions for installing the Fibre Channel Switch Agent on Windows NT version 4.0. Whenever you install the Fibre Channel Switch Client to operate over a network, you must install the Fibre Channel Switch Agent. The Agent is responsible for establishing communication with the Client and its subsystems on the network.

Compaq recommends that you install the Fibre Channel Switch Agent on only one host system. You may install the Fibre Channel Switch Agent on multiple host systems if you intend to manage separate collections of Fibre Channel Switches.

Instructions on how to install the Fibre Channel Switch Agent on the Windows NT version 4.0 operating system are separated into two sections: "Beginning the Installation" and "Completing the Installation." The first section tells you how to access the installation file from the CD-ROM. The second section describes how to configure your Agent during the installation.

CAUTION

Before you install the Fibre Channel Switch Agent, do the following:

- Verify that you have the Windows NT SNMP service installed on your computer. If this service is not installed, the Fibre Channel Switch Agent will be unable to monitor your Fibre Channel switches. The Windows NT SNMP service is available on your Windows NT installation CD-ROM. To check, double-click *Services* in *Start|Settings|Control Panel*. After SNMP service has been installed, re-install Windows NT service pack 3.

- Read the `fcswitch.txt` file for more information on the installation.

CAUTION

- Verify that you are logged into an account that is a member of the administrator group. All Fibre Channel Switch Agent installations on Windows NT must be done locally. Do not attempt to install the Agent over the network.

- For the Alpha computers running Windows NT, verify that you have only Windows NT service pack 3 (*My Computer|Help|About*).

- For the Intel computers running Windows NT, verify that you have Windows NT service pack 3 or 4 (*My Computer|Help|About*).

3.3.1 Beginning the Fibre Channel Switch Agent Installation

To install the Fibre Channel Switch Agent, perform the following tasks:

1. Place the *Compaq StorageWorks Fibre Channel Switch* CD, (part number AG-RHC1A-TE, packaged with your Fibre Channel Switch), into the CD-ROM drive.

2. Run *File Manager* or *Windows Explorer.*

3. Navigate to the folder: *drive_letter:\SWCC\NTIntel\agent*

4. Double-click on *Setup*.

3.3.2 Completing the Fibre Channel Switch Agent Installation

1. Click *Next* to continue. The Command Console License Agreement window appears.

2. Click *Next* to continue. The Command Console Setup window appears.

3. Select the destination directory for the Agent software, then click *Next.* The default is the following: `C:\Program Files\swcc\FibreAgent\`

   The Folder Selection window appears.

4.   Type a name for the folder, then click *Next.* The default name is
StorageWorksFibreAgent. This name will appear in the Add/Remove Program
Properties window.

The software will check for configuration files in the directory that you specified to
install the Agent. You will not have these files if you are installing the Agent for the
first time. The software will do one of the following:

- If any of the configuration files are missing, the installation The Fibre Agent
  Configuration Utility appears. Go to step 5.

- If all three configuration files are found, the program will tell you where it found
  the configuration files. To change the configuration, run the Fibre Agent
  Configuration Utility after the installation. Click *OK.* The program finishes the
  installation and starts the Agent. The following steps do not apply because the
  software has been installed and the Agent has a previous configuration.

5.   Click *Next* to configure your Agent. The Agent Setting window appears.

6.   Type an interval to poll each element in the fabric. The default is 5 minutes;
      however, you can enter an interval from 1 to 60 minutes.

7.   Click *Next*. The Client List window appears.

8.   Type the Client's name and monitor mode: TCP/IP and/or SNMP or none. The
      notification scheme defines the network method that the Agent will use when
      notifying the selected Client of a change in the state in a subsystem, which is
      comprised of switches. You can select one, both, or none of the notification options.
      The following table lists the definitions of the notification schemes:

Notification Scheme	Function
Transmission Control Protocol/ Internet Protocol (TCP/IP)	Is the Client's notification scheme. If you do not select TCP/IP, the Client will only display a subsystem change when you refresh the Fabric window.
Simple Network Management Protocol (SNMP)	Must have an SNMP-compatible monitoring program (ServerWORKS, for example) running on the Client's system.

9.   Click *Add* to add the Client.

10. Click *Finish* if satisfied or <u>*Back*</u> for changes. The program prompts you with the message:

    ```
    Changes have been made to the Fibre Agent Configuration.
    The Agent is not currently running. Start the Agent now?
    ```

11. Click <u>*Yes*</u> to run Agent. The Agent is started and the program finishes loading.

3.3.3 Removing the Fibre Channel Switch Agent from Windows NT

The following instructions tell you how to remove the Agent:

1. Click <u>*Settings*</u> under the *Start* menu, then click <u>*Control Panel*</u>. Double-click *Services*, and select *SWCC Fibre Agent*.

2. Click <u>*Stop*</u> to stop the Agent.

3. Click the *Add/Remove Programs* icon in the Control Panel. The *Add/Remove Program Properties* window appears.

4. Select *StorageWorkFibreAgent* located in the window, and click *Add/<u>Remove</u>*. The computer asks:

    ```
    Are you sure you want to completely remove the selected
    application and all of its components?
    ```

5. Click *Yes*. The Agent is removed.

6. To remove the configuration, Client list, and the parameter files, change to the `C:\Program Files\SWCC\FibreAgent` directory and delete the files listed in the following table. These files may be in a different directory if you did not install the Agent in `C:\Program Files\SWCC\`. If you delete these files, you will need to reconfigure the Agent during reinstallation.

    > **CAUTION**
    >
    > Other Agents use the `client.ini` and `storage.ini` file names, but their files are in different directories. If you cannot locate the files for the Fibre Channel Switch Agent, use the Find command in Windows Explorer to find the `FibreAgent.cfg` file. The `client.ini` and `storage.ini` files for the Fibre Channel Switch Agent are usually in the same directory as the `FibreAgent.cfg` file

File Name	File Type
FibreAgent.cfg	Configuration File
client.ini	Client List File
storage.ini	Parameter File

3.4   Configuring the Fibre Channel Switch Agent

This section contains instructions on how to configure the Fibre Channel Switch Agent on Windows NT. You configure the Fibre Channel Switch Agent by using the Fibre Agent Configuration program. Topics in this chapter include:

- Client Notification Options

- Adding a Client System Entry

- Modifying a Client System Entry

- Removing a Client System Entry

- Restarting the Agent

- Changing the Polling Interval

- Starting and Stopping the Agent

3.4.1 Client Notification Options

The notification scheme defines the network method that the Agent will use when notifying the selected Client of a change in the state in a subsystem, which is comprised of Fibre Channel Switches. You can select one, both, or none of the notification options. The following table lists the definitions of the notification schemes:

Notification Scheme	Function
Transmission Control Protocol/ Internet Protocol (TCP/IP)	Is the Client's notification scheme. If you do not select TCP/IP, the Client will only display a subsystem change when you refresh the Fabric window.
Simple Network Management Protocol (SNMP)	Must have an SNMP-compatible monitoring program (ServerWORKS, for example) running on the Client's system.

3.4.2 Adding a Client System Entry

To add a Client system entry, follow these instructions:

1. Select the Fibre Agent Configure entry in *Start*/*Programs*/*StorageWorks*, as shown in Figure 3–1.

Figure 3–1 Launching the Fibre Agent Configuration Utility
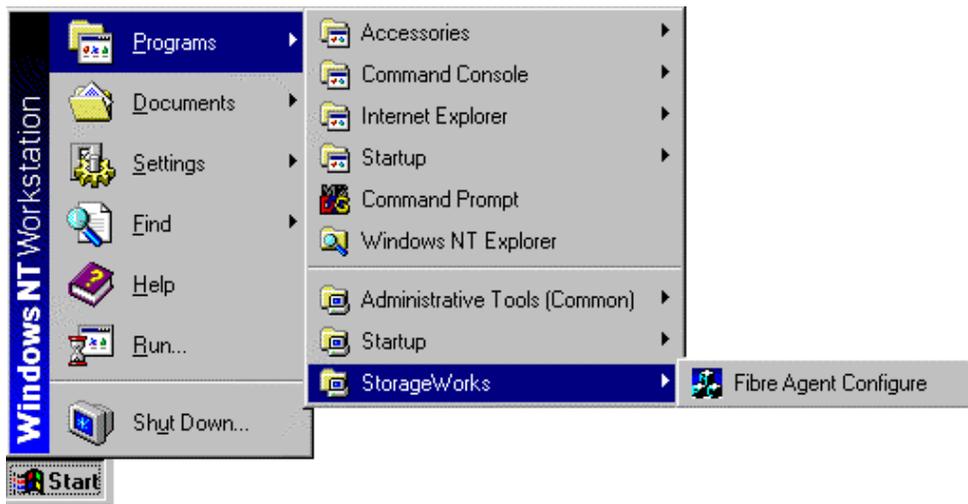
2. When the StorageWorks Command Console Fibre Agent Configuration window appears, click on the *Clients* tab as shown in Figure 3–3.

Figure 3–2 Fibre Agent Configuration Window
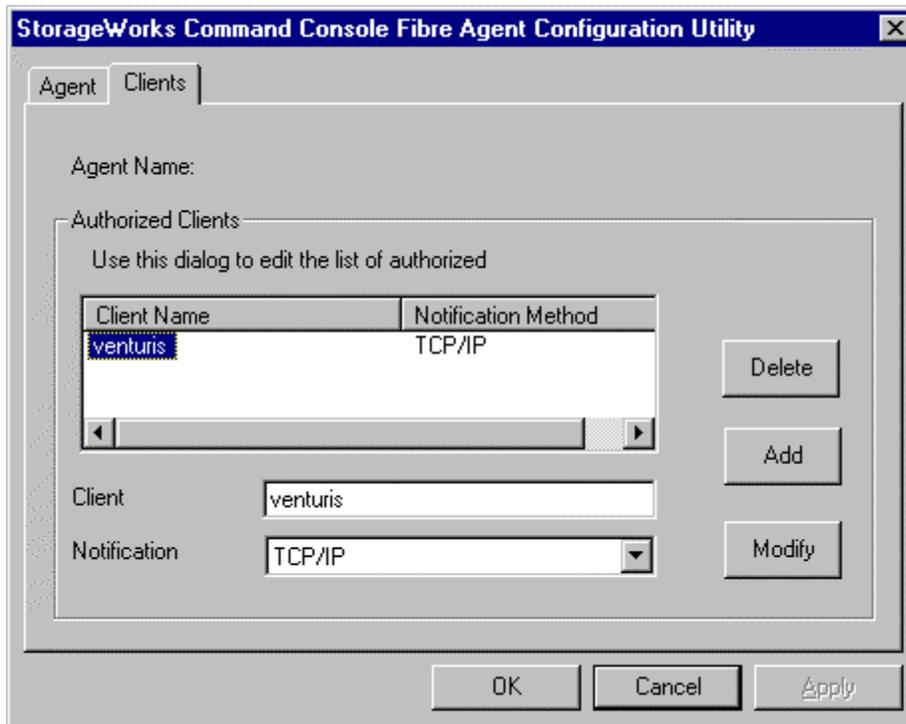
1. Type the Client's name.

2. Select the *Notification Scheme*: TCP/IP and/or SNMP or none.

3. For a definition of the notification schemes, please read the section, "Client Notification Options," located at the beginning of this chapter.

4. Click *Add Client* to add the Client system entry to the Client list.

5. Click *OK* or *Apply* to confirm your addition. If you click *OK*, you will leave the configuration program after you are asked to restart the Agent. If you click *Apply*, you will stay in the configuration program after you are asked to restart the Agent.

6. The software asks you if you want to restart the Agent.

7. Click *Yes*. The Agent is restarted.

3.4.3 Modifying a Client System Entry

This section contains instructions on how to modify a Client system entry on
Windows NT.

1. Click the Fibre Agent Configure entry in *Start/Programs/StorageWorks*. The
    StorageWorks Command Console Fibre Agent Configuration window appears.

2. Click the *Clients* tab. The Clients window appears.

3. Select the Client that you want to modify in the Client list.

4. Change the notification scheme: TCP/IP and/or SNMP or none.

5. For a definition of the notification schemes, please read the section, "Client
    Notification Options," located at the beginning of this chapter.

6. Click *Modify Client*.

7. Click *OK* or *Apply* to confirm your changes. If you click *OK*, you will leave the
    configuration program after you are asked to restart the Agent. If you click *Apply*,
    you will stay in the configuration program after you are asked to restart the Agent.

8. The software asks you if you want to restart the Agent.

9. Click *Yes*. The Agent is restarted.

3.4.4 Deleting a Client System Entry

This section contains instructions on how to delete a Client system entry on
Windows NT.

1. Click the Fibre Agent Configure entry in *Start/Programs/StorageWorks*. The
    StorageWorks Command Console Fibre Agent Configuration window appears.

2. Click the *Clients* tab. The Clients window appears.

3. Select the Client system entry to delete in the list.

4. Click *Delete Client*. The Client system entry is deleted.

5. Click *OK* or *Apply* to confirm your deletion. If you click *OK*, you will leave the
    configuration program after you are asked to restart the Agent. If you click *Apply*,
    you will stay in the configuration program after you are asked to restart the Agent.

6. The software asks you if you want to restart the Agent.

7. Click *Yes* to restart the Agent.

3.4.5 Changing the Polling Interval

This section contains instructions on how to change the polling interval of the Agent.

1. Click the Fibre Agent Configure entry in *Start*/*Programs*/*StorageWorks*. The StorageWorks Command Console Fibre Agent Configuration window appears (Figure 3–3).

2. Type an interval to poll each element in the fabric. The default is 5 minutes; however, you can enter an interval from 1 to 60 minutes.

3. Click *OK* to confirm your choice. You are asked to restart the Agent

4. Click *Yes* to restart the Agent.

Figure 3–3 SWCC Fibre Agent Configuration Window
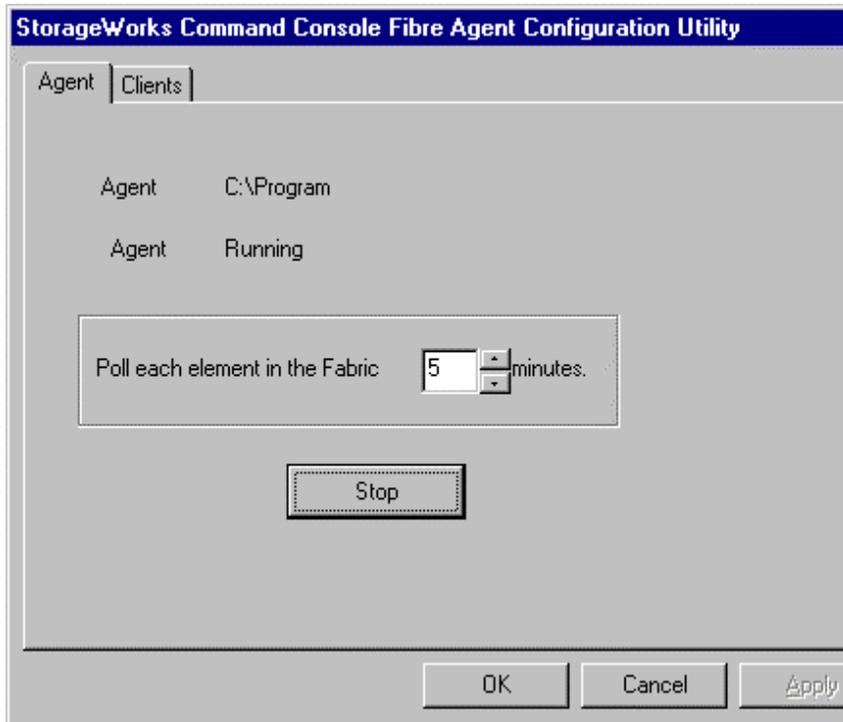
3.4.6 Stopping and Starting the Fibre Channel Switch Agent

This section contains instructions on how to stop and then start the Fibre Channel Switch Agent:

1. Click the Fibre Agent Configure entry in *Start/Programs/StorageWorks*. The StorageWorks Command Console Fibre Agent Configuration window appears (Figure 3–3).

2. Click the *Agent* tab. The Agent window appears. You are given the Agent's location and its status.

3. Click *Stop Agent* to stop the Agent or click *Start Agent* to start the Agent.

3.5 Connecting to the Fibre Channel Switch using the SWCC Client

1. From the START menu, select Command Console, then StorageWorks Command Console (reference Figure 3–4).

Figure 3–4 Launching the StorageWorks Command Console Navigation Window
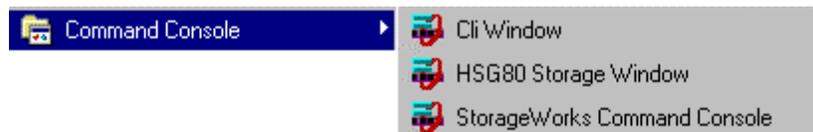
2. At the Command Console Navigation Window, use the *File* pull-down menu to select the *Add System* option, as shown in Figure 3–5.

Figure 3–5 Selecting the *Add System* Option

3. In the *Add System* Window (Figure 3–6), enter the name or the TCP/IP address of the Client host system on which you are running the software, then click *Apply*.

NOTE

The Client host name or address entered here must be the same as the host name added to the agent, as described in Section 3.3.2, *Adding a Client System Entry.*

Figure 3–6 The *Add System* Window
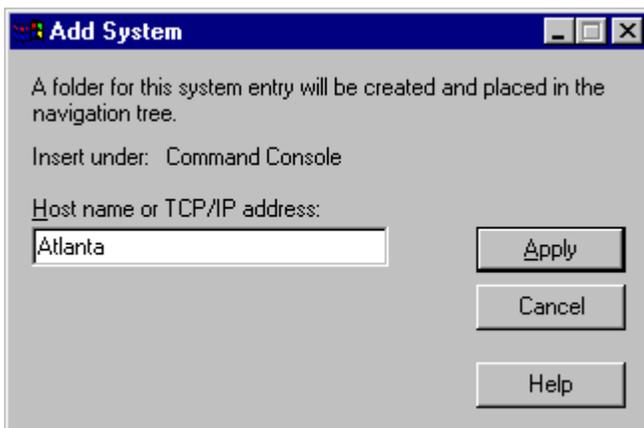
The Client name (or TCP/IP address) entered appears as a folder in the Command Console Navigation Window, indicating that the client is communicating with the Agent across the network.

4. Click on the + icon next to the Client host system folder. The Fibre Channel Network is displayed, along with the Fabric Window icon, as shown in Figure 3 –7.

