StorageWorks™ Solutions

StorageWorks FDDI Server Installation Guide

Order Number: EK–HS1xx–IG. A01

This guide describes the procedures necessary to install, configure, and use the StorageWorks FDDI server. Mechanical cabinet configuration and power configuration as well as software installation and configuration are covered in this guide.

Digital Equipment Corporation
Maynard, Massachusetts

February 1995

NOTE: 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 the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation.

Any changes or modifications made to this equipment may void the user's authority to operate the equipment.

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.

This document was prepared using VAX DOCUMENT Version 2.1.

Contents

4 Customizing Server Operating System Parameters

5 Configuring the StorageWorks FDDI Server Storage Devices

6 Integrating the Server into the VMScluster System

7 Supporting and Operating the StorageWorks FDDI Server

8 System Software Maintenance

A Working with RAIDsets

B Device Channel Processor Command Line Interpreter

Examples

Figures

Tables

Preface

This manual describes the procedures necessary to install, configure, and use the StorageWorks FDDI server. Mechanical cabinet configuration and power configuration as well as software installation and configuration are covered in this manual.

Intended Audience

This manual is intended for use by customers and Digital™ Multivendor Customer Services personnel responsible for installing and configuring a StorageWorks FDDI server.

Structure

This manual is organized as follows:

Chapter 1	Provides an overview of the StorageWorks FDDI servers and a description of their major components.
Chapter 2	Provides site preparation information.
Chapter 3	Describes the unpacking and placement of a StorageWorks FDDI server.
Chapter 4	Describes installing StorageWorks FDDI server software.
Chapter 5	Describes configuring the StorageWorks FDDI server storage devices.
Chapter 6	Describes how to connect the server to the FDDI interconnect.
Chapter 7	Describes supporting and operating the StorageWorks FDDI server in the VMScluster system.
Chapter 8	Describes system software maintenance tasks.
Appendix A	Describes RAIDset fundamentals.
Appendix B	Describes the device channel processor Command Language Interpreter commands.
Glossary	Defines the acronyms and specialized terms used in the StorageWorks environment.

Related Documents

Table 1 lists the StorageWorks-related user documents organized by use, system, or product. Table 2 lists the OpenVMS Alpha™ operating system-related documents.

Table 1 StorageWorks Related Documentation

Document Title	Order Number
StorageWorks Primary Publications	
StorageWorks Solutions Configuration Guide	EK–BA350–CG
StorageWorks Solutions Shelf and SBB User's Guide	EK–BA350–UG
StorageWorks Solutions BA356–SB 16-Bit Shelf User's Guide	EK–BA356–UG
StorageWorks Enclosures	
StorageWorks Solutions SW500 and SW800 Cabinet Metric Shelf Bracket Kit Installation Guide	EK–35XRD–IG
StorageWorks RETMA Shelf Rail Kit Installation Guide	EK–35XRB–IG
StorageWorks Solutions SW800-Series Data Center Cabinet Cable Distribution Unit Installation Guide	EK–SWCDU–IS
StorageWorks Solutions SW800-Series Data Center Cabinet Installation and User's Guide	EK–SW800–IG
Storage Devices	
StorageWorks Building Blocks User's Guide	EK–SBB35–UG
StorageWorks Solutions 3½-Inch Storage Device Installation Guide	EK–MC350–IG
StorageWorks Solutions 5¼-Inch Storage Device Installation Guide	EK–MC525–IG
General Reference Publications	
Digital Systems and Options Catalog	†
Small Computer System Interface, An Overview	EK–SCSIS–OV
Small Computer System Interface, A Developer's Guide	EK–SCSIS–DK

† Available from your Digital account representative.

Table 2 OpenVMS Alpha Operating System Related Documentation

Document Title	Order Number
A Comparison of System Management on OpenVMS Alpha and OpenVMS VAX	AA–PV71B–TE
OpenVMS System Manager's Manual: Essentials	AA–PV5MB–TE
DECnet for OpenVMS Guide to Networking	AA–PV5ZA–TK
OpenVMS License Management Utility Manual	AA–PVXUB–TK
VMScluster Systems for OpenVMS	AA–PV5WB–TK
OpenVMS Alpha Version 6.1 Upgrade and Installation Manual	AA–PV6XB–TE
Guide to OpenVMS Performance Management	AA–PV5XA–TK
OpenVMS System Manager's Manual: Tuning, Monitoring and Complex Systems	AA–PV5NB–TK

Documentation Conventions

The following conventions are used in this manual:

boldface type	Boldface type indicates the first instance of terms being defined in text, in the glossary, or both.
italic type	Italic type indicates emphasis and complete manual titles. In the glossary, italic type also is used to indicate cross-references.

Manufacturer's Declarations

Following are manufacturer's declarations applicable to the StorageWorks FDDI server:

CAUTION

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 die Benutzer für entsprechende Gegenmaßnahmen verantwortlich sind.

ATTENTION !

Ceci est un produit de Classe A. Dans un environment domestique, ce produit risque de créer des interférences radiélectriques, il appartiendra alors à l´utilisateur de prendre les mesures spécifiques appropriées.

Note

The equipment described in this manual is listed by the Underwriters Laboratories Incorporated and bears the UL Listing mark. The StorageWorks FDDI server cabinets also are certified by the Canadian Standards Association and TUV Product Service GmbH and bear both the CSA certification and TUV GS marks.

Table 3 Acoustics—Preliminary Declared Values per ISO 9296 and ISO 7779

	Sound Power Level L_{WAd}, B‡		Sound Pressure Level L_{pAm}, dBA (Bystander Positions)	
Product†	**Idle**	**Operate**	**Idle**	**Operate**
HS1 with only cabinet fans operating	7.6	7.6	59	59
HS1 with 2 BA350–MA shelves and 12 BA350–SA shelves, each containing 6 RZ26–VA disk drives	7.6	7.6	59	59
Per device when installed in HS1				
BA350–SA shelf containing 6 RZ26–VA disk drives	5.7	5.7	39	39
BA350–MA shelf	5.6	5.6	39	39

† Current values for specific configurations are available from Digital representatives.
‡ 1 B = 10 dBA.

Note

Table 4 is a translation of the English language specifications in Table 3 into the German language.

Table 4 Schallemissionswerte—Vorläufige Werteangaben nach ISO 9296 und ISO 7779/DIN EN27779

	Schalleistungspegel L_{WAd}, B‡		Schalldruckpegel L_{pAm}, dBA (Beistehende Position)	
Gerät†	**Leerlauf**	**Betrieb**	**Leerlauf**	**Betrieb**
HS1 nur mit kabinett Lüftern in Betrieb	7,6	7,6	59	59
HS1 mit 2 BA350–MA shelves und 12 BA350–SA shelves, jedes bestückt mit 6 RZ26–VA disk drives	7,6	7,6	59	59
Pro Gerät installiert im HS1				
BA350–SA shelf mit 6 RZ26–VA disk drives	5,7	5,7	39	39
BA350–MA shelf	5,6	5,6	39	39

† Aktuelle Werte für spezielle Ausrüstungsstufen sind über die Digital Equipment Vertretungen erhältlich.
‡ 1 B = 10 dBA.

1
Introduction

This chapter presents a system overview of the StorageWorks FDDI server, describes each of its major components, and lists the upgrade options. The final section of this chapter provides information on how to use the remainder of the guide.

1.1 System Overview

The StorageWorks FDDI servers integrate Digital's Alpha technology with StorageWorks modular design to meet the storage requirements of large FDDI-based VMScluster™ systems. The StorageWorks FDDI servers support a wide range of StorageWorks solid state disk, magnetic disk, tape, optical, and loader devices, for configuring the exact FDDI storage solution to meet your application's needs. The StorageWorks FDDI server family extends the StorageWorks line with both standard and highly available servers to meet the I/O needs of even the largest FDDI-based VMScluster systems. The StorageWorks FDDI servers currently available are the following:

- HS111–AA, AB
- HS121–AA, AB
- HS110–AA

Each StorageWorks FDDI server comprises part of a FDDI-based VMScluster environment.

1.1.1 HS111 Description

Figure 1–1 shows a conceptual model of the HS111 StorageWorks FDDI server. It provides managed storage services to as many as 95 other host computers or workstations in a VMScluster system. The model HS111 StorageWorks FDDI server consists of a single server processor with a FDDI interface, a device channel processor (HS1CP) capable of connecting the server processor to as many as 42 StorageWorks devices, and your choice of supported StorageWorks devices and packaging options. (See Software Product Description for the HS1CP for a complete list of supported StorageWorks devices.) Using the currently available StorageWorks disk devices, as much as 88 GB of online storage per server becomes available. The model HS111 StorageWorks FDDI server and its attached storage devices are housed in a single StorageWorks SW800 cabinet.

The HS111's server processor runs the OpenVMS Alpha operating system. This means that, in addition to serving online storage to hosts, the HS111 StorageWorks FDDI server can use the full range of OpenVMS storage management software to provide a optimally managed storage environment. For enhanced data protection, a user-configurable RAID capability is available for the HS111's device channel processor.