Figure 3–7 "Atlanta" Client Host Added (Folder Open)
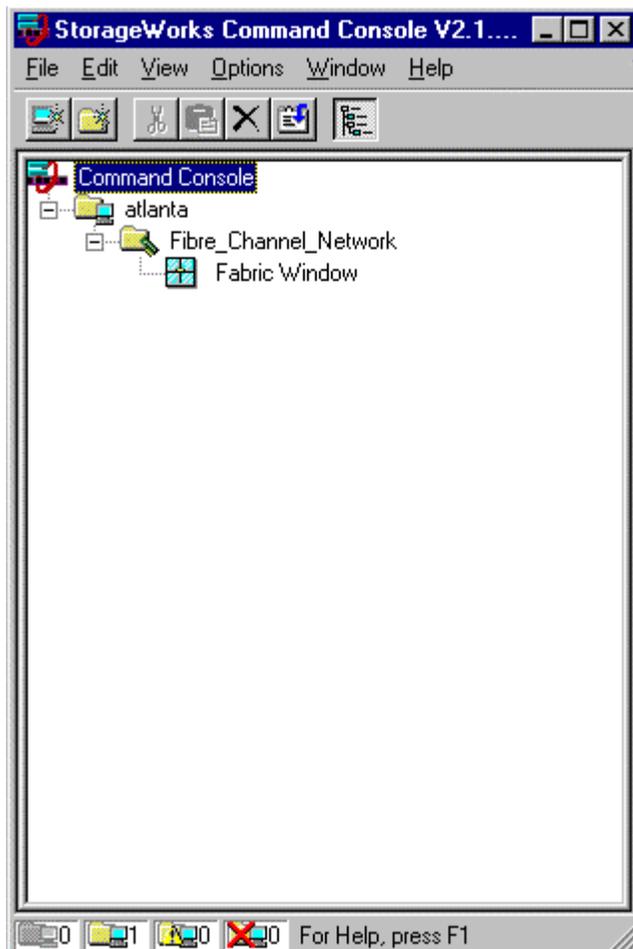
5. Double-click on the Fabric Window icon to launch the Fibre Channel Switch Client.

The Add Fabric Window will appear (Figure 3 –8).

Figure 3–8 Add Fabric Window
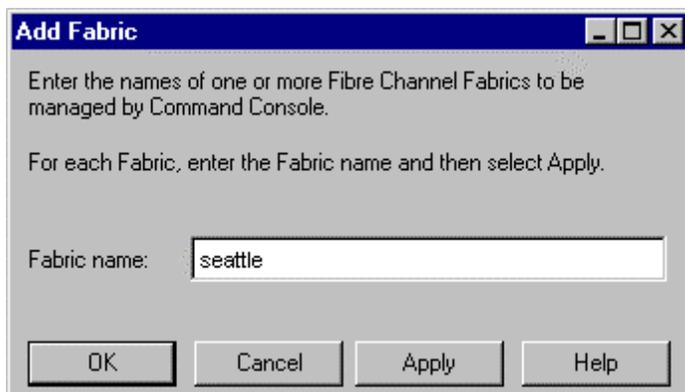
6. Enter a name for the Fabric Network, (*Example: seattle*), and click *Apply*.

The Fabric Window for your Client host system will appear, as shown in Figure 3 –9.

Figure 3–9 Fabric Window for Client Host, Fibre Channel Network
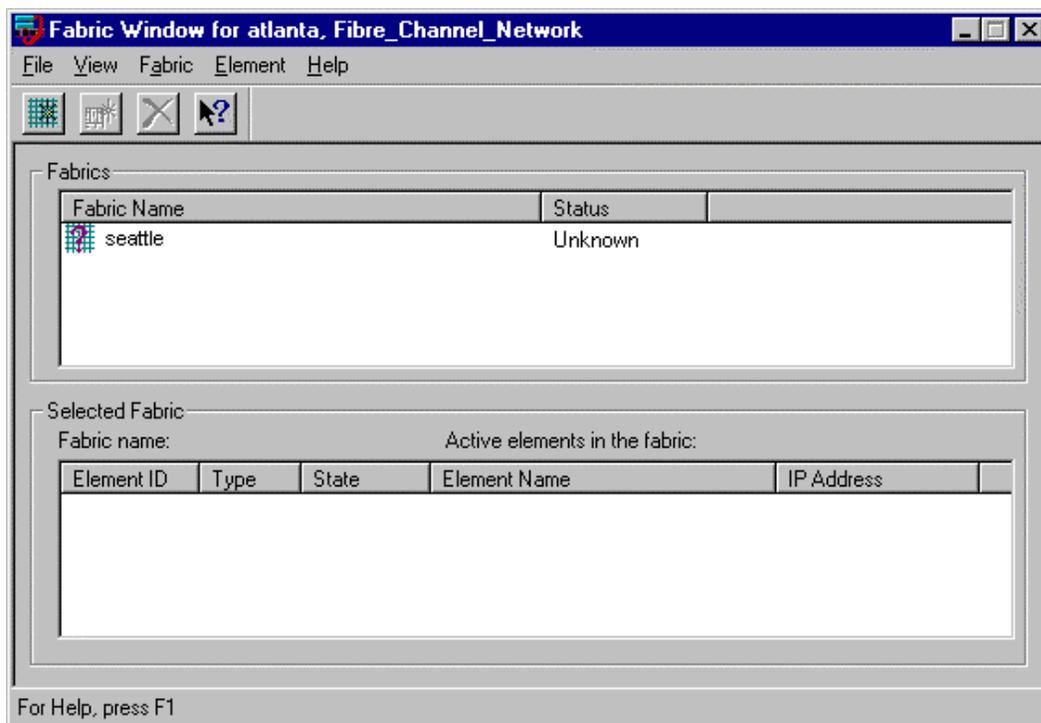
7.   In the Fabrics box of the Fabric Window, click on the Fabric icon (*Example:* seattle).

The Add Element box will appear, as shown in Figure 3 –10.

Figure 3–10 Adding a Fibre Channel Element to the Fabric
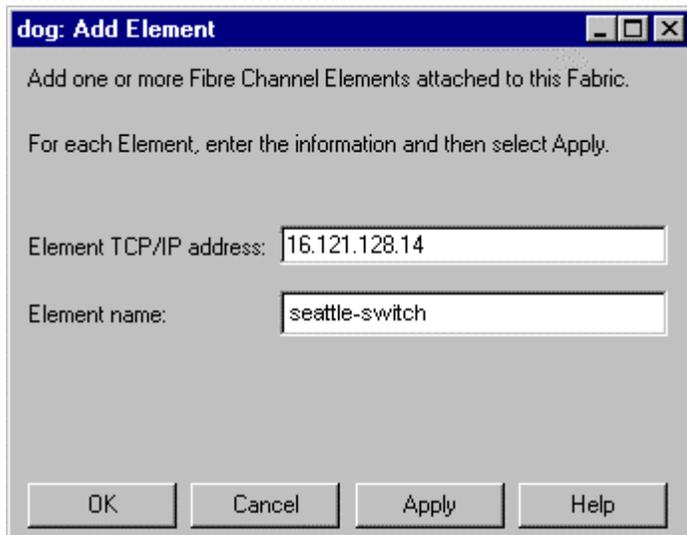
8. In the first field, enter the TCP/IP address of your Fibre Channel Switch. (This must be the same address entered during Initial FC Switch Configuration. Refer to your specific platform *Installation Reference Guide* for this information.). Tab to the second field and enter a name for your switch, (*Example:* seattle-switch), then click *Apply*.

The switch address and name entered will appear in the Fabric Window, as shown in Figure 3–11.

Figure 3–11 Fabric Window shows Switch Name and Address
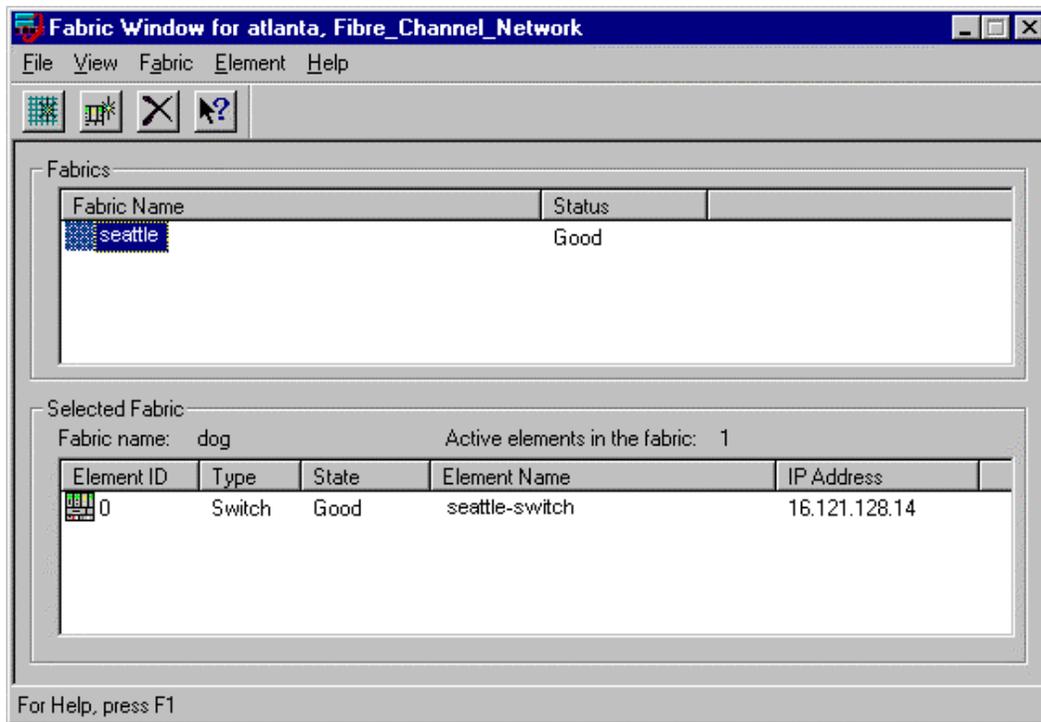
9. To launch the Switch Management Application, double -click on the Element ID icon (located in the Selected Fabric half of the Fabric Window).

When the Fibre Channel Switch image appears (Figure 3 –12), it indicates that you are now ready to configure and monitor the switch, as described in your *Fibre Channel Switch User's Guide*.

Figure 3–12 Switch Manangement Application Window
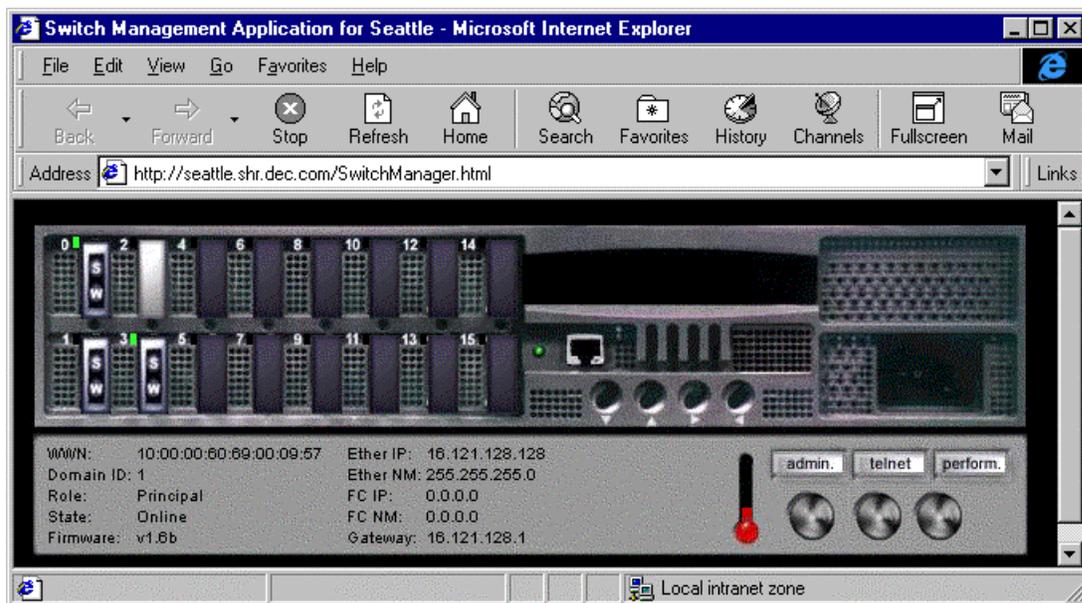
3.6 Basic Fibre Channel Switch Management

3.6.1 Admin., Telnet, and Performance Button

Admin. Button

This button allows you to change switch passwords and download fimware. Double -click on this button and a "Network Password Dialogue Box" appears (Figure 3 –14). Enter your user name and password.

Figure 3–13 Network Dialogue Password Box

Click *OK* and a System Admin. Window appears for the switch. In the example given in Figure 3–14, the switch is *seattle*.

Figure 3–14 System Admin. Window for Seattle
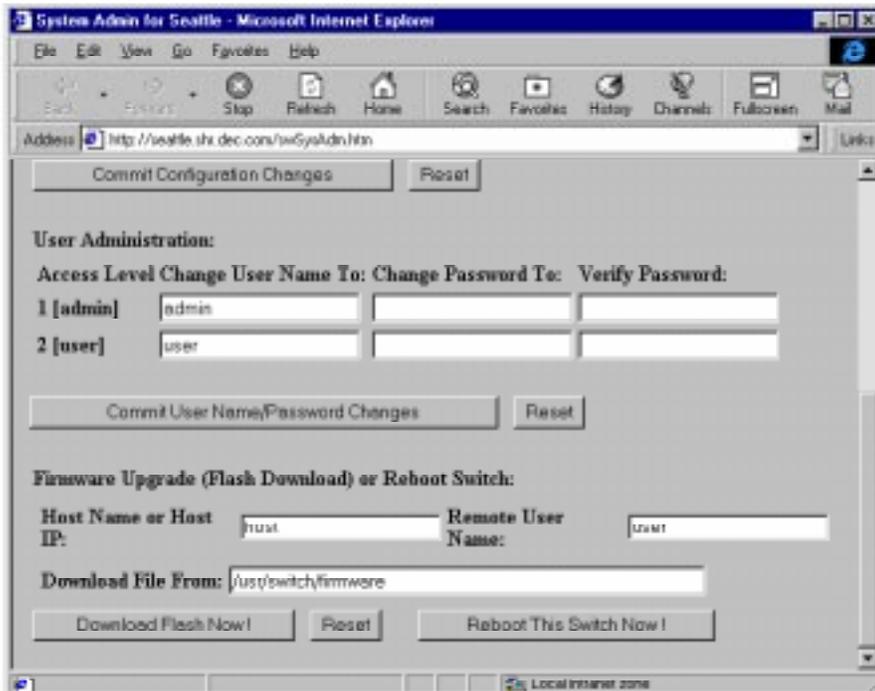
Telnet Button

This button allows you to access the Command Line Interpreter (CLI) by double -clicking on the button labeled *Telnet* in the window. For more information about using CLI, see Chapter 2, Section 2.15 of this user's guide.

Figure 3–15 Telnet Button Window for Seattle

```
Telnet - seattle.shr.dec.com                                    _ □ ✕

Connect  Edit  Terminal  Help

agtcfgSet              Set SNMP agent configuration
agtcfgShow             Print SNMP agent configuration
aliasShow              Print Alias Server information
date                   Print/set the system date and time
dateShow               Print system and firmware dates
errDisplayFilter       Print/set level of error filtering
errDump                Print error log (no page breaks)
errShow                Print error log
fabricShow             Print fabric membership info
fastboot               Reboot this switch, bypassing POST
flashDefault           Reset config to factory default
flashDownload          Download firmware into flash
flashSet               Set switch config parameters
flashShow              Print switch config parameters
h                      Print shell history
help                   Print this list
i                      Print task summary
ifShow                 Print network interface information
ipAddrSet              Set ethernet and FC IP addresses
ipAddrShow             Print ethernet and FC IP addresses
login                  Login as a new user
logout                 Logout from remote session

Type <CR> to continue, Q<CR> to stop: _
```

Performance Button

The Performance button shows you the performance of each individual port of a particular switch. Double-click on the *Performance* button on the switch in the Switch Management Application Window for the Performance Window to open.

Figure 3–16 Performance Button Window for Seattle

4
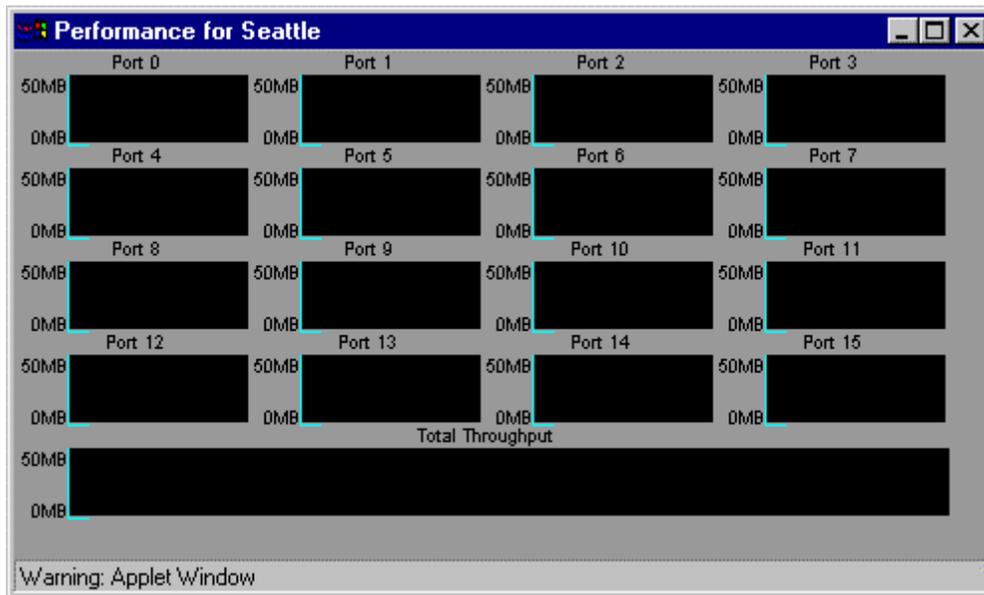
Command Console Client's Windows, Menus, and Tools

This Chapter describes the Command Console's Client Windows, Menus, and Tools.

4.1 Introduction

Command Console is a Windows-style graphical user interface, that uses standard Windows navigation, command selection and configuration features. Folders are used to arrange storage subsystems and a wizard creates virtual disks. Options and commands are selected using drop-down menus, tabs, dialog boxes, radio buttons, and check boxes. Windows can be sized and moved or minimized. Command Console Client has two main windows, the Navigation Window and the Storage Window, shown in Figure 4 –1.

Navigation Window: A navigation tool for adding host computers and storage subsystems to the network and for management of storage subsystems over a network.

CLI Window: A text-based utility for configuring storage subsystems.

Storage Window: Every storage subsystem has its own Storage Window. The Storage Window provides detailed information on your devices, controller(s) and virtual disks. This window lets you configure your controller(s) and create and manage virtual disks.

Figure 4–1 Command Console's Main Windows
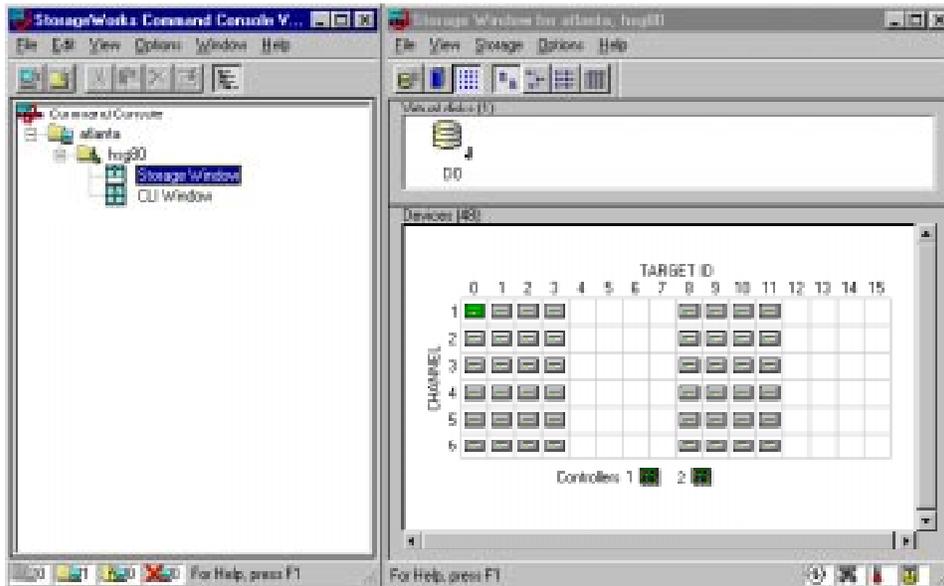
Drop-down Menus

Both the Navigation Tree and Storage Window have menu bars and drop-down menus. Each drop-down menu presents a group of related functions. For example, file management functions are located in the File menu.

Menu and Command Access Methods

You can use the mouse or the keyboard to access menus and commands. For keyboard access press the Alt key and then the underlined letter in the menu's name followed by the underlined letter in the command's name.

Shortcut Menus

Instead of using the standard menus to find the command you need, use the right mouse button to click a system, virtual disk, device, or folder. The menu that appears shows the most frequently used commands for that item.

4.2 Navigation Window

The Navigation Window (Figure 4–2) provides a list of all the host computers and storage subsystems to which Client is connected. The Navigation Window is made up of host system folders, storage subsystem and general folders. It has its own Menu Bar, Toolbar, and a Status Bar.

You can use the Navigation Window to see all the systems and storage managed by Command Console. This is especially useful for monitoring storage by watching the Navigation Window for problems or failures. The host systems and subsystems are displayed in a "tree."

Command Console will monitor your network connection and storage subsystem and report status by changing the icons in the Navigation Window.

Figure 4–2 RAID Array 8000/ESA12000 Navigation Window

Menu Bar

Tool Bar

Host System

Storage Subsystem

Storage Window Icon

Status Bar
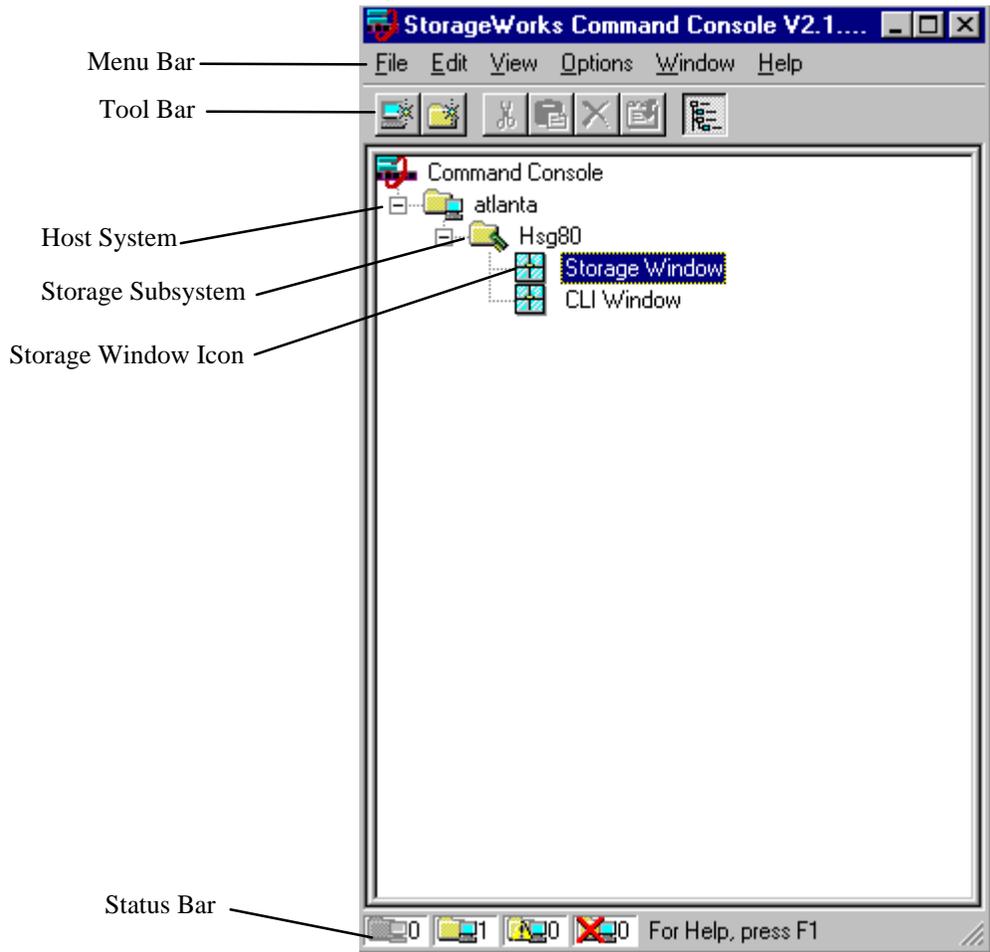
The Navigation Window uses the following icons:

Icon Name	Icon	Function/Description
System Folder		Shows the node of the host system containing controllers, virtual disks, and devices.
General Folder		This is a folder that you create. A general folder may be used to group other folders.
Controller Folder		Shows the Storage Window associated with a controller.
Storage Window		Double-clicking on the a Storage Window icon in the Navigation Window launches a Storage Window. The Storage Window shows all the virtual disks and devices associated with a storage subsystem.

System Folders

A System Folder shows the node containing storage. The following table shows the four states that System Folders can have in the Navigation Tree:

	Connection to the system has been established but there is a critical condition.
	Connection to the system is working properly.
	A system warning, subsystem may not be fully functional or in rebuild state.
	A grayed out folder means no connection to the storage subsystem.

General Folders

A General Folder is a container for grouping folders or systems. You can create general folders with the Add Folder command.

All controllers, virtual disks, and devices inside this folder are operating properly.

A system warning, subsystem within the folder may not be fully functional or in rebuild state.

A grayed out folder indicates no connection to the subsystem.

A controller, virtual disk, or device inside this folder may have failed.

Controller Folders

A Controller Folder shows all the storage associated with a controller.

All virtual disks and devices contained in this folder are working properly.

A virtual disk or device contained in this folder may have failed.

A grayed out folder indicates no connection to the controller.

A virtual disk or device contained in this folder has failed.

4.3   Navigation Window Menu Bar

The Menu Bar (Figure 4–3) is at the top of the Navigation Window and provides drop-down menu selections for Navigation Window tasks.

Figure 4–3 Menu Bar

The Menu Bar contains the following drop-down menus:

- **File Menu**
- **Edit Menu**
- **View Menu**
- **Options Menu**
- **Window Menu**
- **Help Menu**

These menus and their commands are described in the online help.

4.4 Navigation Window Toolbar

The Navigation Window Toolbar (Figure 4–4) is located directly under the Menu Bar and provides buttons to activate some of the more often-used Command Console functions. To display or hide the toolbar, you can toggle the Toolbar option on the View Menu.

Figure 4–4 Toolbar

Navigation Window Toolbar Buttons

Frequently used Navigation Window functions can be accessed by clicking a button on the Navigation Tree Toolbar. Toolbar buttons are:

Add System

Add New Folder

View system properties

Cut folder

Paste folder

Delete folder

View Navigation Tree

4.5 Navigation Window's Status Bar

The Status Bar (Figure 4–5) provides information about the number of systems currently connected to the Command Console client and the state of those connections.

The Status Bar provides colored status indicators to indicate the current state of system connections. Red indicates a failure and yellow a warning (possible failure).

Figure 4–5 Status Bar
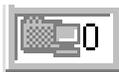
  A grayed out icon indicates the number of disconnected systems. In this example 0 systems are disconnected from Client.
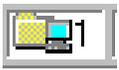
  A full-color icon indicates the number of fully functional "good" systems. In this example 1 system is good and online.

  A yellow triangle with an exclamation mark indicates the number of degraded systems. In this example 0 systems are degraded.

  A red X indicates a critically failed system. In this example there are 0 failed systems.

4.6 Storage Window

The Storage Window (Figure 4–6) provides detailed information on a particular storage subsystem that you have added to the Navigation Tree. You can open a separate Storage Window for each subsystem you have added in the Navigation Window.

The Storage Window has a Menu Bar, Toolbar, Virtual Disk Window, Device Window and Status Bar.

Figure 4–6 Multiple Storage Windows Opened from the Navigation Window

Three storage windows are opened:: one for host system sgiorg2000, one for sgiorg200, and one for atlanta. The storage window for atlanta is active.
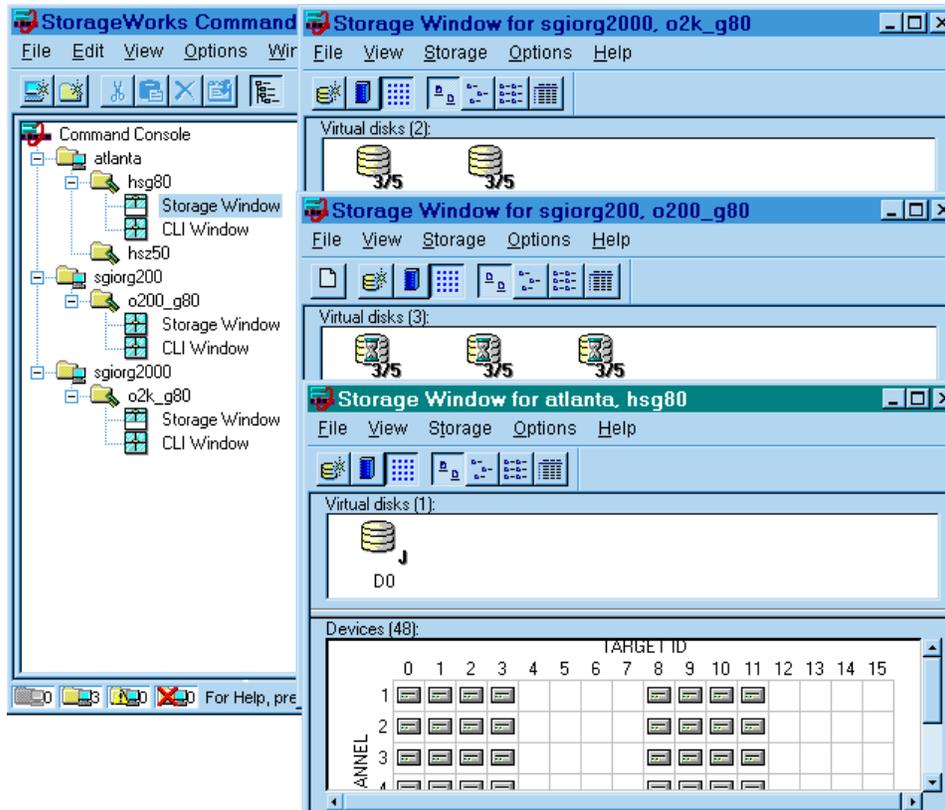
4.7   Storage Window Menu Bar

The Menu Bar (Figure 4–7) provides drop-down menu selections for the functions you can perform in the Storage Window.

Figure 4–7 Storage Window Menu Bar
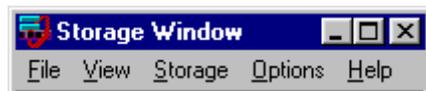
There are five, drop-down menus in the Menu Bar:

- **File Menu**
- **View Menu**
- **Storage Menu**
- **Options Menu**
- **Help Menu**

These menus and their commands are described in the online help.

4.8   Storage Window Toolbar

The Toolbar (Figure 4–8) provides buttons with which you can perform many of the functions available in the Storage Window. You can alternately display or hide the toolbar by clicking the Toolbar entry in the View menu.

Figure 4–8 Storage Window Toolbar

Icon Name	Icon	Function/Description
New Connection		Enables you to connect to a new storage subsystem.
Add Virtual Disk button		Enables you to add a new virtual disk to your storage subsystem.
Cabinet button		Displays your devices in a pictorial view of your cabinet.
Grid button		Displays your devices as a logical array of icons.
Large Icons button		Displays your virtual disks as large icons.
Small Icons button		Displays your virtual disks as small icons.
List button		Displays your virtual disks as an array of small icons.
Details button		Displays your virtual disks as a detailed list.

4.9   Virtual Disk Window

The Virtual Disk Window (Figure 4–9) provides detailed information about the virtual disks you have configured in your storage subsystem.

Figure 4–9 Virtual Disk Window

View Options

You have four options, shown in Figure 4–10, for viewing the information in the Virtual Disk Window. You can select any one of these options in the *Virtual Disk* selection in the View menu.

Figure 4–10 Virtual Disk Window View Choices
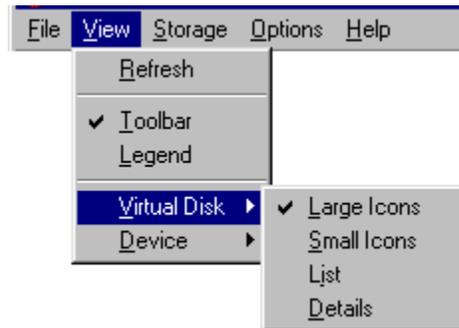
Large Icons, Small Icons, and List Views

If you choose the *Large Icons, Small Icons,* or *List* options in the *Virtual Disk* option in the View menu, the program uses icons to indicate the underlying structure (that is, the type) of the virtual disks in the virtual disk window, as follows:

Indicates a JBOD (that is, a single-disk) virtual disk.

Indicates a RAID 0+1virtual disk.

Indicates a RAID 0 virtual disk.

Indicates a RAID 5 virtual disk.

Indicates a RAID 1 virtual disk.

The program also uses icons to indicate the status of the virtual disks in the virtual disk window, as follows:

Indicates that the virtual disk is completely operational

Indicates that the controller is in an initializing state.

Indicates that the virtual disk is not operational

Indicates that the virtual disk is in the process of regenerating the data on one of its members

Indicates that the virtual disk is reduced

Details View

If you choose the *Details* option (Figure 4–11) in the View menu, the Virtual Disk
Window displays four information fields:

Figure 4–11 Virtual Disk Details View

Virtual disks (3):				
Name	State	Capacity	RAID Level	
U00	Good	3.995 GB	Single Disk (JBOD)	
U01	Good	1.960 GB	Single Disk (JBOD)	
U02	Good	3.910 GB	Striped Parity (RAID3/5)	

Name: The name Client has given your virtual disk.

State: The current operational state of the virtual disk. There are five possible states:
- Good – The virtual disk is completely operational.
- Failed – The virtual disk is completely non-functional.
- Initializing – The controller is initializing any data on the virtual disk and is writing
  its own file structure.
- Reduced – The virtual disk has one failed member.
- Reconstructing – The virtual disk is in the process of regenerating the data on one of
  its members.

Capacity: The maximum amount of data, in bytes, that the virtual disk can contain.

RAID level: The underlying structure of the virtual disk. There are five, possible levels:

- Single disk – The virtual disk consists of a single disk drive.
- RAID 0 – The virtual disk is a striped disk.
- RAID 1 – The virtual disk is a mirrored disk.
- RAID 0/1 – The virtual disk is a striped mirrored disk.
- RAID 5 – The virtual disk is a striped disk with a floating parity disk.

NOTE

You can change the width of an information field by
dragging on one of its borders.

Virtual Disk Properties

You can view virtual disk properties by accessing their property sheets. You can access a virtual disk's property sheets directly by double-clicking the virtual disk's icon in any view. You can also right-click the icon and select *Properties*.

Locating Virtual Disk Members

You can locate the members of any virtual disk by clicking on its icon in the Virtual Disk Window. Its members are shown as shaded device icons in the Device Window.

4.10 Device Window

The Device Window (Figure 4–12) provides detailed information about the devices you have configured in your storage subsystem. You can display a menu of device icons and their meanings by clicking *Legends* on the View Menu.

Figure 4–12 Device Window (Grid View Shown)
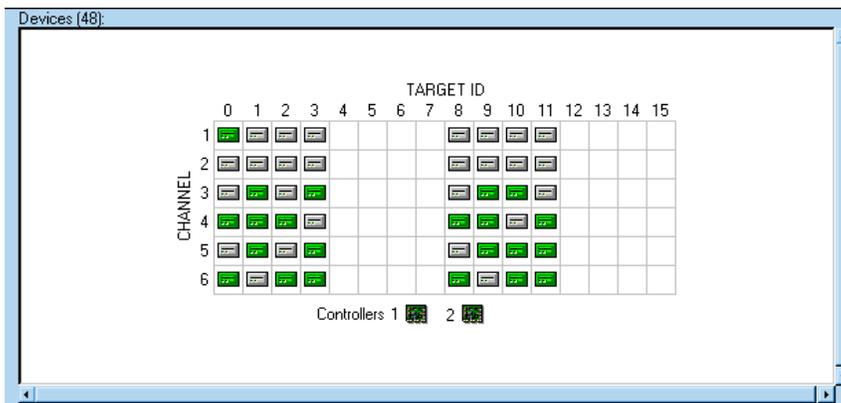
Device Properties

You can view device properties by accessing their property sheets. You can access a device's property sheets directly by double-clicking on it. You can also right-click on a device icon and select *Properties*. There are three device property sheets: *General*, *Membership* and *Support Information*.