Figure 1–1 HS111 Diagram

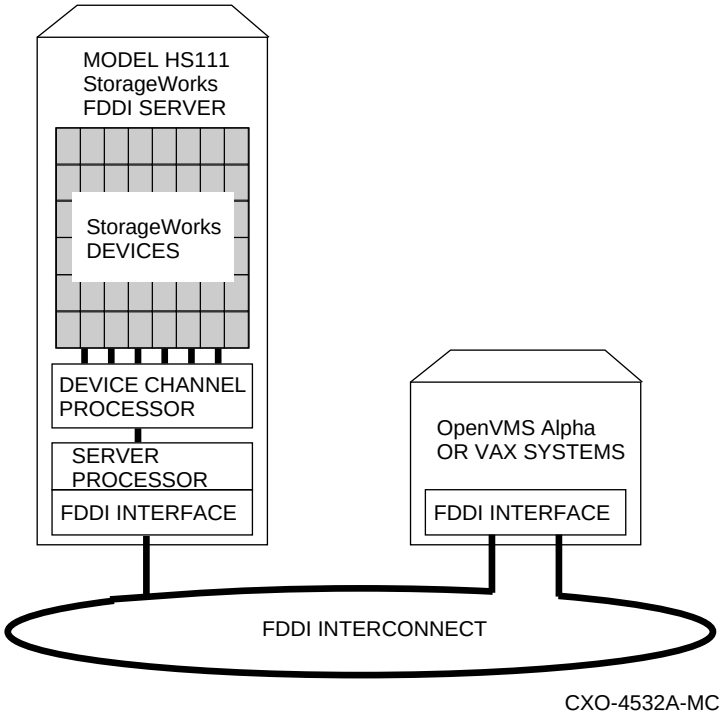

1.1.2 HS121 Description

Figure 1–2 shows a conceptual model of the HS121 StorageWorks FDDI server. It is a complete high-available FDDI storage solution, offering full protection against any single component failure. The HS121 StorageWorks FDDI server includes redundant server processors, two paths to the FDDI bus, two paths to every storage device, and redundant power and cooling. Expansion space is available for the addition of dual redundant device channel processors on each path. Redundant power controllers are also a standard feature. The full range of StorageWorks redundant power and cooling features for devices is also available.

The basic HS121 StorageWorks FDDI server provides fully redundant access to up to 36 StorageWorks device connections. (See Software Product Description for the HS1CP for a complete list of supported StorageWorks devices.) Expansion space is available for the addition of a second pair of device channel processors, which increases the capacity to 72 device connections. Using the currently available StorageWorks disk devices, as much as 151 GB of online storage per server with fully redundant FDDI access becomes available. All data protection and storage management options are available for the HS121 StorageWorks FDDI server. The model HS121 StorageWorks FDDI server and its attached storage devices are housed in a single StorageWorks SW800 cabinet.

During normal operation, I/O performance is enhanced by balancing the I/O load across the redundant system components in a customer-specifiable way. If a component failure occurs, its redundant partner takes over for it, providing continued service until the failed component can be repaired or replaced. All HS121 StorageWorks FDDI server components can be hot swapped, or replaced without interrupting I/O service.

Figure 1–2 HS121 Diagram

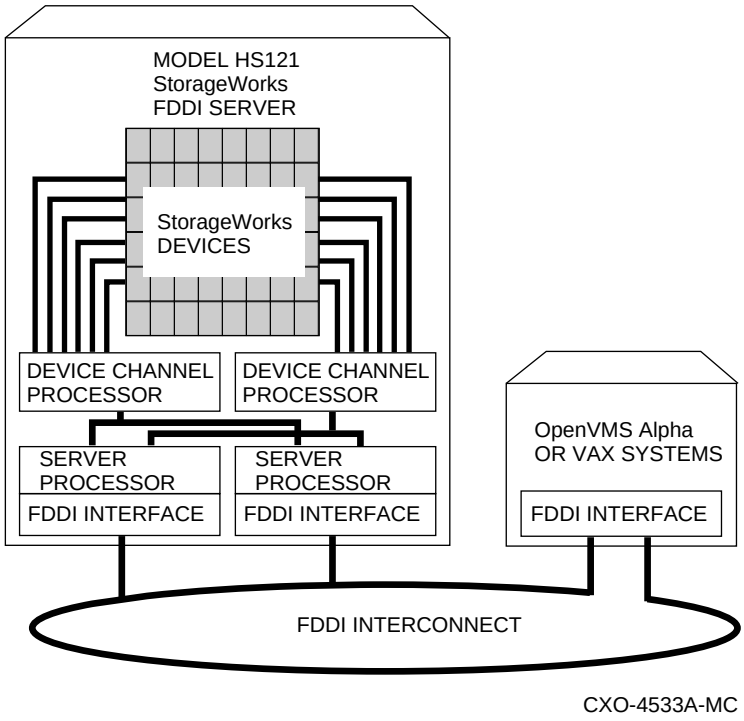

1.1.3 HS110 Description

The HS110 contains all the major components of a HS111 except for the SW800 cabinet and the StorageWorks devices. The HS110 provides for the conversion of an existing SW800 Cabinet to a StorageWorks FDDI server. Refer to the *StorageWorks FDDI server Upgrade Manual* for conversion details.

1.1.4 Major Components

Figure 1–3 shows a StorageWorks FDDI server calling out each of the major components and subcomponents. The major components are as follows:

- StorageWorks FDDI server processor
- device channel processor (HS1CP)
- StorageWorks storage devices

The subcomponents are explained in the sections that follow Figure 1–3.

Figure 1–3 StorageWorks FDDI Server Major Components

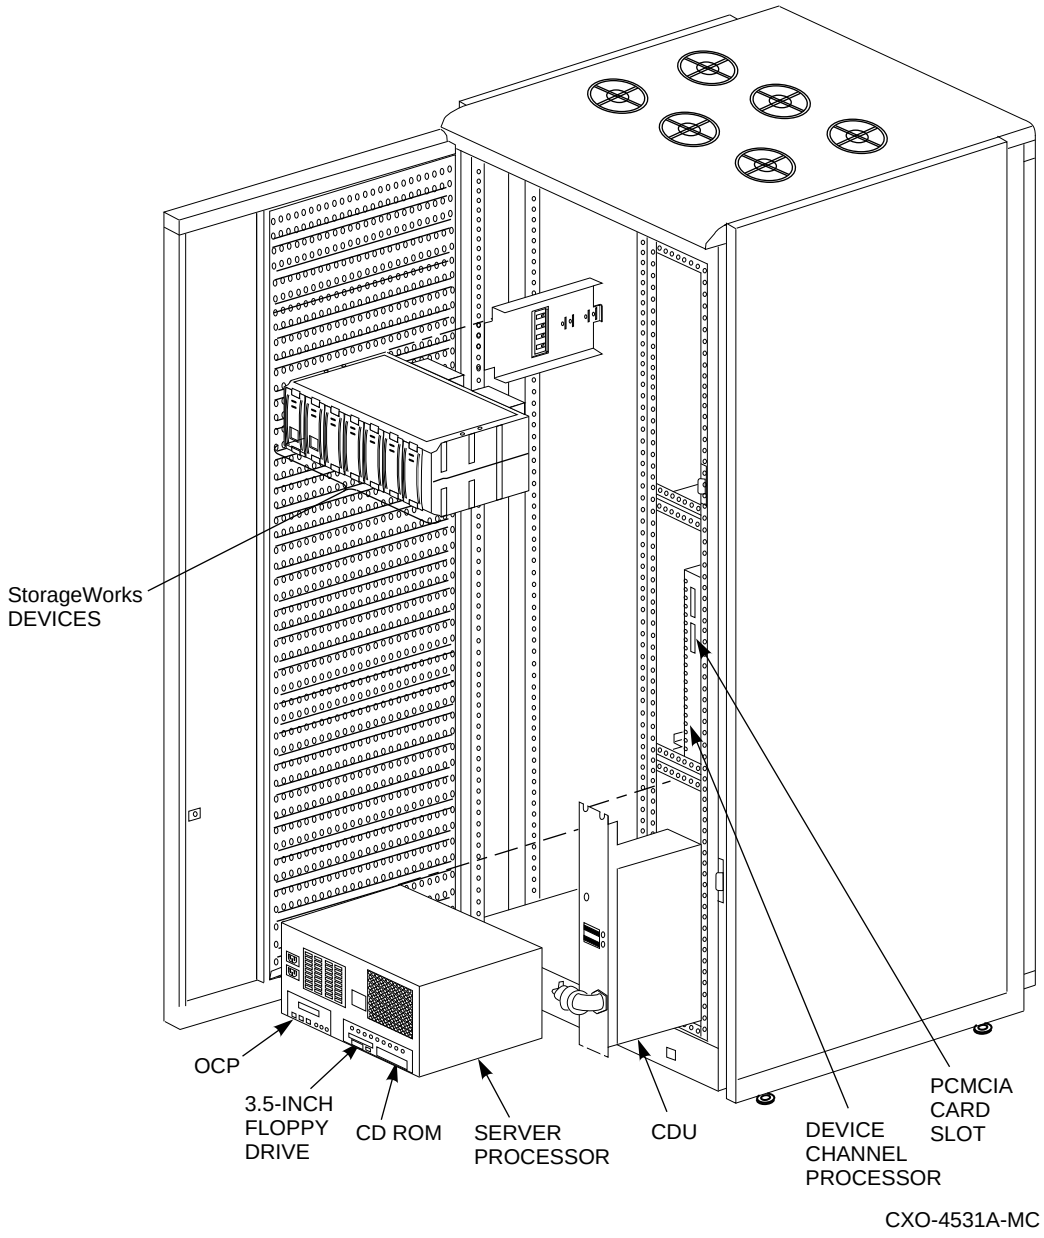

1.1.4.1 StorageWorks FDDI Server Processor

The StorageWorks FDDI server processor combines Digital's advanced Alpha technology with 64 MB RAM, to provide the internal performance required to handle the I/O load of FDDI-based VMScluster systems. A single server processor with dual HS1CPs can process up to 1,500 I/O requests per second and transfer up to 12 MB per second to and from the FDDI interconnect.