The *General* property sheet (Figure 4–13) provides general device information such as target ID, capacity and transfer rate.

Figure 4–13 General Property Sheet
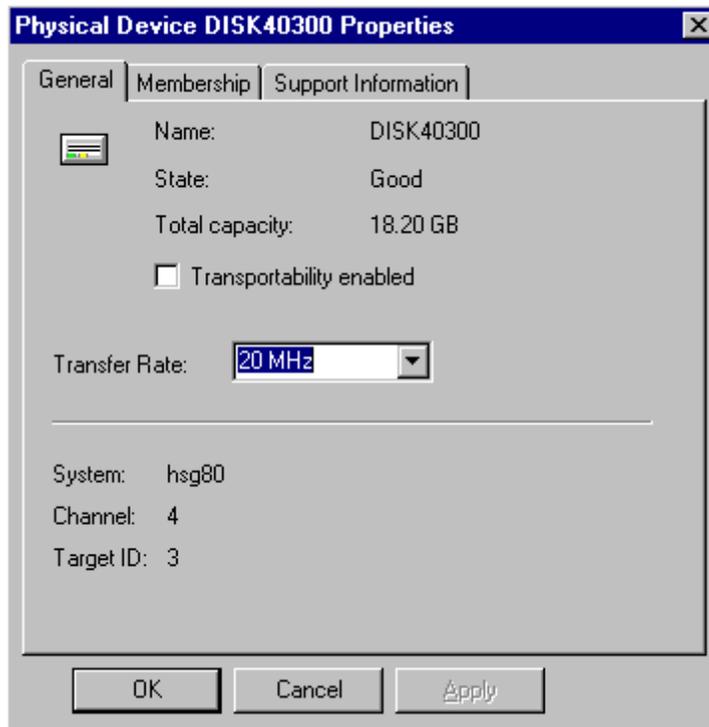
The *Membership* sheet (Figure 4–14) shows which virtual disk the device is part of. In this example the device belongs to virtual disk U02 which is currently operating in a reduced state.

Figure 4–14 Membership Property Sheet
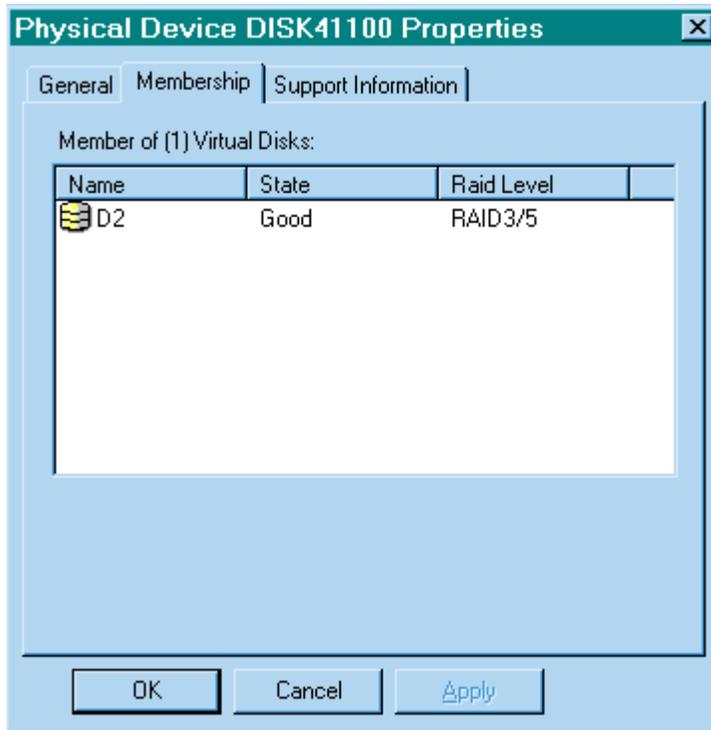
The *Support Information* sheet (Figure 4–15) identifies the device by name and manufacturer. Its firmware revision is also provided.

Figure 4–15 Support Information Property Sheet
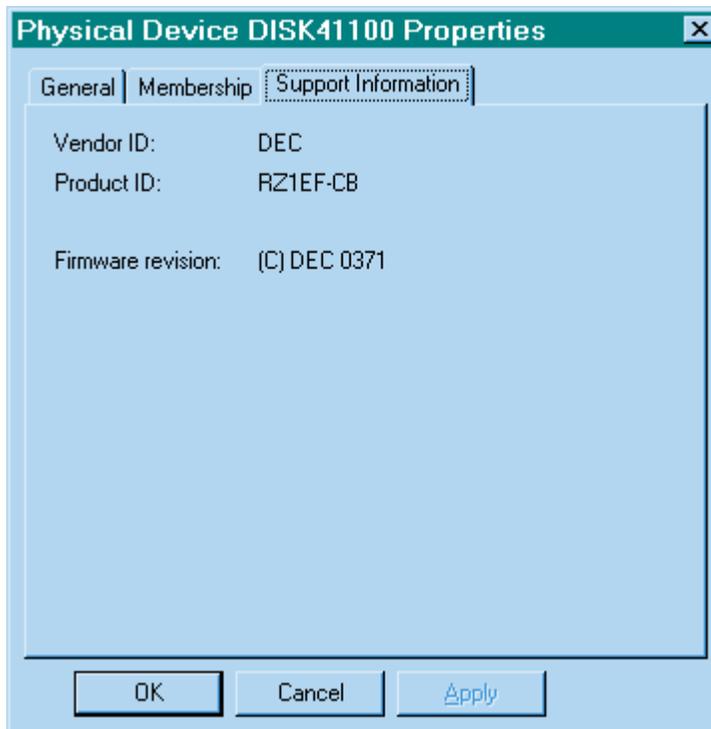
View Options

You have two options: a grid view or a cabinet view for displaying the devices in your storage subsystem. You can select either of these options in the *Device* selection in the View Menu.

Grid View

If you choose the *Grid* option in the *Device* selection in the View Menu, the Device Window displays your storage subsystem as an array of icons. Your devices are arranged vertically in the grid by device bus and horizontally by target on the bus. The program displays an icon for each physical device that you have added to your configuration. A blank slot indicates that no device is installed or that you have not added an installed device.

The program uses icons to indicate the status of the disks in the disk window, as follows:

	Indicates an operational disk used in a virtual disk
	Indicates a disk that has been added to the configuration and is available for use in a virtual disk
	Indicates a disk that is a spare and is available for automatic replacement of a failed member in a RAID virtual disk
	Indicates a disk that is not operational
	Indicates a disk that is reconstructing

Your controllers are arranged just below the device array. The program uses icons to indicate the status of your controllers, as follows:

	Indicates an operational controller
	Indicates a controller that is not operational

Cabinet View

If you choose the *Cabinet* option (Figure 4–16) in the *Device* selection in the View Menu, the program displays a pictorial representation of your cabinet.

Figure 4–16 Cabinet View
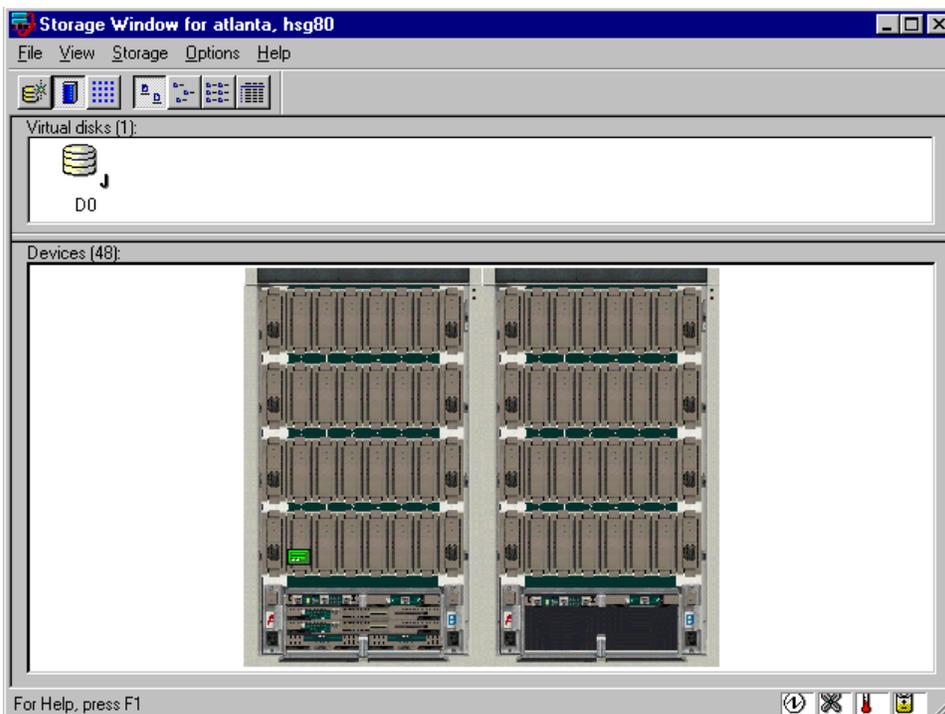
Disk and Controller Properties

You can view very detailed attributes of your disks and controllers by accessing their property sheets. You can access the property sheets directly by double-clicking the object's icon in either view. You can also right-click the icon and select Properties.

Locating Virtual Disk Members

You can locate the members of any virtual disk by clicking on its icon in the Virtual Disk Window. Its members are shown as shaded device icons in the Device Window.

4.11 EMU Icon

The EMU icon (Figure 4-17) provides environmental status of the pedestal and UPS. The EMU icon is located in the lower right corner of the status window. Double-click on the icon to view the environmental properties.

Figure 4–17 Storage Window Status Bar
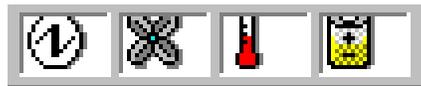
Indicates that the cabinet fans are all operating normally

Indicates a problem with the cabinet fans

Indicates that cabinet temperature is within its acceptable range

Indicates that cabinet temperature is out of its acceptable range

Indicates that all power supplies are operating normally

Indicates a problem with one or more power supplies

Indicates that the cache backup batteries are fully charged and operating normally

Indicates a low charge condition on the cache backup batteries

Indicates that the cache backup batteries have failed or are discharged

4.12 CLI Window

The Command Line Interpreter (CLI) is a command-line user interface for monitoring and configuring your storage controller. Because the configuration of the storage subsystem can be changed using CLI commands, this feature is password-protected. If you are familiar with CLI syntax, you can enter commands into the command line area directly underneath the Menu Bar. Error messages and system responses are displayed in the message area directly beneath the command line area.

Figure 4–18 CLI Window
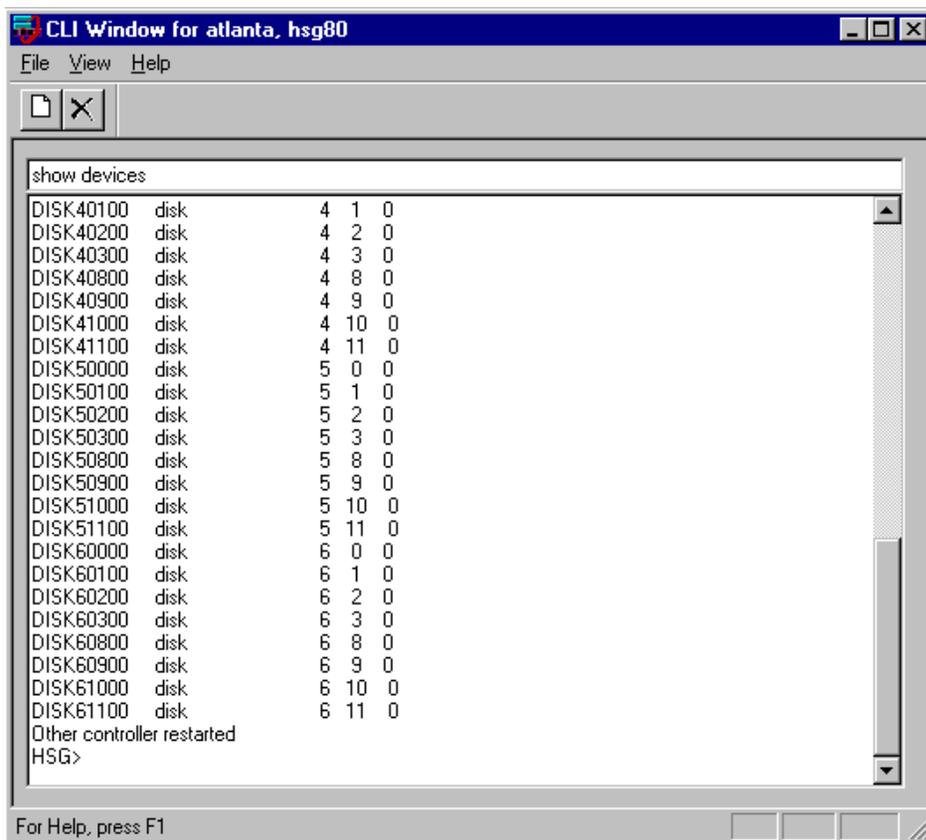
4.13 Virtual Disk Wizard

Command Console Client feature a powerful Virtual Disk Wizard to create virtual disks. The Virtual Disk Wizard guides you through the process of adding a virtual disk to your configuration. The Add Virtual Disk Wizard is opened by selecting *Add Virtual Disk* from the Storage menu. The first wizard screen (Figure 4 –19) prompts you to select the type of virtual disk you want to create. Click on one of the radio buttons, then click the *Next* button to begin the process.

Figure 4–19 Opening Virtual Disk Window, Step 1
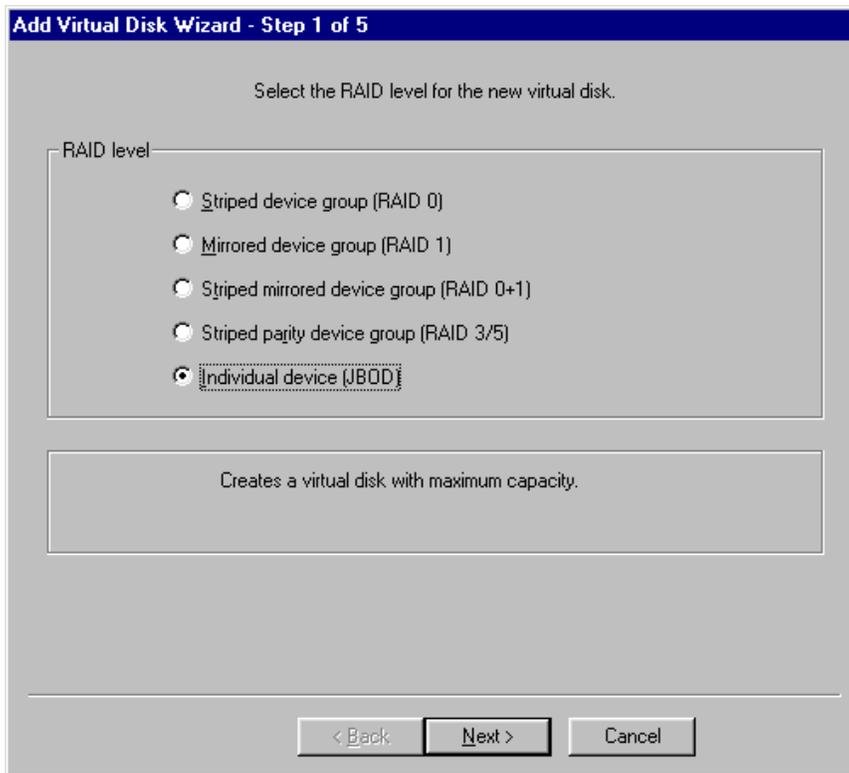
There are five wizard steps:

Step 1: Select the type of virtual disk you want to create.

Step 2: Select the devices you want to include in your virtual disk.

Step 3: Specify the capacity of the virtual disk.

Step 4. Specify the target ID, LUN and other characteristics for your virtual disk.

Step 5: Confirm the selections you have made and create the virtual disk.

5

Maintenance and Troubleshooting

This chapter provides system backup recommendations, instructions on how to upgrade your controller's software, and troubleshooting tips.

5.1   Backing up Files

Backing up your Configuration File

You should keep a backup copy of your current configuration file to facilitate restoration if the storage subsystem configuration file is corrupted, if you want to restore a previous configuration setting, or if you need to reconfigure a failed controller.

Backing up Data Files

You should consistently back up your data to ensure its integrity. Make sure you have a consistent and comprehensive subsystem backup policy in place while using Command Console. If you do not maintain consistent backups of your data, you put it at risk.

Following good subsystem management practice, you should back up the data on your disk drives to aid in recovery if there is a primary system failure, if the integrity of a storage device becomes questionable, or if you experience gross user error.

Client does not automatically back up your data for you. Use the appropriate host backup utilities for this purpose.

5.2   Updating Controller Firmware

See the release notes for your version of Client for information and restrictions on downloading various versions of controller software. You can upgrade your firmware using a network or a SCSI connection. To update your controller's firmware, proceed as follows:

1. From the Storage menu (Figure 5–1), select *Update Controller Software*.

Figure 5–1 Update Firmware Command
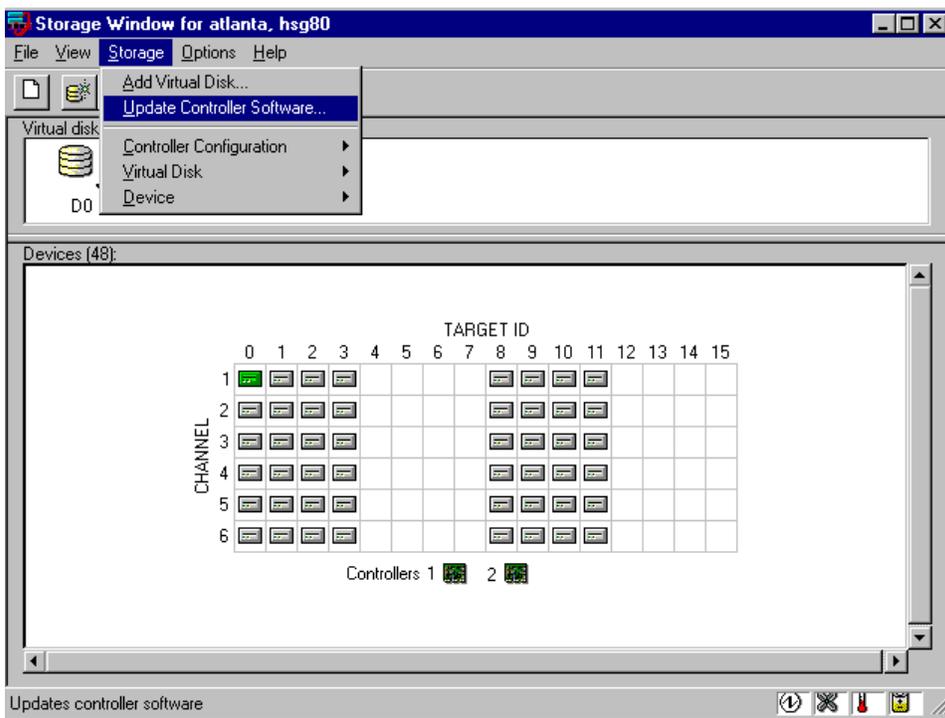
2. A window will appear (Figure 5–2), asking you to specify the firmware file that you want to load. You can easily identify Firmware software by its .fdi extension. Enter the fdi file name , then click Start Update. Upon completion of the Firmware update, the system will automatically reboot and update the controller's firmware.

Figure 5–2 Firmware Update Dialog Box
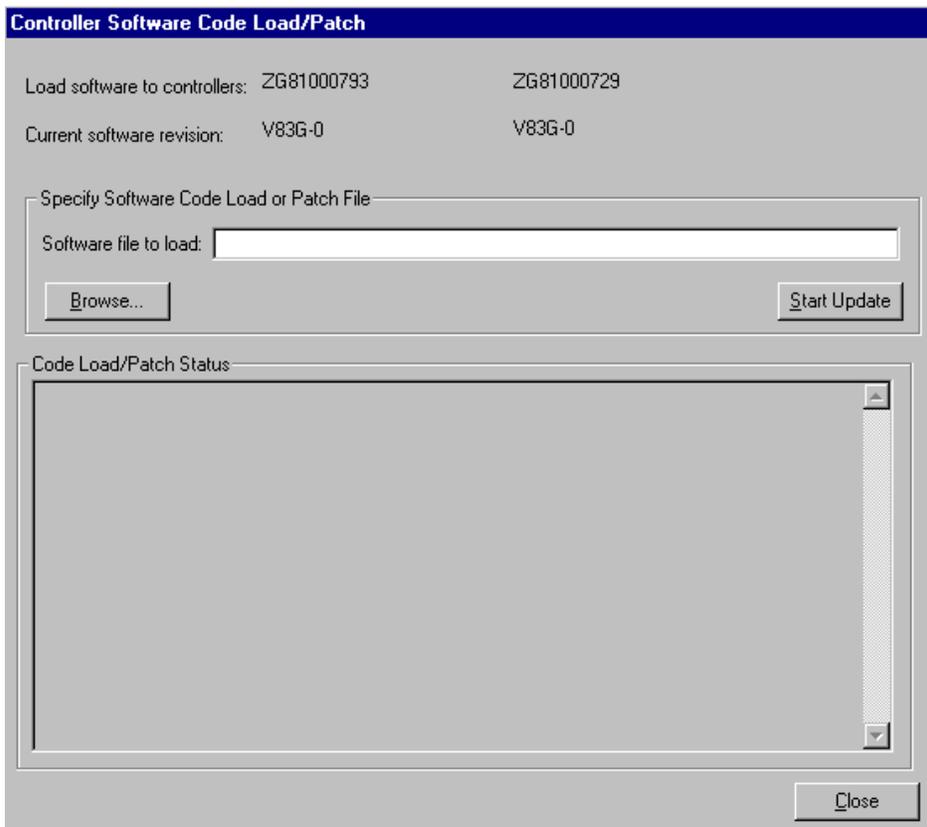
5.3   SNMP Trap Information

When a physical storage device, a logical storage unit, or a physical enclosure component changes state, Agent sends out an SNMP trap (that is, an asynchronous event notification) to any hosts that are listed in the *client.ini* file with this option enabled. (See "Editing the Client Access File" in this document for more information.)

This section describes the format of the SNMP trap.

5.3.1 Information Fields

An Agent SNMP trap consists of the following information fields:

- Source IP address
- GENERIC = 6 (6 = ENTERPRISE-SPECIFIC VALUE TO FOLLOW)
- SPECIFIC = 0 or 1, where:
  - "0" indicates that the storage component transitioned from a bad to a good state
  - "1" indicates that the storage component transitioned from a good to a bad state
- OID Variable of the offending storage component
- OCTET STRING Value returning the name of the storage subsystem containing the storage component specified by the OID Variable

5.3.2 OID Variables

Each object ID (OID) is prefixed by a Base Enterprise Number (BEN) for the Command Console product. The default BEN for Command Console Agents is as follows:

> 1.3.6.1.4.1.36.2.15.21
> (iso.org.dod.internet.private.enterprises.dec.ema.sysobjects.raidmanager)

Following are the OID values used by Command Console:

- Disk Status = BEN + .3.2.1.4
  (BEN + .subsys.ssStatusTable.ssEntry.ssDiskStatus)
- Power Supply Status = BEN + .3.2.1.5
  (BEN + .subsys.ssStatusTable.ssEntry.ssPowerStatus)
- Cooling Status = BEN + .3.2.1.6
  (BEN + .subsys.ssStatusTable.ssEntry.ssFanStatus)
- Cache Battery Charge Status = BEN + .3.2.1.7
  (BEN + .subsys.ssStatusTable.ssEntry.ssCacheBatteryStatus)OverTemperature
  Status = BEN + .3.2.1.8
  (BEN + .subsys.ssStatusTable.ssEntry.ssTemperatureStatus)
- Communication Status = BEN + .3.2.1.9
  (BEN + .subsys.ssStatusTable.ssEntry.ssCommStatus)
- EMU External Input Status = BEN + .3.2.1.10
  (BEN + .subsys.ssStatusTable.ssEntry.ssEmuExtInputStatus)
- THIS Controller Status = BEN + .3.2.1.12
  (BEN + .subsys.ssStatusTable.ssEntry.ssController1Status)
- OTHER Controller Status = BEN + .3.2.1.19
  (BEN + .subsys.ssStatusTable.ssEntry.ssController2Status)
- LUN Status = BEN + .3.2.1.27
  (BEN + .subsys.ssStatusTable.ssEntry.ssLUNStatus)

If your SNMP Management Utility has the Agent MIB compiled, the utility should display the character string values in parentheses above. Note that some OID values specified in the MIB are not actually used for SNMP traps.

5.3.4 Octet String Values

The octet string value attached to a trap takes the form:

  subsystem_name controller_type ID:*state*

where:
subsystem_name = the name of the subsystem to which the trap applies
controller_type = the model number of the controller in the subsystem
ID = the identification of the storage component that caused the trap
     (Note that a "storage component" can be a physical storage device, a logical storage unit, or a
       physical enclosure component)
state = the status of the storage component that caused the trap

ID and State Values

The range of values in the ID and status value fields varies, depending upon the type of storage component causing a trap. Values for various component types are as follows:

- For disk drives:
  disk_name:*state,* where:

  > *disk_name* = the name of the disk drive in the subsystem (for example, "disk20100"
  > *state* = either of the following:
  > > 1 = drive went from bad to good
  > > 2 = drive went from good to bad

- For temperature sensors:
  cabinet_number:sensor_type:state, where:

  > *cabinet_number* = the cabinet ID, from 0–3, in a BA370 pedestal enclosure
  > *sensor_type* = EMU sensor, sensor 1, or sensor 2 in a BA370 pedestal enclosure
  > *state* = either of the following:
  > > 1 = temperature went from bad to good
  > > 2 = temperature went from good to bad

- For cabinet fans:
  cabinet_number:fan_position:state, where:

  > *cabinet_number* = the cabinet ID, from 0–3, in a BA370 pedestal enclosure
  > *fan_position* = fan location, from 1–8
  > *state* = either of the following:
  > > 1 = fan went from bad to good
  > > 2 = fan went from good to bad

- For power supplies:
  cabinet_number:pwr_position:state, where:

  > *cabinet_number* = the cabinet ID, from 0–3, in a BA370 pedestal enclosure
  > *pwr_position* = power supply location, from 1–8
  > *state* = any of the following:
  > > 1 = power supply went from bad to good
  > > 2 = supply went from good to bad
  > > 3 = power supply not present

- For batteries:
  controller_device_scsi_id:*state,* where:

  > *controller_device_scsi_id* = the SCSI ID of the reporting controller on the controller's device bus
  > *state* = any of the following:
  > > good = battery is good
  > > low = battery voltage is low
  > > fail = battery has failed

- For LUNs:
  virtual_disk_number:*state,* where:

  > *virtual_disk_number* = the unit number of the virtual disk
  > *state* = any of the following:
  > > 0 = good
  > > 1 = reduced
  > > 2 = reconstructing
  > > 4 = failed

5.3.5 Trap Example

An example SNMP trap for storage subsystem "subsys1" with a failing power supply on an Agent at IP address "16.82.16.01", would contain the following information:

> Source address = 16.82.16.01
> GENERIC = 6
> SPECIFIC = 1
> OID Value = 1.3.6.1.4.1.36.2.15.21.3.2.1.5
> (iso.org.dod.internet.private.enterprises.dec.ema.sysobjects.raidmanager.subsys.
> ssStatusTable.ssEntry.ssPowerStatus)
> OCTET STRING = subsys1 HSZ70 1:3:2

5.3.6 SNMP MIB Specification

This section presents the Command Console SNMP MIB specification, for reference only. This version of the specification applies to Version 2.0 of Command Console.

```
-- Digital Equipment Corporation SWCC Agent Private MIB
--
-- MIB for release 2.0 of StorageWorks Command Console Daemon Agent.
--
-- Copyright (c) 1996,1997 Digital Equipment Corporation.  All Rights Reserved.
-- Reproduction of this document is authorized on condition that the
-- foregoing notice is included.
--
-- This SNMP MIB Specification contains Digital SWCC's confidential
-- and proprietary intellectual property.  Digital retains all title and
-- ownership in the Specification, including any revisions.
--
-- It is Digital's intent to encourage the widespread use of this
-- Specification in connection with the management of storage subsystems
-- using the Command Console or Enterprise Storage Manager Graphical User Interface or
-- other SNMP management utilities.  Digital grants vendors, end-users,
-- and other interested parties a non-exclusive license to use this
-- Specification in connection with the management of Digital StorageWorks
-- RAID Controller products.
--
-- This Specification is supplied AS IS, and Digital makes no warranty,
-- either expressed or implied, as to the use, operation, condition, or
-- performance of the Specification.
--
--    StorageWorks Enterprise Array Manager (STEAM) Management Information Base (MIB)
--
-- iso(1) org(3) dod(6) internet(1) private(4) enterprises(1) dec(36)
-- ema(2) sysobjids(15) raidmanager(21)
```

```
    RAIDMANAGER-MIB DEFINITIONS ::= BEGIN

  IMPORTS
     enterprises
             FROM RFC1155-SMI
     OBJECT-TYPE
             FROM RFC-1212
     DisplayString
             FROM RFC-1213;


  dec          OBJECT IDENTIFIER ::= { enterprises 36 }
  ema          OBJECT IDENTIFIER ::= { dec 2 }
  sysobjids    OBJECT IDENTIFIER ::= { ema 15 }
  raidmanager    OBJECT IDENTIFIER ::= { sysobjids 21 }

-- raidmanager MIB
  agent         OBJECT IDENTIFIER ::= { raidmanager 1 }
  workstation    OBJECT IDENTIFIER ::= { raidmanager 2 }
  subsys        OBJECT IDENTIFIER ::= { raidmanager 3 }

--     Agent Information Group

  agManufacturer OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "The name of the Enterprise Storage Manager Agent manufacturer."

    ::= { agent 1 }

  agMajVersion OBJECT-TYPE
    SYNTAX  INTEGER
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "Enterprise Storage Manager Agent Major Version Number (e.g., 3
      for 3.0)."

    ::= { agent 2 }

  agMinVersion OBJECT-TYPE
    SYNTAX  INTEGER
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "Enterprise Storage Manager Agent Minor Version Number (e.g., 0
```

```
    for 3.0)."

   ::= { agent 3 }

agHostName OBJECT-TYPE
   SYNTAX  DisplayString
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "The Host System Network Name where the agent resides."

   ::= { agent 4 }


agEnterprise OBJECT-TYPE
   SYNTAX  OBJECT IDENTIFIER
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "The Enterprise ID subtree for Enterprise Storage Manager Agent MIB
      is registered."

   ::= { agent 5 }

agDescription OBJECT-TYPE
   SYNTAX  DisplayString
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "The Enterprise Array Storage Agent description."

   ::= { agent 6 }


-- Workstation Group

   wsCPU OBJECT-TYPE
      SYNTAX  DisplayString
      ACCESS  read-only
      STATUS  mandatory
      DESCRIPTION
        "The workstation CPU type (e.g., 80486)."

      ::= { workstation 1 }

   wsComputerType OBJECT-TYPE
      SYNTAX  DisplayString
```

```
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "The workstation Computer type (e.g., PC/AT)."

    ::= { workstation 2 }

wsModel OBJECT-TYPE
    SYNTAX  INTEGER
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "The workstation model number."

    ::= { workstation 3 }

wsSubModel OBJECT-TYPE
    SYNTAX  INTEGER
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "The workstation submodel number."

    ::= { workstation 4 }

wsBiosVersion OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "The workstation BIOS Version."

    ::= { workstation 5 }

wsOS OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "The workstation operating system name (e.g., WINNT)."

    ::= { workstation 6 }

wsOSMajVersion OBJECT-TYPE
    SYNTAX  INTEGER
    ACCESS  read-only
    STATUS  mandatory
```

```
        DESCRIPTION
          "The workstation OS major version number
          (e.g., 3 for WINNT 3.51)."

        ::= { workstation 7 }

     wsOSMinVersion OBJECT-TYPE
        SYNTAX  INTEGER
        ACCESS  read-only
        STATUS  mandatory
        DESCRIPTION
          "The workstation OS minor version number (e.g., 51 for
          WINNT 3.51)."

        ::= { workstation 8 }


-- Subsystem Group

     ssTotalSubSystems OBJECT-TYPE
        SYNTAX  INTEGER
        ACCESS  read-only
        STATUS  mandatory
        DESCRIPTION
          "The total number of subsystems presently serviced
           by this agent."

        ::= { subsys 1 }


-- Subsystem Status Group

     ssStatusTable OBJECT-TYPE
        SYNTAX  SEQUENCE OF SsEntry
        ACCESS  not-accessible
        STATUS  mandatory
        DESCRIPTION
          "This table holds the status information for
           each subsystem."

        ::= { subsys 2 }

     ssEntry OBJECT-TYPE
        SYNTAX  SsEntry
        ACCESS  not-accessible
        STATUS  mandatory
        DESCRIPTION
```

```
        "The subsystem information entry."
        INDEX   { ssEntryIndex }

        ::= { ssStatusTable 1 }

    SsEntry ::=
     SEQUENCE{
      ssEntryIndex
              INTEGER,
      ssSubsysName
              DisplayString,
      ssOverallStatus
              INTEGER,
      ssDiskStatus
              INTEGER,
      ssPowerStatus
              INTEGER,
      ssFanStatus
              INTEGER,
      ssCacheBatteryStatus
              INTEGER,
      ssTemperatureStatus
              INTEGER,
      ssCommStatus
              INTEGER,
      ssEmuExtInputStatus
              INTEGER,
      ssEmuPresent
              INTEGER,
      ssController1Status
              INTEGER,
      ssController1SerNum
              DisplayString,
      ssController1RDAC
              INTEGER,
      ssController1Type
              DisplayString,
      ssController1ProdID
              DisplayString,
      ssController1FwRev
              DisplayString,
      ssController1HwRev
              DisplayString,
      ssController2Status
              INTEGER,
      ssController2SerNum
              DisplayString,
```

```
    ssController2RDAC
            INTEGER,
    ssController2Type
            DisplayString,
    ssController2ProdID
            DisplayString,
    ssController2FwRev
            DisplayString,
    ssController2HwRev
            DisplayString,
    ssUpsStatus
            INTEGER,
    ssLunStatus
            INTEGER
   }

ssEntryIndex OBJECT-TYPE
   SYNTAX  INTEGER
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "The index into ssStatusTable ."

   ::= { ssEntry 1 }

ssSubsysName OBJECT-TYPE
   SYNTAX  DisplayString
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "The local Subsystem Name."

   ::= { ssEntry 2 }

ssOverallStatus OBJECT-TYPE
 SYNTAX  INTEGER  {
        GREEN(0),
        YELLOW(1),
        RED(2)
      }
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "This variable reports the overall status of the
      storage subsystem.
      GREEN: Normal Operating Condition
      YELLOW: Warning Condition
```

RED: Subsystem Communication Failure; No response to Agent requests."

   ::= { ssEntry 3 }

ssDiskStatus OBJECT-TYPE
 SYNTAX INTEGER {
        ALL_OK(0),
        AT_LEAST_1_FAILURE(1)
      }
   ACCESS read-only
   STATUS mandatory
   DESCRIPTION
     "This variable reports the overall status for the
      physical devices in the subsystem."

   ::= { ssEntry 4 }

ssPowerStatus OBJECT-TYPE
SYNTAX INTEGER {
        ALL_OK(0),
        AT_LEAST_1_FAILURE(1)
      }
   ACCESS read-only
   STATUS mandatory
   DESCRIPTION
     "This variable reports the status for the
      Power Supply system in the subsystem."

   ::= { ssEntry 5 }

ssFanStatus OBJECT-TYPE
SYNTAX INTEGER {
        ALL_OK(0),
        AT_LEAST_1_FAILURE(1)
      }
   ACCESS read-only
   STATUS mandatory
   DESCRIPTION
     "This variable reports the status for the
      FAN system in the subsystem."