Each server processor supports the full range of OpenVMS standard and optional storage management utilities, providing a managed storage environment that the OpenVMS system administrator will find familiar.

In the model HS121 StorageWorks FDDI server, two server processors share the I/O load under normal circumstances. When a failure occurs, the operational server processor takes over the load of its failed partner, and continues to deliver data from all devices until repair can be effected. Because server processors are separately powered and cooled, a power supply or cooling failure does not interrupt data access.

The following paragraphs describe the storage components of the server processor, which are the:

- System disk
- CD–ROM
- 3.5-Inch floppy drive

System Disk

The system disk contains the OpenVMS Alpha Operating System, the Software Customization Procedure (SCP), various relevant documents, and required layered software components. It consists of an RZ28 disk drive located inside the server processor.

CD–ROM

The CD–ROM is used for system disk rebuild operations if there is no system disk backup or for updating the server operating system to a new version of OpenVMS Alpha operating system. The CD–ROM contains the following:

- Files to install the OpenVMS Alpha operating system
- A bootable subset of the OpenVMS Alpha operating system that permits a limited system boot directly from the CD–ROM
- OpenVMS Alpha volume shadowing interoperability patch kit (I14Y)
- The Software Customization Procedure (SCP)
- OpenVMS Alpha documentation in Bookreader, PostScript®, and ASCII text formats
- Applicable Software Product Descriptions (SPDs)

For a description of the CD–ROM hardware and contents of the CD–ROM media, see Section 1.2.

3.5-Inch Floppy Drive

The 3.5-inch floppy drive is used to run the EISA Configuration Utility (ECU) whenever configuration changes are made to the server processor's internal bus, such as during a hardware upgrade.

1.1.4.2 Device Channel Processor

The device channel processor can provide nonredundant connections between a server processor and up to 42 industry-standard SCSI–2 devices on six independent fast (10 MB/second) SCSI–2 buses. The device channel processor supports a wide range of StorageWorks magnetic, solid state, and optical disks, tapes, and media loader devices.

To support the I/O demands on this much storage, each device channel processor is capable of processing up to 1000 I/O requests per second. The device channel processor firmware automatically balances stripeset I/O across member disks as well as enabling the following features:

- Redundant access
- Read/Write caching
- RAIDset configuration support (optional feature)

Each of these features is described later in this section.

The device channel processor firmware resides on a PCMCIA card that ships with the StorageWorks FDDI server. The PCMCIA card is located in the slot located on the front of the device channel processor.

Redundant Access

In model HS111 StorageWorks FDDI servers, two device channel processors can optionally be configured as a dual-redundant pair connected to the same storage devices. In this configuration, each device channel processor can assume control of all storage devices in the event of its partner's failure, including flushing any unsaved data from the partner's write cache to storage media. This feature is standard on the model HS121 StorageWorks FDDI server.

Read/Write Cache

Each device channel processor includes a 32 MB nonvolatile read/write cache that reduces I/O request execution times. Firmware allocates the cache dynamically based on actual I/O activity. For optimal resource utilization, management facilities allow you to enable or disable both read and write caching for each storage unit attached to a device channel processor. In the event of a power failure, on-board batteries keep write cache contents intact for up to 100 hours, so no data is lost.

Note

The write-back cache module installed in your StorageWorks FDDI server contains batteries that were completely charged at the factory. It is normal for these batteries to discharge slightly in shipment. The server's write-back cache and RAID features require fully-charged batteries to maintain absolute data integrity. After installation, these advanced features may not be available until the batteries have had an opportunity to completely recharge. The charging process may take up to 4 hours to complete.

RAIDset Configuration Support

As an optional upgrade, the device channel processor provides for Digital's implementation of parity RAID technology, which dynamically adjusts to the I/O workload, optimizing I/O performance in both high data transfer rate and high transaction rate environments. Using distributed parity, the StorageWorks FDDI server RAID option economically provides continued access to data in the event of a disk failure. In the fully redundant model HS121 StorageWorks FDDI server, this increases data availability to a level comparable to that of shadowing at significantly lower cost.

1.1.4.3 StorageWorks Storage Devices

Because StorageWorks FDDI servers support a wide selection of Digital's StorageWorks components, storage subsystems can be custom configured to meet virtually any application need. For increased reliability, the SW800 cabinet that houses both StorageWorks FDDI server models supports the full suite of StorageWorks redundant power and cooling features.

StorageWorks enclosure components and supported storage devices are all available on a configure-to-order basis. Consult the Software Product Description for the HS1CP for an up-to-date list of supported StorageWorks devices.

1.1.5 Available Upgrades

The StorageWorks FDDI servers have a number of applicable upgrades: some are StorageWorks FDDI server specific, others are applicable to the SW800 enclosure or StorageWorks shelves and support the StorageWorks FDDI server upgrades. Each upgrade is briefly described in the Table 1–1. A complete description of each upgrade can be found in the manual titled *Upgrading the StorageWorks FDDI Server*.

Table 1–1 StorageWorks FDDI Server Upgrades

Upgrade Order Number	Description
StorageWorks FDDI Server-Specific Upgrades	
HS110–AA	Upgrades an HS111–AA, AB to an HS121–AA, AB or upgrades an existing SW800 cabinet to an HS111. If updating an HS111–AA, AB, the upgrade provides redundant access to the installed StorageWorks storage devices. If updating an SW800 cabinet containing only StorageWorks storage devices, the update provides nonredundant access to the installed StorageWorks devices.
HS1MM–AA	StorageWorks FDDI server 64 MB server memory expansion
HS1CP–AF	StorageWorks FDDI server device channel processor including 32 MB nonvolatile read/write cache
HS1AD–AA	StorageWorks FDDI server processor to device channel processor bus adapter
HS1PW–AA	StorageWorks FDDI server processor redundant power supply
Supporting Upgrades	
BA350–JA	StorageWorks device shelf with mounting kit and power supply. Provide mounting for up to seven StorageWorks devices
BA35X–HA	StorageWorks device shelf power supply
SW8XP–AX	Additional power controller for the SW800 cabinet
QL–3HYA9–AA	device channel processor RAID firmware license (one required for each device channel processor attached to a RAID array)

1.1.6 StorageWorks FDDI Server Specifications

Table 1–2 provides the StorageWorks FDDI server specifications.

Table 1–2 StorageWorks FDDI Server Specifications

	Model HS111	Model HS121
Storage Capacity		
Maximum number of disks	42 (36 if dual-redundant device channel processors are installed)	36 (36 to 72 if two dual-redundant device channel processor pairs are installed)
Maximum disk capacity using RZ28 disks	88 GB (75 GB if dual-redundant device channel processors are installed)	75 GB (75 to 151 GB if two dual-redundant device channel processor pairs are installed)
Maximum number of tapes	See the *StorageWorks Solutions Configuration Guide* (EK–BA350–CG) for configuration guidelines.	
Maximum number of Solid State disks	See the *StorageWorks Solutions Configuration Guide* (EK–BA350–CG) for configuration guidelines.	
I/O Performance		
Maximum single-block disk I/O request rate	1,500	2,200
Maximum data transfer rate (to/from FDDI interconnect)	12 MB/second	12 MB/second
I/O response time for typical OpenVMS load	12 ms (cache disabled) 7 ms (read and write cache enabled)	12 ms (cache disabled) 7 ms (read and write cache enabled)
Power Requirements		
Voltage	100 to 240 V ac	100 to 240 V ac
Frequency	50/60 Hz	50/60 Hz
Device channel processor battery backup	Standard	Standard
Redundant Power	Optional	Standard
Environmental		
Operating temperature	10 C to 40 C	10 C to 40 C
Non-operating temperature	-40 C to 66 C	-40 C to 66 C
Relative humidity	10% to 80%	10% to 80%
Altitude	Sea level to 2400 m (8000 ft)	Sea level to 2400 m (8000 ft)

(continued on next page)

Table 1–2 (Cont.) StorageWorks FDDI Server Specifications

	Model HS111	Model HS121
Physical Dimensions		
Height	1700 mm (67 in)	1700 mm (67 in)
Width	800 (31 in)	800 mm (31 in)
Depth	875 (34.5 in)	875mm (34.5 in)
Weight (without devices)	310 kg (685 lbs)	355 kg (780 lbs)

1.2 Using the CD–ROM

The StorageWorks FDDI server ships with the current version of OpenVMS Alpha software installed on the system disk; consequently, the CD–ROM is not needed for installation or day-to-day operation. However, it is needed to perform upgrades of the system disk or to restore or rebuild the system disk in the event of a system disk failure if no backup is available. It is also useful in performing complete backups of the server's system disk. A simple standalone-mode menu, executed from the CD–ROM, offers you a quick and easy way to perform the appropriate task.