   ::= { ssEntry 6 }

ssCacheBatteryStatus OBJECT-TYPE
SYNTAX INTEGER {
        NORMAL_CHARGE(0),

```
        LOW_CHARGE(1)
     }
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "This variable reports the status for the
      Cache Battery in the subsystem. This status object is currently
      non-functional but in place for future use."

   ::= { ssEntry 7 }

ssTemperatureStatus OBJECT-TYPE
SYNTAX  INTEGER  {
        NORMAL(0),
        OVER_TEMP_CONDITION(1)
     }
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "This variable reports the status for the
      Temperature in the subsystem."

   ::= { ssEntry 8 }

ssCommStatus OBJECT-TYPE
SYNTAX  INTEGER  {
        NORMAL(0),
        COMMUNICATION_FAILURE(1)
     }
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "This variable reports the status for the
      communication to the Primary Controller."

   ::= { ssEntry 9 }

ssEmuExtInputStatus OBJECT-TYPE
SYNTAX  INTEGER  {
        ALL_OK(0),
        AT_LEAST_1_FAILURE(1)
     }
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "This variable reports the status for the
      EMU external input."
```

```
   ::= { ssEntry 10 }

ssEmuPresent OBJECT-TYPE
SYNTAX  INTEGER  {
        EMU_NOT_PRESENT(0),
        EMU_PRESENT(1)
      }
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "This variable reports the availability of the
      Environment Monitoring Unit for the storage subsystem."

   ::= { ssEntry 11 }


ssController1Status OBJECT-TYPE
SYNTAX  INTEGER  {
        NORMAL(0),
        FAILED(1)
      }
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "This variable reports the status for the
      Primary Controller."

   ::= { ssEntry 12 }

ssController1SerNum OBJECT-TYPE
   SYNTAX  DisplayString
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
     "The Primary Controller's Serial Number."

   ::= { ssEntry 13 }

ssController1RDAC OBJECT-TYPE
SYNTAX  INTEGER  {
        SINGLE_MODE(0),
        DUAL_MODE(1)
      }
   ACCESS  read-only
   STATUS  mandatory
   DESCRIPTION
```

"The Primary Controller MODE."

   ::= { ssEntry 14 }

ssController1Type OBJECT-TYPE
   SYNTAX DisplayString
   ACCESS read-only
   STATUS mandatory
   DESCRIPTION
     "The Primary Controller's Type indicator."

   ::= { ssEntry 15 }

ssController1ProdID OBJECT-TYPE
   SYNTAX DisplayString
   ACCESS read-only
   STATUS mandatory
   DESCRIPTION
     "The Primary Controller's Product ID."

   ::= { ssEntry 16 }

ssController1FwRev OBJECT-TYPE
   SYNTAX DisplayString
   ACCESS read-only
   STATUS mandatory
   DESCRIPTION
     "The Primary Controller's Firmware Revision."

   ::= { ssEntry 17 }

ssController1HwRev OBJECT-TYPE
   SYNTAX DisplayString
   ACCESS read-only
   STATUS mandatory
   DESCRIPTION
     "The Primary Controller's Hardware Revision."

   ::= { ssEntry 18 }

ssController2Status OBJECT-TYPE
SYNTAX INTEGER {
        NORMAL(0),
        FAILED(1)
       }
   ACCESS read-only
   STATUS mandatory

```
    DESCRIPTION
      "This variable reports the status for the
       Secondary Controller."

    ::= { ssEntry 19 }

ssController2SerNum OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "The Secondary Controller's Serial Number."

    ::= { ssEntry 20 }

ssController2RDAC OBJECT-TYPE
SYNTAX  INTEGER  {
        SINGLE_MODE(0),
        DUAL_MODE(1)
      }
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "The Secondary Controller MODE."

    ::= { ssEntry 21 }

ssController2Type OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "The Secondary Controller's Type indicator."

    ::= { ssEntry 22 }

ssController2ProdID OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
      "The Secondary Controller's Product ID."

    ::= { ssEntry 23 }

ssController2FwRev OBJECT-TYPE
    SYNTAX  DisplayString
```

```
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "The Secondary Controller's Firmware Revision."

    ::= { ssEntry 24 }

ssController2HwRev OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "The Secondary Controller's Hardware Revision."

    ::= { ssEntry 25 }

ssUpsStatus OBJECT-TYPE
SYNTAX  INTEGER  {
        UPS_NOT_PRESENT(0),
        UPS_PRESENT(1)
      }
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "Uninterruptable Power Supply status (Currently for future use)."

    ::= { ssEntry 26 }

ssLunStatus OBJECT-TYPE
SYNTAX  INTEGER  {
        UPS_NOT_PRESENT(0),
        UPS_PRESENT(1)
      }
    ACCESS  read-only
    STATUS  mandatory
    DESCRIPTION
       "LUN status."

    ::= { ssEntry 27 }

 END
```

5.4   General Troubleshooting Guidelines

Before you call customer service, review the following information and the latest
Command Console release notes to determine if you have a problem that you can easily
correct. Also, keep in mind that system performance characteristics such as scan and
initialization times directly relate to the capacity of you storage subsystem. If you are
using Client over a network, response times are related to network speed and load.

Initializing Virtual Disks

The initialization time for a virtual disk is dependent on the number of members in the
virtual disk and the size of the virtual disk. The virtual disk is not available for I/O during
this time.

Error Displayed but Current Status Good

If you notice that an error is displayed in Client but the current subsystem status is good,
an error may have occurred in the past with a subsequent break in the communications
between Client and the controller. To clear the error, you must exit and restart Client.

5.5   Troubleshooting Client and Agent Connections

- Make sure the system names used for the Client and Agent match each other consistently. As a general rule, always use lower-case letters when entering names.
- If you have entered the host name in the NT client's \windows\system32\ drivers\ etc\hosts file, or the Windows 95 client's \windows\hosts file, make sure it is spelled correctly using lower-case characters. The host name is not required in this file, but TCP/IP will look here first when resolving host names.
- Make sure you have the correct port numbers in the client's services file. On Windows NT, look in the \windows\system32\drivers\ etc\services file. On Windows 95, look in the \windows\services file. Be sure the port entries are followed by a carriage return if they are the last entries in the file:

spgui	4998/tcp
Spagent	4999/tcp

- Make sure the Command Console Agent on the host system is using a correct Domain Name Service (DNS) name to recognize the Client. If the host system is Windows NT, run Agent Configure (Config.exe) or examine the client.ini file on the host system. On other host systems, look at the client.ini file.
- If you are using a Windows NT Client, run netstat at the MS-DOS prompt to check the TCP/IP active connections.
- Check the Client name that the Agent's host system sees is the same as the name specified in the Agent's client.ini file. Run netstat at the MS-DOS prompt on the Agent, if you are using a Windows NT host system, and check the TCP/IP active connections. On Unix agents, run netstat with the correct options (see *netstat man* pages in Unix help information).
- Make sure the Agent's host system has been configured to recognize the Client by its assigned name rather than by an alias name. If the host system has been configured to recognize the Client by an alias name, remove the alias name from the system's "hosts" file.

NOTE

- Always reboot the Client system after making changes to the hosts or services file.

- Command Console does not support DHCP (Dynamic Host Configuration Protocol) connections, where Clients are configured for TCP/IP at system startup time.

5.6 Windows NT Event Logging

Command Console uses the event logging facility in WindowsNT to create a history of error and other subsystem events. You can use system and application event logs to help you troubleshoot subsystem problems.

5.7 Contacting Customer Support

If you have a problem with Command Console that you cannot resolve, contact your sales representative for information on how to receive technical support. You can also find the latest product information at: *http://www.compaq.com/products/storageworks*

A

Creating Storagesets

This appendix provides supplemental guidelines for determining and planning storage requirements. For a complete discussion of RAID, refer to The RAIDBOOK—A Source Book for Disk Array Technology.

A.1   Introduction - What are Storagesets

Storagesets are implementations of RAID technology, also known as a "Redundant Array of Independent Disks." Every storageset shares one important feature: each one looks like a single storage unit to the host, regardless of the number of drives it uses.

You can create storage units by combining disk drives into storagesets, such as stripesets, RAIDsets, and mirrorsets, or by presenting them to the host as single-disk units, as shown in Figure A–1.

- Stripesets (RAID 0) combine disk drives in serial to increase transfer or request rates.

- Mirrorsets (RAID 1) combine disk drives in parallel to provide a highly-reliable storage unit.

- RAIDsets (RAID 3/5) combine disk drives in serial—as do stripesets—but also store parity data to ensure high reliability.

- Striped mirrorsets (RAID 0+1) combine mirrorsets in serial to provide the highest throughput and availability of any storage unit.

Controllers can support the number of RAIDsets as listed in Table A-1.

Figure A–1 Units Created from Storagesets, Partitions, and Drives
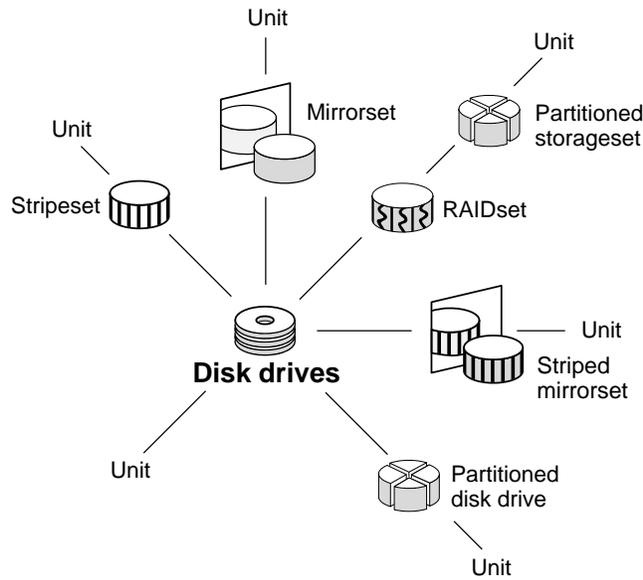
CXO5368B

Table A–1 Controller Limitations for RAIDsets

RAIDset Type	Limit
Total Number of RAID5	20
Total Number of RAID5 + RAID1	30
Total Number of RAID5 + RAID1 + RAID0	45

NOTE

Limitations on the number and size of RAID0, RAID1, and RAID5 storagesets as well as other HSG80 controller capabilities are detailed in the ACS section of the Release Notes.

A.2   Determining Storage Requirements

Start the planning process by determining your storage requirements. Here are a few of the questions you should ask yourself:

- What applications or user groups will access the subsystem? How much capacity do they need?

- What are the I/O requirements? If an application is data-transfer intensive, what is the required transfer rate? If it is I/O-request intensive, what is the required response time? What is the read/write ratio for a typical request?

- Are most I/O requests directed to a small percentage of the disk drives? Do you want to keep it that way or balance the I/O load?

- Do you store mission-critical data? Is availability the highest priority, or would standard backup procedures suffice?

Use your responses to these questions along with Table A–2 to determine the types of storagesets you should create to satisfy your organization's requirements.

A.3   Choosing a Storageset Type

Different applications may have different storage requirements, so you will probably want to configure more than one kind of storageset in your subsystem.

All of the storagesets described in this book implement RAID (Redundant Array of Independent Disks) technology. Consequently, they all share one important feature: each storageset, whether it contains two disk drives or ten, looks like one large, virtual disk drive to the host.

Table A–2 compares different kinds of storagesets to help you determine which ones satisfy your requirements.

Table A-2 A Comparison of Different Kinds of Storagesets

Storageset	Relative Availability	Request Rate (Read/Write) I/O per second	Transfer Rate (Read/Write) MB per second	Applications
Array of disk drives (JBOD)	Equivalent to a single disk drive	Identical to single disk drive	Identical to single disk drive	
Stripeset (RAID 0)	Proportionate to number of disk drives; worse than single disk drive	Excellent if used with large chunk size	Excellent if used with small chunk size	High performance for non-critical data
Mirrorset (RAID 1)	Excellent	Good/Fair	Good/Fair	System drives critical files
RAIDset (RAID 3/5)	Excellent	Excellent/Fair	Good/Poor	High request rates, read-intensive, data lookup
Striped Mirrorset (RAID 0 + 1)	Excellent	Excellent if used with large chunk size	Excellent if used with small chunk size	Any critical response-time application

A.3.1 Using Stripesets to Increase I/O Performance

Stripesets enhance I/O performance by spreading the data across multiple disk drives. Each I/O request is broken into small segments called *chunks.* These chunks are then *striped* across the disk drives in the storageset, thereby allowing several disk drives to participate in one I/O request to handle several I/O requests simultaneously.

For example, in a three-member stripeset that contains disk drives 10000, 20000, and 30000, the first chunk of an I/O request is written to 10000, the second to 20000, the third to 30000, the fourth to 10000, and so forth until all of the data has been written to the drives.

Figure A–2 Striping Lets Several Disk Drives Participate in Each I/O Request
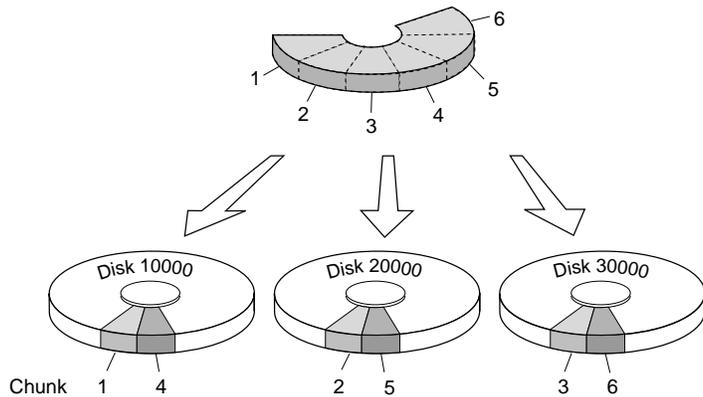
CXO5507A

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rate. You can set the chunk size or let the controller set it automatically. Refer to "Chunk Size" in Chapter 3 of the *HSG80 Array Controller ACS Version 8.3 User's Guide* .

A major benefit of striping is that it balances the I/O load across all of the disk drives in the storageset. This can increase the subsystem's performance by eliminating the hot spots, or high localities of reference, that occur when frequently-accessed data becomes concentrated on a single disk drive.

Considerations for Planning a Stripeset

Keep the following points in mind as you plan your stripesets:

- A controller can support up to 45 storagesets, consisting of stripesets, mirrorsets and RAIDsets (refer to Table A–1).

- Reporting methods and size limitations prevent certain operating systems from working with large stripesets. See the *HSG80 Array Controller ACS Version 8.3 Release Notes* and the *Getting Started Guide* that came with your platform kit for details about these restrictions.

- A storageset should only contain disk drives of the same capacity. The controller limits the capacity of each member to the capacity of the smallest member in the storageset when the storageset is initialized (the base member size). Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, the 4 GB disk drive will be the base member size, and you will waste 5 GB of capacity on each 9 GB member.

- If you need high performance and high availability, consider using a RAIDset, striped mirrorset, or a host-based shadow of a stripeset.

- Striping does not protect against data loss. In fact, because the failure of one member is equivalent to the failure of the entire stripeset, the likelihood of losing data is higher for a stripeset than for a single disk drive.

For example, if the mean time between failures (MTBF) for a single disk is one hour, then the MTBF for a stripeset that comprises N such disks is l/N hours. As another example, if a single disk's MTBF is 150,000 hours (about 17 years), a stripeset comprising four of these disks would only have an MTBF of slightly more than four years.

For this reason, you should avoid using a stripeset to store critical data. Stripesets are more suitable for storing data that can be reproduced easily or whose loss does not prevent the system from supporting its critical mission.

- Evenly distribute the members across the device ports to balance load and provide multiple paths as shown in Figure A–3.

Figure A–3 Distribute Members across Device Ports

CXO6235B
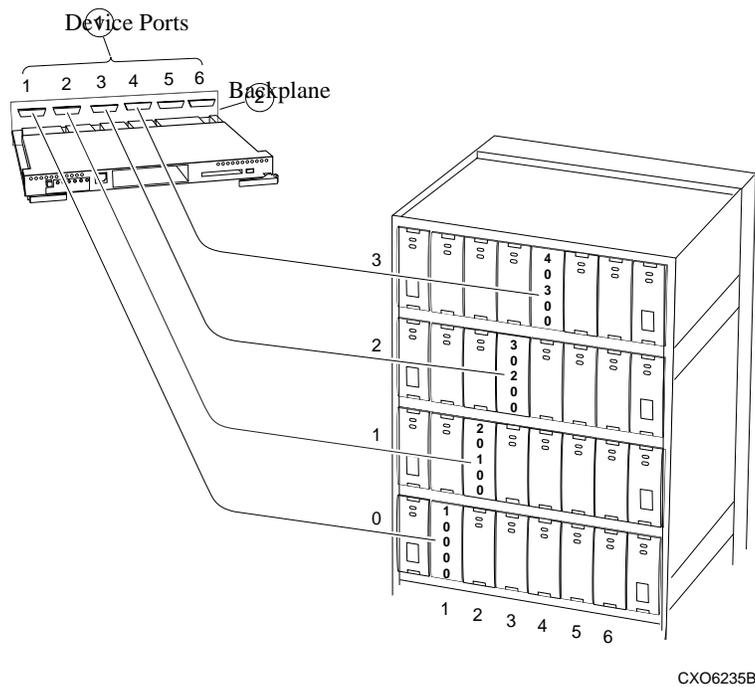
- Stripesets contain between 2 and 14 members.

- Stripesets are well suited for the following applications:

  ⇒ Storing program image libraries or run-time libraries for rapid loading.

  ⇒ Storing large tables or other structures of read-only data for rapid application access.

  ⇒ Collecting data from external sources at very high data transfer rates.

- Stripesets are not well-suited for the following applications:

  ⇒ A storage solution for data that cannot be easily reproduced or for data that must be available for system operation

  ⇒ Applications that make requests for small amounts of sequentially-located data

  ⇒ Applications that make synchronous random requests for small amounts of data

By spreading the traffic evenly across the buses, you will ensure that no bus handles the majority of data to the storageset.

A.3.2 Using Mirrorsets to Ensure Availability

Mirrorsets use redundancy to ensure availability, as illustrated in Figure A –4. For each primary disk drive, there is at least one mirror disk drive. Thus, if a primary disk drive fails, its mirror drive immediately provides an exact copy of the data.

Figure A–4 Mirrorsets Maintain Two Copies of the Same Data

Disk 1010_0 — A

Disk 1000_0 — A'

Disk 2010_0 — B

Disk 2000_0 — B'

Disk 3010_0 — C

Disk 3000_0 — C'

Mirror drives contain copy of data
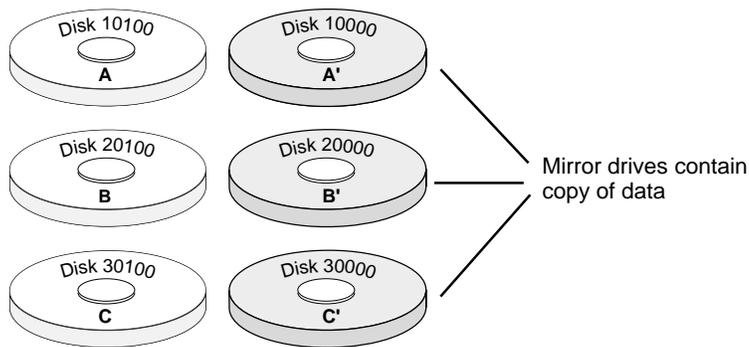
CXO5511A

Considerations for Planning a Mirrorset

Keep these points in mind as you plan your mirrorsets:

- A controller can support up to 30 storagesets, consisting of mirrorsets and RAIDsets. Mirrorsets that are members of a stripeset count against this limitation (refer to Table A–1).

- Data availability with a mirrorset is excellent but costly —you need twice as many disk drives to satisfy a given capacity requirement. If availability is your top priority, consider using redundant power supplies and dual-redundant controllers.

- You can configure up to 30 mirrorsets per controller or pair of dual-redundant controllers. Each mirrorset contains a minimum of one and a maximum of six members.

- A write-back cache module is required for mirrorsets, but write-back cache need not be enabled for the mirrorset to function properly.

- Both write-back cache modules must be the same size.

- If you're using more than one mirrorset in your subsystem, you should put the first member of each mirrorset on different buses as shown in Figure A-8. (The first member of a mirrorset is the first disk drive you add.)

When a controller receives a request to read data from a mirrorset, it typically accesses the first member of the mirrorset. Read access depends upon the read source switches. If you have several mirrorsets in your subsystem and their first members are on the same bus, that bus will be forced to handle the majority of traffic to your mirrorsets. When a controller receives a request to write data to a mirrorset, it accesses and writes to all members.

Figure A–5 First Mirrorset Members on Different Buses

First member
of mirrorset 1

First member of
mirrorset 2
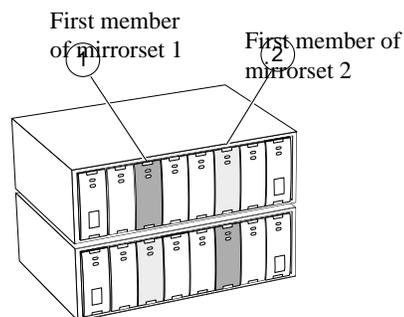
CXO5506B

To avoid an I/O bottleneck on one bus, you can simply put the first members on different buses. Additionally, you can set the read-source switch to Round Robin. Refer to "Mirrorset Switches" in Chapter 3 of the *HSG80 Array Controller ACS Version 8.3 User's Guide* for more information about this switch.

- Place mirrorsets and RAIDsets on different ports to minimize risk in the event of a single port bus failure.

- Mirrorset units are set to WRITEBACK_CACHE by default which increases a unit's performance.

- A storageset should only contain disk drives of the same capacity. The controller limits the capacity of each member to the capacity of the smallest member in the storageset. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, the 4 GB disk drive will be the base member size, and you waste 5 GB of capacity on each 9 GB member.

- Evenly distribute the members across the device ports to balance load and provide multiple paths.

- Mirrorsets are well suited for the following:

  ⇒ Any data for which reliability requirements are extremely high.

  ⇒ Data to which high-performance access is required.

  ⇒ Applications for which cost is a secondary issue.

- Mirrorsets are not well suited for the following applications:

  ⇒ Write-intensive applications (JBODs) are better for this type of application, but mirrorsets are preferred over Raid 5 RAIDsets.)

  ⇒ Applications for which cost is a primary issue

A.3.3 Using RAIDsets to Increase Performance and Availability

RAIDsets are enhanced stripesets—they use striping to increase I/O performance and distributed-parity data to ensure data availability. Figure A -6 illustrates the concept of RAIDsets and parity data.

Figure A–6 Parity Ensures Availability; Striping Provides Good Performance
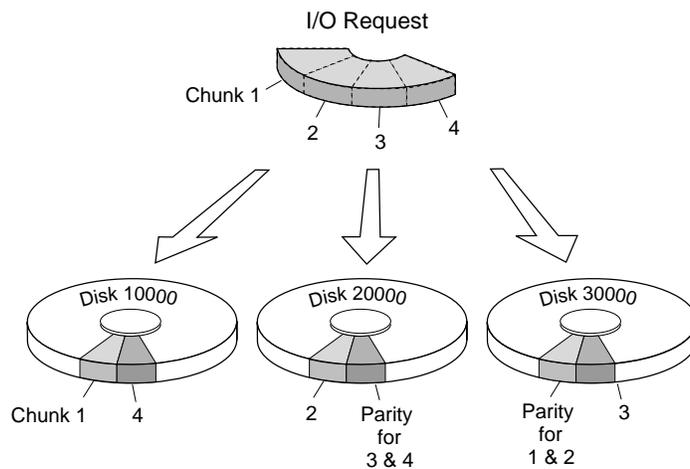
CXO5509A

Just as with stripesets, the I/O requests are broken into smaller *chunks* and striped across the disk drives until the request is read or written. But, in addition to the I/O data, chunks of parity data—derived mathematically from the I/O data —are also striped across the disk drives. This parity data enables the controller to reconstruct the I/O data if a disk drive fails. Thus, it becomes possible to lose a disk drive without losing access to the data it contained. (Data could be lost if a second disk drive fails before the controller replaces the first failed disk drive.)

For example, in a three-member RAIDset that contains disk drives 10000, 20000, and 30000, the first chunk of an I/O request is written to 10000, the second to 20000, then parity is calculated and written to 30000; the third chunk is written to 30000, the fourth to 10000, and so on until all of the data is saved.

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rates. You can set the chunk size or let the controller set it automatically. Refer to "Chunk Size" in Chapter 3 of t he *HSG80 Array Controller ACS Version 8.3 User's Guide* for more information on setting the chunk size.

Considerations for Planning a RAIDset

Keep these points in mind as you plan your RAIDsets:

- A controller can support up to 20 storagesets, consisting of RAIDsets (refer to Table A–1).

- Reporting methods and size limitations prevent certain operating systems from working with large RAIDsets. See the *HSG80 Array Controller ACS Version 8.3 Release Notes* or the *Getting Started Guide* that came with your platform kit for details about these restrictions.

- A cache module is required for RAIDsets, but write-back cache need not be enabled for the RAIDset to function properly.

- A RAIDset must include at least 3 disk drives, but no more than 14.

- Evenly distribute the members across the device ports to balance load and provide multiple paths as shown in Figure A–3.

- A storageset should only contain disk drives of the same capacity. The controller limits the capacity of each member to the capacity of the smallest member in the storageset when the storageset is initialized (the base member size). Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you will waste 5 GB of capacity on each 9 GB member.

- RAIDset units are set to WRITEBACK_CACHE by default which increases a unit's performance.

- RAIDsets and mirrorsets on different ports to minimize risk in the event of a single port bus failure.

- RAIDsets are particularly well-suited for the following:
  - ⇒ Small to medium I/O requests
  - ⇒ Applications requiring high availability
  - ⇒ High read request rates
  - ⇒ Inquiry-type transaction processing

- RAIDsets are not particularly well suited for the following:
  - ⇒ Write-intensive applications
  - ⇒ Applications that require high data transfer capacity
  - ⇒ High-speed data collection
  - ⇒ Database applications in which fields are continually updated
  - ⇒ Transaction processing

A.3.4 Using Striped Mirrorsets for Highest Performance and Availability

As illustrated in Figure A–10, striped mirrorsets are simply stripesets whose members are mirrorsets. Consequently, this kind of storageset combines the performance of striping with the reliability of mirroring. The result is a storageset with very high I/O performance and high data availability.

Figure A–7 Striping and Mirroring in the Same Storageset

CXO5508A
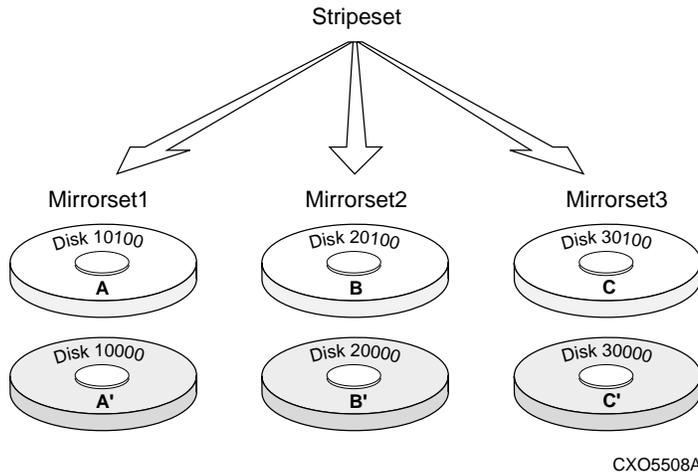
The failure of a single disk drive has no effect on this storageset's ability to deliver data to the host and, under normal circumstances, it has very little effect on performance. Because striped mirrorsets do not require any more disk drives than mirrorsets, this storageset is an excellent choice for data that warrants mirroring.

Considerations for Planning a Striped Mirrorset

Plan the mirrorset members, then plan the stripeset that will contain them. Review the recommendations in Section A.3.1, *Considerations for Planning a Stripeset*, and Section A.3.2, *Considerations for Planning a Mirrorset*.

There are the following limitations for a striped mirrorset:

- A maximum of 14 mirrorsets in a stripeset.

- A maximum of 6 disks in each mirrorset.

- A maximum of 48 disks in the entire striped mirrorset.

A.4 Planning Partitions

Use partitions to divide a storageset or disk drive into smaller pieces, which can each be presented to the host as its own storage unit. Figure A –8 shows the conceptual effects of partitioning a single-disk unit.

Figure A–8 Partitioning a Single-Disk Unit

Partition 1
Partition 2
Partition 3
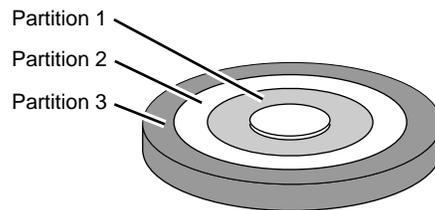
CXO-5316A-MC

You can create up to eight partitions per disk drive, RAIDset, mirrorset, stripeset, or striped mirrorset. Each partition has its own unit number so that the host can send I/O requests to the partition just as it would to any unpartitioned storageset or device. Because partitions are separately-addressable storage units, you can partition a single storageset to service more than one user group or application.

A.4.1 Defining a Partition

Partitions are expressed as a percentage of the storageset or single disk unit that contains them. For mirrorsets and single disk units, the controller allocates the largest whole number of blocks that are equal to or less than the percentage you specify. For RAIDsets and stripesets, the controller allocates the largest whole number of stripes that are less than or equal to the percentage you specify. For stripesets, the stripe size = chunk size x number of members. For RAIDsets, the stripe size = chunk size x (number of members-1).

An unpartitioned storage unit has more capacity than a partition that uses the whole unit because each partition requires five blocks of administrative metadata. Thus, a single disk unit that contains one partition can store n-5 blocks of user or application data.

A.4.2 Guidelines for Partitioning Storagesets and Disk Drives

Keep these points in mind as you plan your partitions:

- You can create up to eight partitions per storageset or disk drive.

- All of the partitions on the same storageset or disk drive must be addressed through the same controller port. This ensures a transparent failover of devices should one of the dual-redundant controllers fail.

- Partitions cannot be combined into storagesets. For example, you cannot divide a disk drive into three partitions, then combine those partitions into a RAIDset.

- Partitioned storagesets cannot function in multiple bus failover dual-redundant configurations. Because they are not supported, you must delete your partitions before configuring the controllers for multiple bus failover.

- Once you partition a container, you cannot unpartition it without reinitializing the container.

- Just as with storagesets, you do not have to assign unit numbers to partitions until you are ready to use them.

B

Fibre Channel Switch Installation and Basic Set-up

This Appendix is designed to give you basic instructions for installing the Fibre Channel switch and for making the necessary connections. For complete details of this procedure refer to Compaq StorageWorks Quick Setup Guide StorageWorks Fibre Channel Switch, AA-RHCOA-TE and the Quick Setup Guide RAID Array 8000/ESA12000 Fibre Channel Storage Subsystem for Compaq Tru64 UNIX (ACS V8.4 for Fibre Channel Switch), AA-RFASA-TE/387388-001.

B.1 Procedure

The following steps are to be used in conjunction with Figure B –1.

1. Select a location and mounting method (refer to the Installation Card or User's Guide for details).

2. Connect the StorageWorks Subsystem:

   - Install a GBIC module in switch port 0.
   - Connect the fibre-optic cable from the StorageWorks subsystem to the GBIC (1).

3. Connect the Host:

   - Install a GBIC module in switch port 1.
   - Connect the fibre-optic cable from the host to the GBIC (7).

4. Connect the Ethernet Network:

   - Connect the Ethernet cable to the front panel RJ-45 connector (6).

5. Connect the AC Power:

   - Connect the power cord to the AC connector (4).
   - Turn on the powers switch (3).

Figure B–1 Connecting to the Fibre Channel Switch

SHR-1309
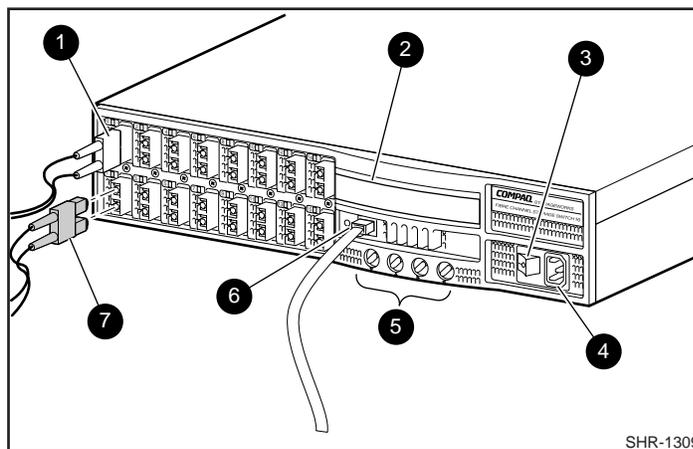
6.   Verify successful POST:

  •   Ensure power LED shows steady green.

  •   Ensure Port 0 and 1 LEDs show steady green (providing the ports are connected).

7.   Assign an IP Address to the Switch:

  •   Press the front panel *Down* button (5).

  •   When *Select Menu: Configuration Menu* appears on the display (2), use the down button to select *Ethernet IP Address*.

  •   Press the *Enter* button to display the switch Ethernet IP Address.

  •   Use the *Up* and *Down* buttons to increment or decrement the displayed value; use the *Tab/Esc* button to move the cursor between fields.

  •   When you have set the address, press the *Enter* button to store the address.

Glossary

This Chapter provides a glossary of Storage Subsystem terms.

adapter

A device that converts the protocol and hardware interface of one bus type into that of another without changing the functionality of the bus.

agent

Command Console's secure, TCP/IP network communication program. Agent runs on your storage subsystems' host computers and is available in a number of popular operating system versions. An Agent running on a host system can be accessed by one or more client programs. One Agent can communicate with multiple storage subsystems on the same host system using the system's SCSI bus. You can configure Agent with password protection to prevent unauthorized users from accessing your storage subsystems.

array controller

A hardware and software device that facilitates communications between a system and one or more devices organized in an array.

available

A device state in which the device is operational but not yet used as a member of a virtual disk.

bad block replacement (BBR)

The substitution of defect-free device blocks for those found to have defects. The procedure used to locate a replacement block, mark the bad block as replaced, and move the data from the bad block to the replacement block. This process takes place automatically in some controllers.

block

A number of consecutive bytes of data stored on a storage device. In most storage subsystems, a block is the same size as a physical disk sector.

blower

An airflow device mounted in a device or controller shelf.

cache

A fast, temporary, storage buffer in a controller or computer.

capacity

The total amount of data capable of being stored in a physical or virtual disk.

channel

A SCSI device bus. Narrow SCSI device buses support up to 7 devices, and wide buses up to 15. Sometimes also called "port".

chunk

HS-series controller term for strip.

chunk size

HS-series controller term for strip size.

client

The Command Console program that contains the graphic user interface and most of the functionality of the product. Client uses its associate program, Agent, for network connections to multiple storage subsystems. Client can also connect to one storage subsystem via either a serial maintenance port or host SCSI bus connection.
Client consists of two major components:

- The Navigation Window---provides an upper-level view of your host systems and storage subsystems.
- One or more Storage Windows---a detailed graphical window into one storage subsystem.

client list

List of Client network nodes that are allowed to connect to an Agent over the network. The list is maintained in the client.ini file on the host where the Agent is installed. A client list entry contains the TCP/IP address or network name of each Client along with its allowed access level for the host's storage subsystems.

cold swap

A method of device replacement that requires that power be removed from the controller and all shelves in a cabinet. This method is used when conditions preclude the use of the warm swap or hot swap methods.

configuration file

A file in which information describing the configuration of a storage subsystem is stored. The file can be used to automatically configure the subsystem.

container

The HS-series controller term for device group.

controller

A firmware-driven hardware device that expedites communication between a host and one or more storage devices. A controller translates bus protocols and hardware interfaces and adds functionality to the storage subsystem. Some controllers add additional functionality by mapping storage devices to host-addressable, virtual disk with specific performance and availability characteristics using RAID techniques.

controller shelf

A shelf designed to contain controller and cache memory modules. A controller shelf can contain two controllers in a dual-redundant configuration.

copy speed

For HS-series controllers, the rate at which the controller writes mirrored data in a mirrored virtual disk. There are two copy speeds:
- Normal--uses relatively few controller resources to perform the copy, and has little impact on controller performance.
- Fast--uses more controller resources, which reduces the time it takes to complete the copy, but also reduces overall controller performance.

device

In its physical form, a magnetic or optical disk, tape, or CD-ROM that can be attached to a SCSI bus. A device provides a host with large amounts of addressable storage capacity.
The term is also used to indicate a physical device that has been made part of a controller's configuration. That is, a physical device that is known to the controller. Virtual disks are created from devices, once they have been made known to the controller.

device driver

A program that processes I/O requests for a particular device type.

device group

A logical, internal controller structure, representing one or more devices linked as a group. In some controllers, you must create device groups before you can create a virtual disk from them. Device groups are sometimes known as containers.
There are many types of device groups:
- Single-device virtual disks (JBODs)
- Striped virtual disks (RAID 0)
- Mirrored virtual disks (RAID 1)
- Striped mirrored virtual disks (RAID 0+1)
- Striped parity virtual disks with fixed parity disk (RAID 4)
- Striped parity virtual disks with floating parity disk (RAID 5)

disk

A storage device that uses rotating, magnetic media to store data.

disk array

A collection of disk devices physically connected in an ordered structure.

dual-redundant configuration

For HS-series controllers, a configuration consisting of a primary and backup controller in one controller shelf. If the primary controller fails, the backup controller assumes control over the failing controller's devices.

environmental monitoring unit (EMU)

Some subsystem enclosures include an environmental monitoring unit to provide increased protection against catastrophic faults. The EMU works with the controller to give warning about various types of impending subsystem failures. The EMU senses such conditions as failed power supplies, failed blowers, elevated temperatures, and external air sense faults.

failed

A device state that indicates a device is inoperable and is unavailable for use in a virtual disk. In some controllers, you can force a device to the failed state, if you want to remove the device from the storage enclosure. In others, a device is marked as failed when you remove it as a member of a virtual disk.

failover

For HS-series controllers, the process that takes place when one controller in a dual-redundant configuration fails, and the other controller takes over direction of the storage subsystem. The other controller continues to direct the storage subsystem until the failed controller becomes operational or is replaced.

good

A device state that indicates that the device is operational and is used in a virtual disk.

host

The primary or controlling computer to which a storage subsystem is attached.

host access ID

For HSZ-series controllers, the target ID of the host computer allowed exclusive access to a particular virtual disk. A virtual disk can be set to allow all hosts or any specific host target ID access to it.

host adapter

A device that connects a host system to a SCSI bus. The device usually performs the lowest layers of the SCSI protocol. This function may be logically and physically integrated into the host system.

host functionality

For HSZ80-series controllers, you can configure your controller's host port targets for optimum performance and compatibility for VMS, UNIX and Windows NT. There are two modes to choose from. The following table shows the options and their corresponding mode designation:

Mode	Host compatibility
A	
B	IBM AIX
C	
D	
E	SGI IRIX
F	Sun Solaris

hot disks

Hot disks occur when the workload is poorly distributed across storage devices. A hot disk is a device with multiple hot spots. On a hot disk, host I/O requests begin to back up because of the concentrated request load. Hot disks cause subsystem performance to suffer.

hot spare

A device configured to automatically replace a failed device in certain types of virtual disks. If a virtual disk member fails, the controller automatically replaces the member with the spare device, and it rebuilds the member's data from the remaining devices in the virtual disk.

hot spots

Hot spots occur when the workload is poorly distributed across storage devices. A hot spot is a file or a group of files located on the same device and which receive a very high concentration of I/O requests.

hot swap

A method of device replacement in which the system remains operational during device removal and reinstallation. The device being removed or reinstalled is the only device that is inoperational during this process.

initialization

For storage subsystems, the process of restarting the controllers and reestablishing the subsystem configuration and its operation. That is, bootstrapping the storage subsystem.
For virtual disks, the process of writing the controller's file structure on the virtual disk's member devices. The file structure is represented on the member disks by metadata.