The CD–ROM contains the following:

- A bootable subset of the OpenVMS Alpha operating system that permits a limited system boot directly from the CD–ROM
- Volume shadowing interoperability patch kit (I14Y) for OpenVMS Alpha and OpenVMS VAX operating systems
- OpenVMS Alpha documentation in Bookreader, PostScript, and ASCII text formats
- Applicable Software Product Descriptions (SPDs)

1.2.1 Handling the CD–ROM

The CD–ROM is enclosed in an antistatic jewel box. Be careful to handle the disc only by the edges; fingerprints or smudges should be cleaned using a soft cloth.

To insert the CD–ROM into the drive, press the button on the right side of the drive. The CD tray should then extend out; it may require pulling to its fully extended position. Place the CDROM in the tray with the printed side facing up; no disc caddy is required. With the CDROM seated properly in the tray, press the drive button again; the tray will retract back into the drive, and the drive light will momentarily illuminate as the drive establishes the presence of the disc.

To remove the CDROM, press the drive button again, and replace the CDROM in its protective jewel box.

1.2.2 Accessing Documentation

Note

The documentation is available on both the server's system disk and the CD–ROM. Their directory structures are the same. Use whichever source you prefer, though the system disk is suggested, because it is accessible during normal operation.

Prior to accessing documentation on the CD–ROM device, you must load the CD–ROM disc into the device, then mount it.

See Section 6.2.1 for more information on how to access the documentation.

Documentation for the StorageWorks FDDI server is available in the following formats:

- Files with the filetype .PS are PostScript files that are printed on a PostScript printer.
- Files with the filetype .TXT are ASCII text files that are printed or viewed on a character cell terminal.
- Files with the filetype .DECW$BOOK are Bookreader files. You must have DECwindows Motif software installed to read these files with Bookreader. For more information on Bookreader, refer to the *DECwindows Motif for OpenVMS Applications Guide*.

The .PS and .TXT documentation for OpenVMS Alpha software is usually found in the [DOCUMENTATION.V0*xx*] directory, where *xx* represents the OpenVMS operating system software version number. For example, for Version 6.1, the directory is [DOCUMENTATION.V061].

Similarly, the Bookreader documentation for the OpenVMS Alpha software is found in the [DOCUMENTATION.V0*xx*.DECW$BOOK] directory.

1.3 How to Use This Guide

Before your StorageWorks FDDI server is ready to use, you must understand and perform the installation tasks outlined in this guide. The tasks and references to them are listed in Table 1–3.

Table 1–3 Installation Tasks

When you . . .	You must . . .	See . . .
Receive the StorageWorks FDDI server	Understand the product description and installation tasks before you proceed	Chapter 1, Introduction.
Unpacking and Installing the Server Hardware		
Prepare the site	Perform environment and service checks to make sure your site is suitable for installing a StorageWorks FDDI server	Chapter 2, Preparing Your Site.
Install the hardware	Remove the StorageWorks FDDI server from its shipping container, place it where you want, connect a terminal and apply power	Chapter 3, Unpacking, Placing, and Applying Power to the StorageWorks FDDI server.
Configuring Server Software and Devices		
First boot the StorageWorks FDDI server processors	Consider VMScluster system issues and execute the Software Customization Procedure to set processor parameters	Chapter 4, Customizing Server Operating System Parameters.
Configure the StorageWorks storage system	Determine and set up the configurations that will meet your needs and configure the system	Chapter 5, Configuring the StorageWorks FDDI Server Storage Devices.
Integrating the Server into the VMScluster System		
Prepare the StorageWorks FDDI server to join the VMScluster system	Connect the StorageWorks FDDI server to the FDDI interconnect and perform other related tasks	Chapter 6, Integrating the Server into the VMScluster System.
Supporting and Operating the Server		
Determine StorageWorks FDDI server processor system management strategies	Consider installing additional software	Chapter 7, Supporting and Operating the StorageWorks FDDI Server.
Receive new StorageWorks FDDI server system software, or recover from a system disk failure	Upgrade or restore data from the CD–ROM	Chapter 8, System Software Maintenance.

2

Preparing Your Site

You must ensure that the site selected for installing a StorageWorks FDDI server satisfies requirements for power, grounding, safety, and service. This chapter discusses preparing your site to meet those requirements.

2.1 Power Requirements

The StorageWorks FDDI server is intended for installation in Class A computer room environments. Before installing the StorageWorks FDDI server, make sure that your site meets the following power requirements:

- The primary power source can supply the required amount of ac power, as specified in the specifications table (See Chapter 1).
- The proper power receptacles are installed for the cabinet. Figure 2–1 specifies the primary power plugs supplied with both 60 Hz and 50 Hz power versions of the StorageWorks FDDI server.

Figure 2–1 StorageWorks FDDI Server Primary Power Plugs

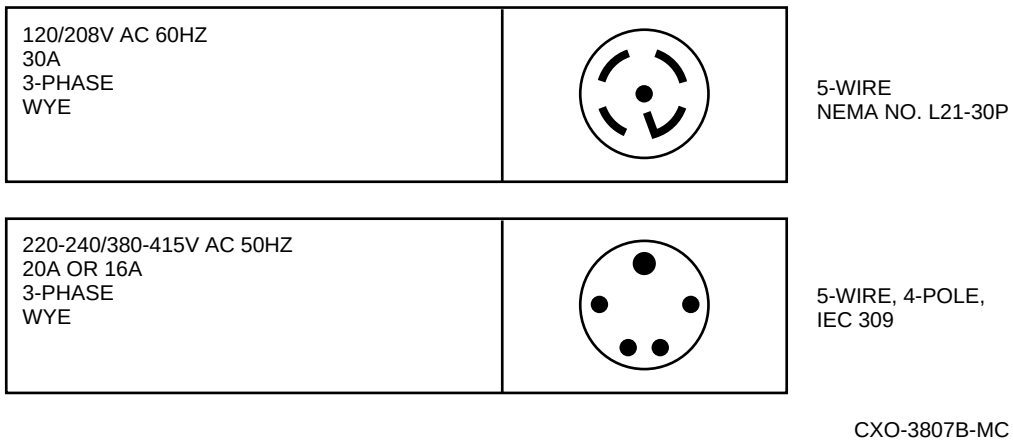

2.2 Grounding Requirements

The StorageWorks FDDI server usually is connected to other equipment by one or more FDDI interconnect cables. For both safe and reliable operation, the cabinet and other equipment require proper grounding.

WARNING

If cabinets are not connected to a common ground, there is a potential for personal injury as a result of electric shock.

Make sure that site power distribution systems meet local electrical codes prior to the installation of StorageWorks FDDI server.

To make sure that the power distribution system will perform satisfactorily, a power system survey should be done before installation. The following areas should be investigated:

- Do all outlets have power ground connections?
- Do the power cords on all equipment at the site have grounding prongs?
- Are all power outlet neutral connections isolated from ground?
- Are the grounds for all outlets connected to the same power distribution panel?
- Are all devices that are connected to the same breaker as the StorageWorks FDDI server UL or IEC approved?

CAUTION

If there is a deficiency found in any area during the power survey, a qualified electrician must correct it before installation may begin. Failure to resolve power survey deficiencies before installing the equipment may result in personal injury as a result of electric shock.

If no problems are found during the survey, the site grounding system may be considered to be adequate for safe and reliable StorageWorks FDDI server operation.

2.3 Safety and Service Requirements

The following general safety and service requirements refer to the floor location chosen for placing the StorageWorks FDDI server:

- The site floor must safely bear the weight of the cabinet, as specified in Chapter 1. The entire weight of the cabinet is supported by the four leveler feet when the cabinet is installed in its final position.
- The space around the cabinet must allow for opening the front and rear doors, for accessing cables, and for adequate airflow. See Figure 2–2 for specific space requirements.

Figure 2–2 Minimum Installation Clearances

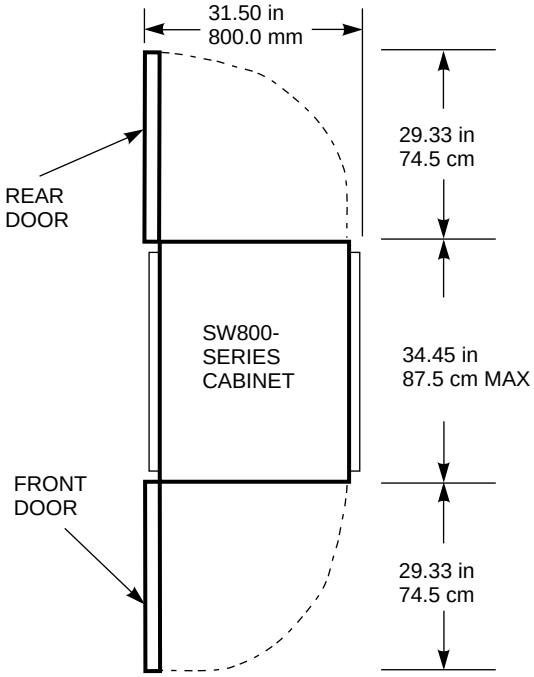

CXO-3805A-MC

- If the cabinet is positioned next to other cabinets, there must be sufficient service loop in any connecting cables to allow the cabinet to be moved out for access.