JBOD (Just a Bunch of Disks)

JBOD is an industry term for a single-device virtual disk. A JBOD virtual disk does not provide any level of data redundancy.

known client list

See client list.

local area network (LAN)

A network that is confined to a single geographic location.

local connection

A direct, hardware connection from Client to a single subsystem using either of two methods:
- A connection from a standalone computer running Client to the serial maintenance port on the controller
- A connection from a host system running Client to a storage subsystem on the host's SCSI bus

A local connection enables you to connect Client to one subsystem controller within the physical range of the serial or host SCSI cable.

logical drive

See virtual disk.

logical storage unit

See virtual disk.

logical unit number (LUN)

A SCSI bus supports a number of target devices, each with a number of logical units within it. A logical unit number is a logical unit's address on a target.

A controller maps one or more SCSI logical units to a virtual disk that is accessible by a host. The term "LUN" is also sometimes used to refer to a virtual disk itself.

mapping

The internal, controller process of organizing single or grouped devices into virtual storage units (that is, virtual disks).

maximum block transfer

For HS-series controllers, a virtual disk's characteristic that specifies the maximum number of blocks of data to be cached before data is written to the virtual disk's members. A data transfer greater than this size is not cached and is written immediately to disk. Valid values are from 1 to 1024 blocks.

member

Any of the devices that is used in a virtual disk based on multiple devices. Any of the devices in the device group on which the virtual disk is based.

metadata

Special data written to a device for the purposes of controller administration. Metadata improves error detection and media defect management for the device. It is also used to define device and virtual disk configuration.

mirrored cache

Some controllers offer high-performance cache hardware that can be set to operate in the mirrored mode. In this mode, the cache data is duplicated and stored in the cache's memory in physically separate memory locations. If one copy of the data is corrupted or unavailable for some reason, the other copy is available for use.

A mirrored cache configuration provides complete protection for cached data. Because the cache memory effectively contains twice as much data for each host request, however, the effective size of the cache is cut in half.

mirrored virtual disk

A group of storage devices organized as duplicate copies of each other. Mirrored virtual disks provide the highest level of data availability at the highest cost. Another name for RAID 1.

mirroring

The simplest form of data redundancy. Two or more devices form a mirrored device group. Each device is an exact copy of the other.
For read requests, Data can be read from either device therefore increasing the throughput. Both devices can handle different requests simultaneously.
For write requests, data is written to both devices. If a device fails, all reads/writes are executed on the mirrored device group's other device. Therefore, mirroring provides excellent availability unless the second device fails before the first device can be repaired or replaced.
Sometimes also called shadowing.
The mirroring technique can also be applied to cache memory.

navigation tree

A window that provides you an upper-level, graphical interface for monitoring your host systems and storage subsystems. One of Client's components.

network connection

A remote connection from Client to multiple storage subsystems using Agent over a TCP/IP-compatible network. Network connections enable you to concurrently monitor thousands of storage subsystems over a LAN or WAN anywhere in the world.

nonredundant

For HS-series controllers, a single-controller configuration.
The term is also sometimes used to refer to a RAID level that does not provide fault tolerance (for example, RAID 0).

normal

In mirrored virtual disks, a member state that indicates that a member contains exactly the same data as that on the other members.

parameter

A setting or value that defines a particular operating characteristic of a device, virtual disk, controller, or storage subsystem.

parity

Binary value calculated from user data and used to detect when associated data becomes corrupted. Parity can also be used to correct corrupted data. Striped parity virtual disks use parity to improve data availability.

parity check

A virtual disk state that indicates that the controller is checking the integrity of a virtual disk's data.

partition

A logical division of a virtual disk, represented itself as a virtual disk to the host. Command Console enables the creation of virtual device partitions, but they are transparent to the user.

password

A code string used by Client with network connections to restrict configuration functions to privileged users. Passwords are set at the Agent at each host.

pedestal

A desk-side, floor-standing, storage enclosure capable of housing one or more controllers and associated devices to make up a complete storage subsystem.

physical device

A storage component installed in a storage subsystem.

physical disk

A storage component with rotating, magnetic media and installed in a storage subsystem.

port

The hardware and software used to connect a host controller to a communications bus, such as a SCSI bus or serial bus. On the device side of the controller, a port is called a channel.
The term is sometimes also used to refer to the logical TCP/IP port that is used to access the Agent over the network.

port/target/LUN (PTL)

For HS-series controllers, the complete address of a physical disk on a device bus, from which the controller derives the device's name. The name, "disk1130", for instance, indicates a physical device at port 1, target 13, and LUN 0.

RAID (Redundant Array of Independent Disks)

RAID is an industry-standard set of techniques for configuring an array of storage devices into host-accessible, virtual disks with various levels of cost, availability, and performance options. Some common RAID levels are: RAID 0, RAID 1, RAID 0+1, RAID 3, RAID 4, RAID 5, and RAID 3/5.

RAID 0

RAID 0 is the industry-standard name for disk striping. A RAID 0 virtual disk is also called a striped virtual disk.
In a RAID 0 virtual disk, host data is broken down into strips spread across the virtual disk's member devices in a stripe. This technique offers much faster read and write performance than does reading and writing to a single device. A three-device, RAID 0 virtual disk has potentially three times the bandwidth of a single device because three, separate, small pieces of host data are moved in parallel.
RAID 0 is the only RAID level that does not provide any level of data redundancy. Because there are more devices to potentially fail, and because there is no way to recover data if one device does fail, the availability of RAID 0 virtual disks is actually lower than an equivalent-sized single disk.

RAID 0+1

RAID 0+1 combines the striping of RAID 0 and the mirroring of RAID 1 to provide the best combination of very high performance and very high availability. A RAID 0+1 virtual disk is also known as a striped mirrored virtual disk.

In a RAID 0+1 virtual disk, each RAID 0 stripe is mirrored to one or more duplicate sets of devices. This technique offers much faster read and write performance than does reading and writing to a single device. A six-device, RAID 0+1 virtual disk has potentially three times the bandwidth of a single device because three, separate, small pieces of host data are moved in parallel.

In addition, because the data is completely mirrored to one or more sets of devices, there is complete data redundancy for very high availability.

A RAID 0+1 virtual disk offers the highest performance and the highest availability of any RAID virtual disk, but it does so at a very high cost. Such a configuration requires at least twice the number of devices of a RAID 0 configuration.

RAID 1

RAID 1 is the industry-standard term for device mirroring. A RAID 1 virtual disk is also known as a mirrored virtual disk.

In a RAID 1 virtual disk, host data is written as a one, large block to one device and the data is mirrored to one or more duplicate disks.

A RAID 1 virtual disk provides very high availability because the data is completely mirrored to one or more devices. Its performance is no better than that of a single device, however, because the data is transferred as one, large block to and from these devices.

RAID 3

RAID 3 virtual disks use parity for data redundancy. A RAID 3 virtual disk is a type of striped parity virtual disk.

In a RAID 3 virtual disk, host data is broken down into strips spread across the virtual disk's member devices in a stripe. An additional strip is appended to each stripe to contain parity information. This technique offers much faster read and write performance than does reading and writing to a single device. A three-device, RAID 3 virtual disk has potentially three times the bandwidth of a single device because three, separate, small pieces of host data are moved in parallel.

In addition, because each stripe of data is protected by parity information, there is a significant level of data redundancy for high availability. Some RAID 3 configurations use a dedicated parity device, but most controllers intersperse the parity strips within the data strips to maximize the read performance.

RAID 3 virtual disks provide high performance and high availability at reasonable cost. They are optimum for use in applications requiring relatively high data transfer rates and having relatively low I/O request rates.

RAID 5

RAID 5 virtual disks use parity for data redundancy. A RAID 5 virtual disk is a type of striped parity virtual disk.

In a RAID 5 virtual disk, host data is written in its entirety as a strip representing one, I/O request into a much larger stripe of I/O requests stored across the virtual disk's member devices. An additional strip is written into each stripe to contain parity information. This technique offers the same read performance as does reading from a single device, but many read I/O requests can be performed in parallel, making the overall subsystem read performance very high.

The write performance, however, is much worse. Because writing a small strip of data into a larger amount of data on a device is really a read-modify-write operation, writing becomes a time-consuming process.

In addition, because each stripe of data is protected by parity information, there is a significant level of data redundancy for high availability.

RAID 5 virtual disks provide high read performance and high availability at reasonable cost. They are optimum for use in applications requiring relatively low data transfer rates and having relatively high I/O read request rates.

read cache

A block of high-speed memory used by a controller to buffer data being read from storage devices by a host. A read cache increases the controller's effective device access speed by satisfying host read requests from its local cache memory when possible instead of from external storage devices.

The controller maintains in the cache copies of data recently requested by the host, and may fetch blocks of data ahead in anticipation that the controller will access the next sequential blocks. In a normal read cache, host write requests are handled as usual, without involving the caching mechanism.

read source

For HS-series controllers, a mirrored virtual disk option that controls the way data is read from the virtual disk's members. There are two Read Source options:

- Least Busy (default)---the normal, virtual disk member with the smallest I/O load is the target of all read operations.
- Round Robin---each Normal, virtual disk member is the target of a read operation in sequential membership order. No preference is given to any member.

rebuild rate

In the HSZ80 controllers, the rate at which the controller reconstructs a failed device on a spare. To devote more or fewer cycles to rebuilding a failed device on a spare, you can adjust the rate using a scale from 1 to 100. A rebuild rate of 100 rebuilds the device at the fastest rate possible.

reconstructing

A physical device state that indicates that the controller is regenerating a failed device's data onto a replacement device that is part of a redundant virtual disk.

All user data remains available during the reconstruction process, but some performance reduction occurs when a request requires access to a device that is being reconstructed.

reconstruction

Process of regenerating all of a failed member's data, writing it to a spare device, and incorporating the spare device as a redundant RAID virtual disk member.

All user data remains available during the reconstruction process, but some performance reduction occurs when a request requires access to a device that is being reconstructed.

reconstruction rate

The speed at which a failed member's data is regenerated. The rate is adjustable. See rebuild rate.

reduced

A virtual disk state that indicates that one member device is missing, failed, or has been physically removed from a virtual disk.

redundant RAID

Any RAID level that offers some level of data protection using redundant information. RAID 1 (mirroring), RAID 0+1 (striped mirroring), RAID 3 (striped parity), RAID 5 (striped parity), and RAID 3/5 (striped parity) virtual disks are all examples of virtual disks using redundant-RAID techniques.

regeneration

The process of recreating all or a portion of the data from a failed device using the surviving data and parity from the other members of the virtual disk.

In most controllers, data from all members remains available during regeneration, and the user is unaware that regeneration is occurring except for some slight reduction in performance.

Regeneration of an entire disk member is called reconstruction.

Data regeneration occurs if one of the following takes place:

- The controller detects a hard read error on a virtual disk member. In this case, the controller transparently corrects the data and continues the read operation.
- The controller detects that a member of a virtual disk has failed. In this case, the controller completely reconstructs the failed member.

SCSI (Small Computer System Interface)

An industry-standard, parallel, bus used to interconnect systems and physical devices.

SCSI device

A computer, a host adapter, a peripheral controller, or any storage element that can be attached to a SCSI bus.

SCSI ID

See target ID.

shadowing

See mirroring.

spare

A device state that indicates a device designated as a hot spare for a failed device within a redundant RAID virtual disk. When a virtual disk member fails, the controller automatically replaces it with a spare device from the pool of spare devices.
The term is also used to refer to a spare device itself.

storageset

HS-series controller term for device group.

storage subsystem

A controller and an array of physical devices attached to a host.

storage window

A window that provides you a graphical operator interface to a selected storage subsystem for configuration and monitoring. One of Client's components.

strip

In any form of RAID virtual disk using striping, host data is stored in pieces called strips. One strip is stored on each member device in the virtual disk. Taken together, the strips make up a stripe.

The strip size can be used in some controllers to tune the striped virtual disk for a specific application:

- If a virtual disk is set up with a small strip size compared to the average host request's size, the controller can break the host data up into strips and can perform its device accesses in parallel, using RAID 3 techniques. This optimizes the virtual disk for applications requiring high data transfer rates.
- If a virtual disk is set up with a large strip size compared to the average host request's size, the controller can perform many device read accesses in parallel, using RAID 5 techniques. This optimizes the virtual disk for applications requiring high I/O request rates.

Also called a chunk, segment, or stripe element.

stripe

In any form of RAID virtual disk using striping, host data is stored in pieces called strips. One strip is stored on each member device in the virtual disk. Taken together, the strips spread across the member devices make up a stripe.

strip size

The number of blocks of data that make up a strip.

In some controllers that offer multiple-RAID-level virtual disks such as RAID 3/5, the relationship between the strip size and the average host I/O request size determines how the controller accesses the devices in the virtual disk, and affects the request and data transfer performance of the subsystem.

striped virtual disk

Another name for a RAID 0.

striped mirrored virtual disk

Another name for a RAID 0 virtual disk.

striped parity virtual disk

Another name for a RAID 3, or RAID 5 virtual disk.

striping

A RAID technique in which host data is stored in pieces called strips. One strip is stored on each member device in the virtual disk. Taken together, the strips spread across the member devices make up a stripe.

subsystem

See storage subsystem.

supported device

A device tested as functionally compatible with a controller in an approved hardware and software configuration.

surviving controller

The controller in a dual-redundant pair of HS-series controllers that assumes service to its companion's devices when the companion controller fails.

target

A SCSI bus device. When one device addresses another device on a SCSI bus, it uses the target device's target ID to uniquely identify it.

target ID

The physical address a bus initiator uses to connect with a bus target. Each bus target is assigned a unique target address.
Also sometimes used to refer to a SCSI device ID itself.

TCP/IP

An acronym for Transport Control Protocol/Internet Protocol. The primary network protocol for the Internet. Uses IP addressing, where each network node has a unique network address.
Client uses the TCP/IP to communicate with Agent.

transfer rate

The speed of transfer of data over a SCSI bus. The transfer rate depends upon the both the bus speed and width. Typical transfer rates for some common SCSI bus types are:

SCSI Bus Type	Width	Maximum Transfer Rate
Fast Narrow	8 bits	10 MB/s
Fast Wide	16 bits	20 MB/s
Ultra	16 bits	40 MB/s

Transfer rates are usually expressed in units of megabytes per second.
Some controllers allow you to set the maximum transfer rate on either their host or device buses. This feature allows you to limit the rate in special situations, such as those that require long bus cabling.

unit

See virtual disk.

volume

See virtual disk.

volume set

See virtual disk.

virtual disk

A controller mapping physical devices on its device buses into logical storage units on its host bus. These logical storage units are called virtual disks and are accessible by the host. Each virtual disk has its own user-configurable operating characteristics.
Some controllers use RAID techniques to provide virtual disks with various combinations of cost, availability, and performance benefits.
Also called logical drives, volumes, volume sets, logical storage units, logical units, units, and LUNs.

warm swap

In some controllers, a feature that allows devices to be added, removed, or replaced while the subsystem remains operational, but with all activity on the controller's device buses stopped.

wide area network (WAN)

A network that spans a large, geographic area.

write-back cache

A cache configuration that increases the performance of host write requests. When the host requests a write operation, the controller writes the host's data first to the cache memory, completing the host's request quickly. It performs the slower operation of flushing the data to the external storage device at a later time. The host sees the write operation as complete when the data has reached the cache.

In most controllers using a write-back cache, the cache also increases performance with read cache techniques.

write only

A device state that indicates that the member was in the process of reconstructing and an error was found on one of the other members of the virtual disk.

write-through cache

A technique for handling host write requests in read caches. When the host requests a write operation, the controller writes data directly to the external storage device and updates the cache memory to ensure that the memory does not contain obsolete data. This technique increases the chances that future host read requests can be filled from the cache. The host sees the write operation as complete only after the external storage device has been updated.

In some controller fault conditions, your write-back cache resorts to write-through cache operation to protect your data.