Note

The StorageWorks FDDI server is not designed to be fastened to adjacent cabinets.

3

Unpacking, Placing, and Applying Power to the StorageWorks FDDI server

This chapter describes unpacking and placing the StorageWorks FDDI server followed by descriptions of connecting a terminal to, and applying power to the StorageWorks FDDI server.

3.1 Unpacking the Cabinet

The StorageWorks FDDI server arrives packed in a corrugated carton attached to a wooden shipping pallet, as shown in Figure 3–1. Unpack the cabinet as follows:

Note

Before unpacking the equipment, inspect the shipping carton for signs of external damage. Report any damage to the local carrier and to Digital Multivendor Customer Services or your local Digital sales office.

1. Remove the cover, the fasteners, and the corrugated board from the pallet.
2. Remove the cartons containing the ramp set and skirt kit and set them aside.
3. Cut the shipping straps. Some cabinets are packaged in a plastic or barrier bag. If the cabinet arrives in a plastic bag, leave the bag in place until the cabinet has adjusted to the local temperature and humidity, as follows:

CAUTION

Failure to thermally stabilize storage subsystems may damage drive media or associated electronics when the unit is turned on. Environmental stabilization begins when the equipment is placed in the room in which it is to be installed.

- To ensure proper operation of Digital storage devices, the StorageWorks building block (SBB) temperature must be within 18°C to 29°C (65°F to 85°F). Table 3–1 specifies the time required to thermally stabilize SBBs based on the ambient shipping temperature.

Table 3–1 Thermal Stabilization Specifications

Ambient Temperature Range °C	Ambient Temperature Range °F	Minimum Stabilization Time
60 to 66	140 to 151	3 hours
50 to 59	122 to 139	2 hours
40 to 49	104 to 121	1 hour
30 to 39	86 to 103	30 minutes
18 to 29	65 to 85	None
10 to 17	50 to 64	30 minutes
0 to 9	32 to 49	1 hour
–10 to –1	14 to 31	2 hours
–20 to –11	–4 to 13	3 hours
–30 to –21	–22 to –5	4 hours
–40 to –31	–40 to –21	5 hours

- If condensation *is visible* on the outside of the storage device, stabilize the device and the SBB in the operating environment for 6 hours or until the condensation is no longer visible, whichever is longer.
- If condensation *is not visible* on the outside of the storage device, thermally stabilize the device for the amount of time specified in Table 3–1.

4. Once the cabinet is unpacked, examine the front and rear doors, right and left side panels, top panel, and undercarriage for any apparent damage. Report such problems immediately to Digital Multivendor Customer Services or your local Digital sales office.
5. Retain the shipping container and all packing materials.

Figure 3–1 Shipping Container Contents

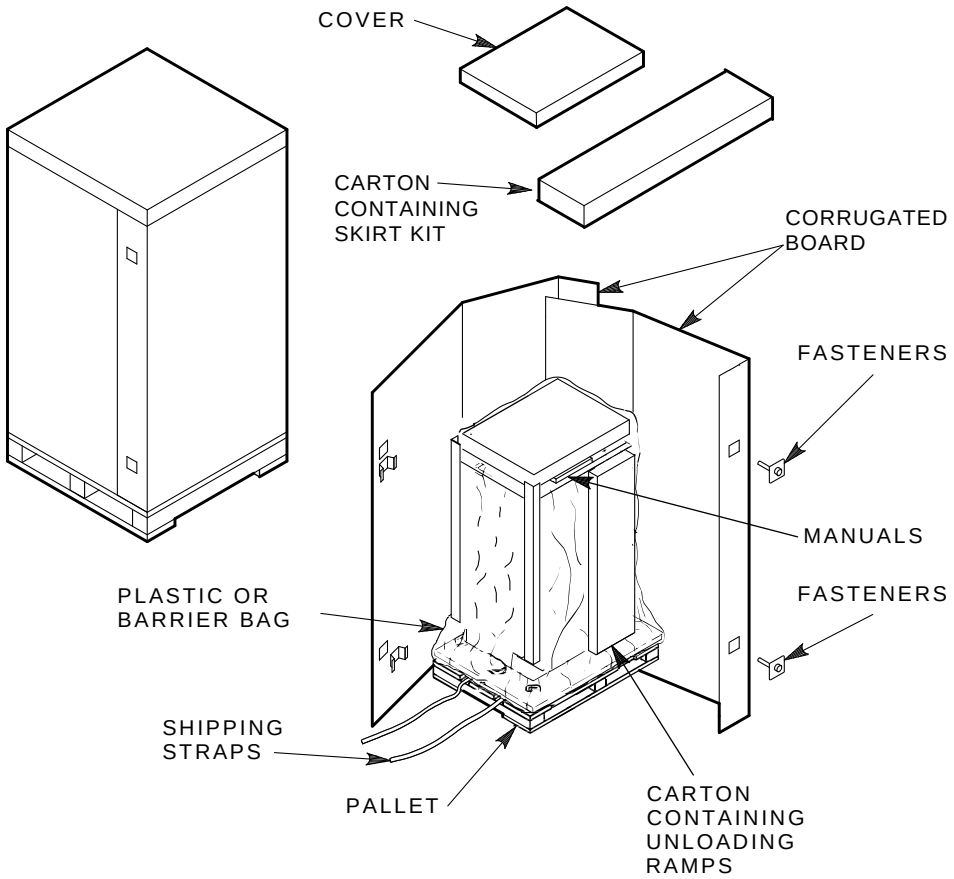

3.2 Removing the Cabinet from the Pallet

WARNING

Three people are required to remove the cabinet from the shipping pallet. Failure to use sufficient personnel and correct safety precautions can result in injury and equipment damage. All personnel should wear safety glasses.

Use the following procedure to remove the cabinet from the shipping pallet:

1. Remove any packing material remaining on the pallet.
2. Remove and inspect the two unloading ramps. The ramps, ramp side rails, and metal hardware should be inspected for the following defects:
   - Cracks more than 25 percent of the ramp depth, either across or lengthwise on the ramp
   - Knots or knotholes going through the thickness of the ramp and greater than 50 percent of the ramp width
   - Loose, missing, or broken ramp side rails
   - Loose, missing, or bent metal hardware

   If any of these defects exist, do not use the ramp. Investigate alternate means of removing the cabinet or order a new ramp. (The part number for the ramp set is 99–08897–05.)
3. Attach the ramps by fitting the metal prongs into the holes on the pallet, as shown in Figure 3–2. Make sure that the arrows on the ramps match up with the arrows on the pallet.
4. Extend the ramps to their full length.
5. See Figure 3–3 for the location of the shipping bolts. Remove the bolts.
6. Remove the shipping brackets, shown in Figure 3–3, from the cabinet leveler feet and set aside.
7. Loosen the leveler locking nuts and screw the four cabinet leveler feet all the way up into the cabinet.

   WARNING

   The leveler feet must be raised fully for the cabinet to roll easily down the unloading ramps. Failure to do so can result in the cabinet tipping off the pallet or ramp.

8. Carefully roll the cabinet off the pallet and down the ramps to the floor as shown in Figure 3–4.

Figure 3–2 Shipping Pallet Ramp Installation

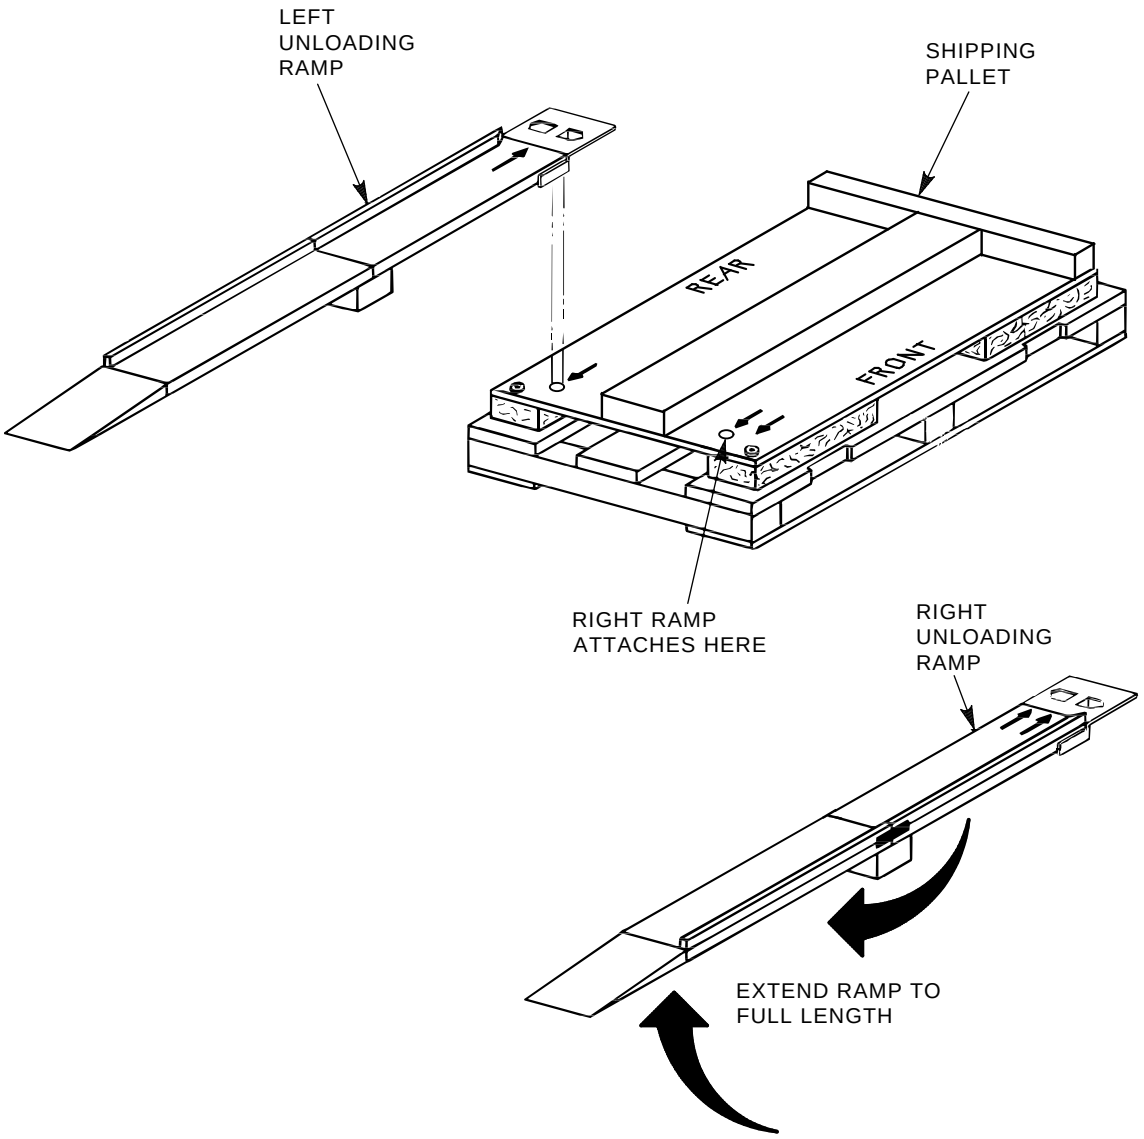

Figure 3–3 Shipping Bolts and Brackets

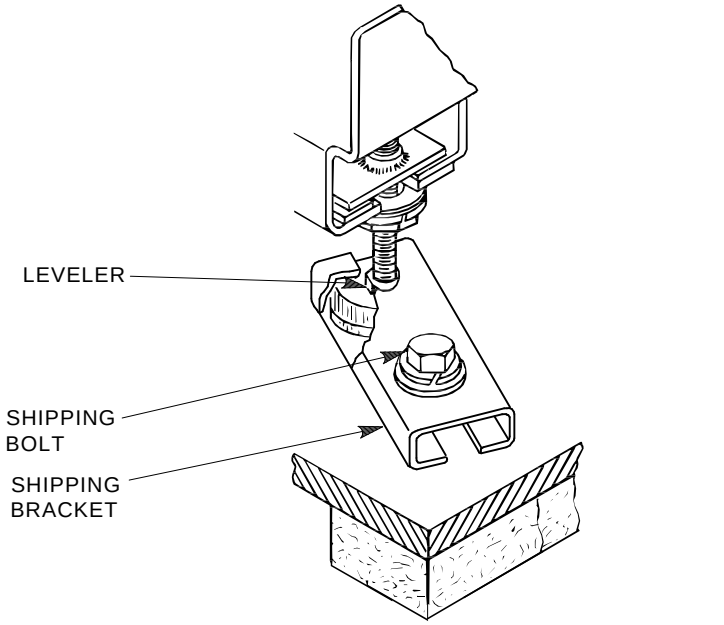

Figure 3–4 Removing the Cabinet from the Pallet

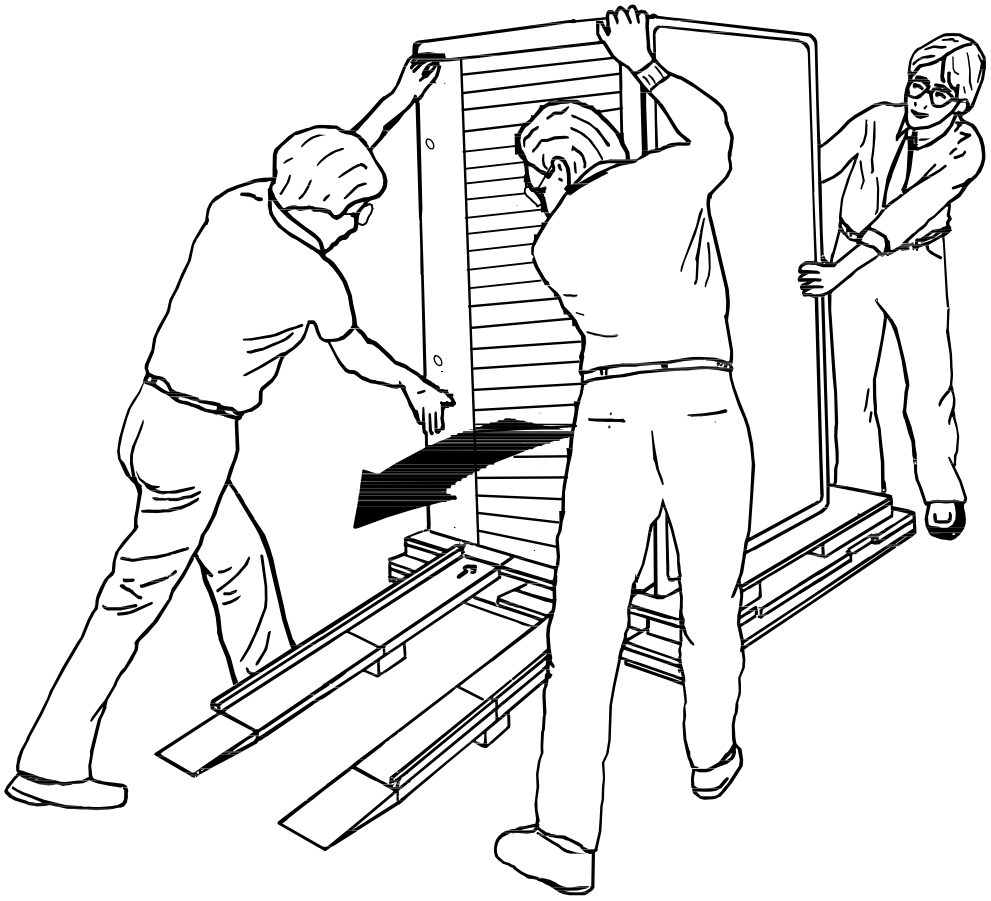

CXO-3808A

3.3 Placing the Cabinet

WARNING

Use extreme caution when rolling the cabinet across the floor. Failure to raise all leveler feet and to provide a clear path for the cabinet's casters can result in the cabinet tipping over and injury to personnel.

Once the cabinet has been removed from the pallet, the cabinet can be rolled to its final installation position. Secure loose cabinet cables up and out of the way when rolling the cabinet.

3.4 Leveling the Cabinet

Level the cabinet in its final position as follows:

1. Loosen the locknuts on all four leveler feet as shown in Figure 3–5.
2. Turn each leveler hex nut clockwise until the leveler foot contacts the floor.
3. Adjust all four leveler feet until the cabinet is level and the load is removed from all casters. Verify that the casters spin freely.
4. Tighten the locknuts on all four leveler feet.

Figure 3–5 Leveler Foot Adjustment

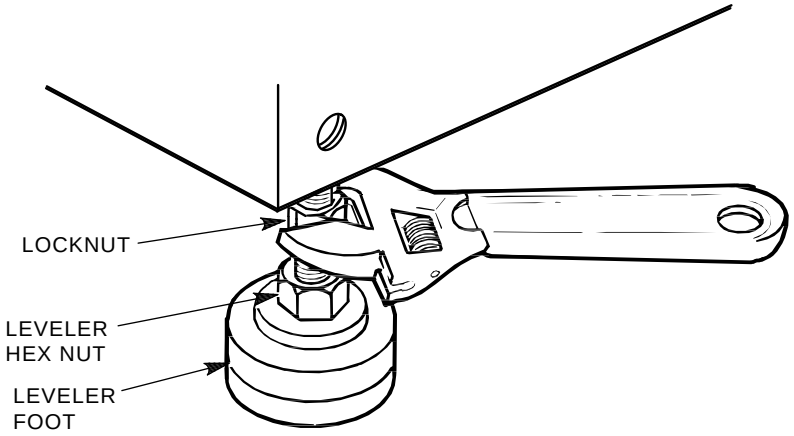

3.5 Installing the Skirt Kit

The skirt kit is packaged separately inside the corrugated carton with the cabinet. Installation of the skirt kit is optional. If you are installing the skirt kit, proceed as follows:

1. Unpack the skirt kit carton, using Figure 3–6 for reference, identify the right, left, front, and rear skirts.
2. Using Figure 3–6 for reference, position the skirts next to the cabinet.
3. The skirt fasteners are quarter-turn fasteners. Position each skirt against the cabinet such that the fasteners mate with the receptacles on the cabinet's base.
4. Using a screwdriver, push each fastener into its mating receptacle on the cabinet base, and turn it one quarter-turn clockwise.

   (Once the skirts are installed, there is a small amount of play that allows them to be adjusted slightly up or down for proper alignment.)

Figure 3–6 Cabinet Skirt Installation

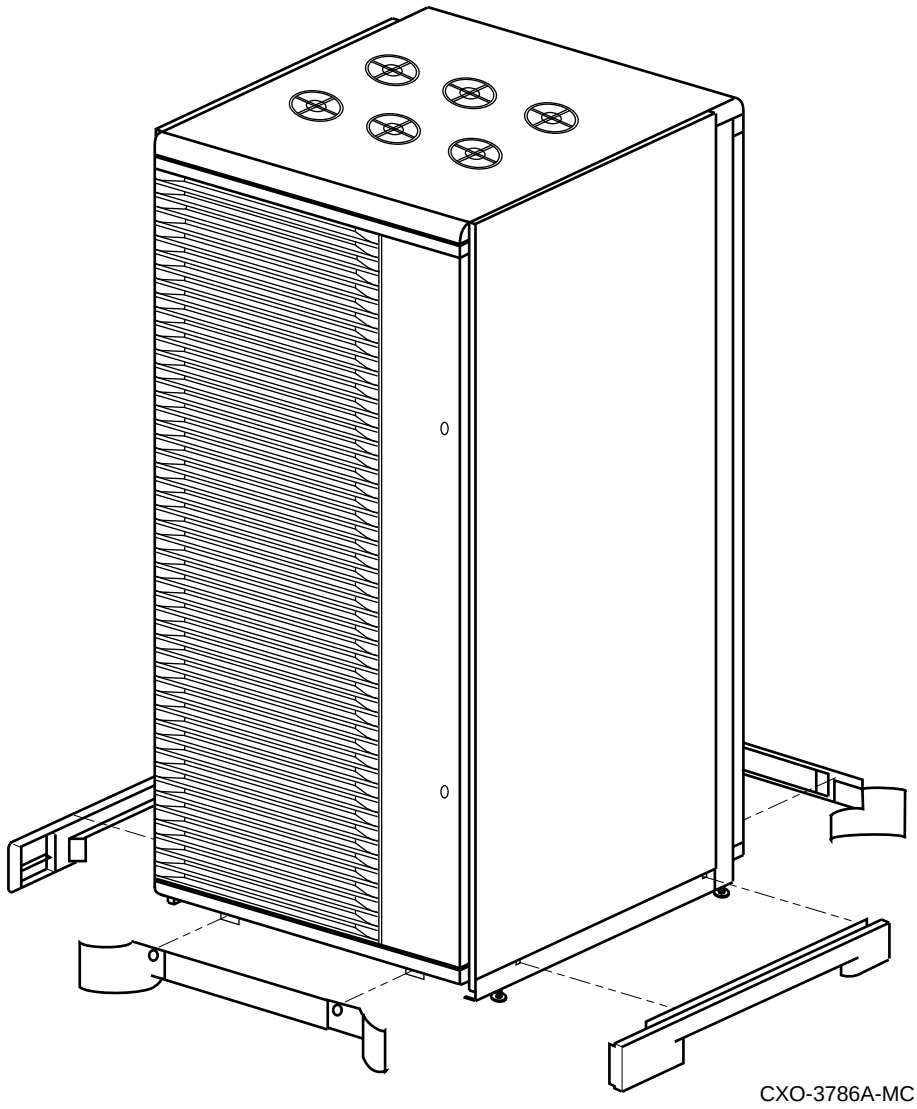

3.6 Inspecting the Cabinet

Inspect the cabinet installation as follows:

1. Make sure that all four leveler feet are lowered to support the full weight of the cabinet, and that the cabinet is level.
2. The remaining steps require access to the cabinet interior. To access the cabinet interior perform the following:
   a. Release the door locks mounted on the smooth vertical panel of each door by turning the locks counterclockwise with a 5/32-inch hex wrench.
   b. Swing the doors open
   c. Move the cabinet away from adjacent cabinets.

d. Remove the side panels.

Note

There are three side panel hanger clips on each side of the cabinet. A matching set of hanger clips are attached to each side panel.

e. Loosen the top cover by pushing up on its front and rear edges until it snaps free of its fasteners.

WARNING

The top cover is heavy and awkward to lift. Removing it requires two people. Failure to use sufficient personnel can result in personnel injury or equipment damage.

f. Using two people, lift the top cover from the cabinet and set it aside.

g. Remove the bolts attaching the side panels to the top side rails of the cabinet.

h. Grasp a panel along its front and rear edges and lift up until the hanger clips disengage. Lift the panel away from the cabinet.

i. Repeat the previous step to remove the other panel.

3. Make sure that the circuit breaker on the CDU(s) is in the ◯ (OFF) position.
4. Make sure that all hardware within the cabinet is fastened securely, and that there are no loose pieces present in the cabinet interior.
5. Check each of the cabinet fans to make sure that the blades turn freely and are not obstructed.
6. Make sure that there are no obstructions to the airflow from the shelf blowers.
7. Check the identification label on the rear of the cabinet to verify that the cabinet is configured to accept the power available at the site.
8. Make sure that all ac power cords connected from the shelves and cabinet fans to the CDUs are firmly seated in their connectors at both ends.
9. Make sure that all SBBs are seated firmly in their shelves.
10. Make sure that any necessary external interface cables are installed and firmly seated in their connectors.
11. Replace the side panels and top cover by reversing the procedures described in step 2.

3.7 Connecting a Terminal to the StorageWorks FDDI Server

Communicating with the StorageWorks FDDI server requires that you connect a terminal to it. Any terminal that supports ANSI control sequences can be used, including graphics displays that provide emulation of an ANSI compatible video terminal.

Use the following procedure to connect most EIA compatible terminals:

1. Make sure the power switch on the back of the terminal is OFF (0).
2. Connect one end of the terminal cable to the back of the terminal.
3. Connect the other end of the terminal cable to the EIA terminal port on the rear of the StorageWorks FDDI server cabinet using a H8571–J adapter as shown in Figure 3–7.

Figure 3–7 Connecting to the Terminal Port of the StorageWorks FDDI Server

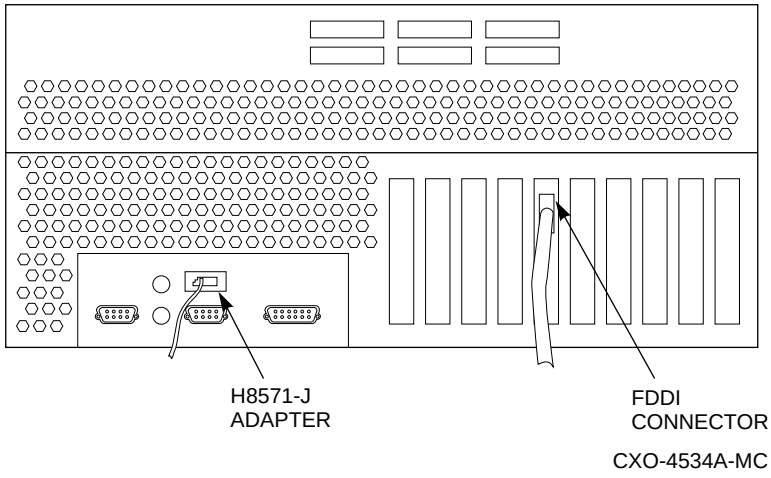

4. Turn the terminal power switch to the ON position.
5. Set the terminal at 9600 baud, with 8 data bits, 1 stop bit, and no parity. Refer to your terminal documentation for terminal setup instructions.

3.8 Applying Power to the Cabinet

Once the cabinet has been inspected and the terminal has been connected to the StorageWorks FDDI server, apply power as follows:

1. Plug the primary power cables from each CDU into the appropriate site power receptacles.
2. Switch the circuit breaker on each CDU to the | (ON) position.
3. Verify that all cabinet fans and shelf blowers are operating and that both status indicators on each shelf power supply SBB are illuminated. Refer to the *StorageWorks Solutions Shelf and SBB User's Guide* for further information on shelf status indicators.
4. Apply DC power to each server processor. The DC power switch is located on the OCP of the server processor. Figure 1–3 shows the location of the OCP.
5. Observe that the server processor power-on self-tests were successful by examining the terminal and looking for the self-test information as shown in Example 3–1.

Example 3–1 Server Processor Self Test Display

```
*** keyboard not plugged in...
ff.fe.fd.fc.fb.fa.f9.f8.f7.f6.f5.
ef.df.ee.f4.ed.ec.eb.....ea.e9.e8.e7.e6.e5.e4.e3.e2.e1.e0.
V1.1-1, built on Nov  4 1994 at 16:44:07
>>>
```

4

Customizing Server Operating System Parameters

This chapter describes the tasks that you perform to customize operating system parameters in the StorageWorks FDDI server.

4.1 Software Configuration Overview

Each StorageWorks FDDI server ships from the factory with the OpenVMS Alpha operating system preloaded on the system disks. You set StorageWorks FDDI server processor operating system parameters the **first time** you boot the servers.

Because each StorageWorks FDDI server processor is a VMScluster node, it is important that you perform the software configuration in the recommended sequence. Table 4–1 describes the configuration tasks in this and following chapters.

Table 4–1 Software Configuration Tasks

Step	Action	Refer to
1	Customize the parameters on the first server processor by running the Software Customization Procedure. Reboot the server processor upon completion.	Section 4.3, Executing the Software Customization Procedure.
2	If your StorageWorks FDDI server has two server processors, move the terminal to the second server processor and customize the parameters on the second server processor by running the Software Customization Procedure (see Section 3.7). Reboot the server processor upon completion.	Section 4.3, Executing the Software Customization Procedure.
3	Configure the StorageWorks storage devices according to your needs.	Chapter 5, Configuring the StorageWorks FDDI Server Storage Devices.

(continued on next page)

Table 4–1 (Cont.) Software Configuration Tasks

Step	Action	Refer to
4	Connect the server to the FDDI interconnect. For each server processor: 1. Shut down the server processor. 2. Reboot into the VMScluster system.	Section 6.1, Connecting to the FDDI Interconnect.
5	Ensure VMScluster system compatibility and server readiness	Section 6.2, Preparing for Operations.

4.2 VMScluster System Considerations

Introducing the StorageWorks FDDI server to a VMScluster system requires you to consider some broad, clusterwide issues. Among them are the VMScluster quorum scheme, enabling you to maintain an available VMScluster system, and Volume Shadowing interoperability. This section addresses each of these topics with some guidelines for you to consider.

4.2.1 Considering the VMScluster Quorum Scheme

The addition of a StorageWorks FDDI server to an existing VMScluster may impact existing VMScluster voting and quorum schemes. Two sample configurations may provide insight into possible VMScluster voting and quorum schemes:

Example 1:

The HS1*xx* is being added to an existing cluster (perhaps on a new FDDI ring) with an existing and working quorum scheme. In this case, the HS1*xx* is providing device serving only and does not contain a VMScluster system disk for other cluster nodes. In this case, the HS1*xx* should have ZERO votes and should contain no quorum disk (which may already exist in another part of the cluster). The expected votes on the HS1*xx* should be set to the current value in the cluster.

Example 2:

The HS1*xx* server provides the VMScluster system disk for other VMScluster satellites (nodes without votes). There are no other voting nodes in the cluster. This case would occur in a new cluster being formed with an HS1*xx* server as the core. In this case, the nodes in the HS1*xx* server must provide all the votes for the cluster. For the HS121 (two-node server), a quorum disk on the HS121 is recommended so that the cluster can continue with only one HS121 node operational. In the voting scheme, each HS121 node would have one vote; the quorum disk would have a single vote; and the expected votes for the cluster would be three. For the HS111 (single-node server), there is no need for a quorum disk, because only the single node in the HS111 will be contributing to quorum. In this case, the HS111 node would have one vote and the expected votes for the cluster would be one. Note that the cluster will hang if the HS111 is unavailable.

Many other cluster configurations are possible, and appropriate quorum and voting schemes should be designed by knowledgeable system and network managers, perhaps with assistance from consulting services. Extensive information regarding quorum and voting issues may be found in the *VMScluster Systems for OpenVMS* manual and the *Guidelines for VMScluster Configurations* manual.

4.2.2 Volume Shadowing Interoperability

Digital recommends that you install the OpenVMS Volume Shadowing Interoperability patch on the StorageWorks FDDI server if Phase II Volume Shadowing is being used in your cluster. Installing this patch on the server has clusterwide implications, because the server is one of the nodes in your VMScluster system. These implications need to be considered by knowledgeable system managers. For more information on this patch and its introduction to your VMScluster nodes, see Section 6.2.2 for references to the appropriate documentation.

The Software Customization Procedure contains more detailed information about the implications of adding this patch to your VMScluster nodes.[1] See Example 4–2 for more information about the Software Customization Procedure and the OpenVMS Volume Shadowing Interoperability kit.

Upon completing the installation of the StorageWorks FDDI server, you will be instructed to install the OpenVMS Volume Shadowing Interoperability kit on other VMScluster nodes if you have chosen to install this patch on your server. See Section 6.2.2 for this instruction and for references to the appropriate documentation.

4.3 Executing the Software Customization Procedure

Note

The Software Customization Procedure should be performed by an experienced cluster manager who is knowledgeable about the clusters at the installation site. Network knowledge also is required to complete the procedure.

Note

You must perform the Software Customization Procedure on each StorageWorks FDDI server processor.

The Software Customization Procedure is self-explanatory and provides immediate help for determining responses to questions asked during the procedure. To get help, type a question mark (?) as your response to a procedure question and press the Return key.

[1] This procedure gives you an opportunity to install the patch in the server during the customization procedure; the patch also can be installed after the server software has been customized if needed.

The Software Customization Procedure requires specific information in order to complete successfully. Before starting, make sure that you have the information described in Table 4–2. (Table 4–2 can be used as a worksheet. Use the underlined space in column 1 to record your response to each of the questions in advance.)

Once completed, the Software Customization Procedure will not run again. However, if you abort the procedure before it completes, it will appear the next time you boot the system.

After completing the Software Customization Procedure, manage the StorageWorks FDDI server using normal OpenVMS Alpha system management procedures. For information on managing the server, consult appropriate documentation on the OpenVMS Alpha operating system.

Table 4–2 Required Information for Customizing the OpenVMS Alpha Operating System

Use the information for the . . .	To respond to the prompts . . .
Current date and time _____	The current system date is dd-mmm-yyyy Is this date correct? [YES] : no Type the current date [ ]: and The current system time is hh:mm:ss Is this time correct? [YES] : no Type the current time [ ]:
System password _____	Type the password for the system account: and To verify the password is correct, type it again now. Password verification :
DECnet node name of the StorageWorks FDDI server _____	What is the server's DECnet node name?
DECnet address of the StorageWorks FDDI server _____	What is the server's DECnet node address?
Cluster group number _____	Enter this cluster's group number:
Cluster password _____	Enter this cluster's password: and Re-enter this cluster's password for verification:
Value of the disk ALLOCLASS parameter _____	Enter a value for *xxxxxx*'s † ALLOCLASS parameter?

†Where *xxxxxx* is the server's name.

(continued on next page)

Table 4–2 (Cont.) Required Information for Customizing the OpenVMS Alpha Operating System

Use the information for the . . .	To respond to the prompts . . .
Value of the tape ALLOCLASS parameter	Enter a value for *xxxxxx*'s † TAPE_ALLOCLASS parameter?
Name of the quorum disk ‡	What is the device name of the quorum disk?
The number of votes the server will contribute toward quorum	Enter a value for the number of votes:
The expected number of votes in the cluster	Enter a value for the number of expected votes:

†Where *xxxxxx* is the server's name.

‡The quorum disk, if specified, must be attached to the server's device channel processor; that is, it must be a DUA disk and not the server's local system disk.

To begin the Software Customization Procedure, type B at the **>>>** system prompt on the terminal attached to the server processor. Examples 4–1 and 4–2 show sample Software Customization Procedure log files and provide sample answers to the questions in the Software Customization Procedure. All user input is shown in bold text except for passwords, which are not echoed, and carriage returns.

Example 4–1 Software Customization Procedure Log File—Part 1

```
>>>boot
(boot dka0.0.0.6.0 -flags 0)
block 0 of dka0.0.0.6.0 is a valid boot block
reading 1007 blocks from dka0.0.0.6.0
bootstrap code read in
base = 12e000, image_start = 0, image_bytes = 7de00
initializing HWRPB at 2000
initializing page table at 120000
initializing machine state
setting affinity to the primary CPU
jumping to bootstrap code

    OpenVMS AXP (TM) Operating System, Version V6.1-1H2

   +-------------------------------------------------------------------+
   |                            Welcome to the                         |
   |                                                                   |
   |                       StorageWorks FDDI Server                    |
   |                                                                   |
   |                   Software Customization Procedure                !
   +-------------------------------------------------------------------+

   Please wait while the Software Customization procedure initializes the
   system.  This will take less than one minute.
```

(continued on next page)

Example 4–1 (Cont.) Software Customization Procedure Log File—Part 1

```
   +---------------------------------------------------------------------+
   |                   Software Customization Procedure                  |
   |                                                                     |
   |                            Introduction                             |
   +---------------------------------------------------------------------+

   The Software Customization procedure will ask you to enter the
   information required for the factory installed software to be
   custom configured for your site.

   To answer a question, type the requested information and press Return.
   To get help at any time, type a question mark (?) and press Return.

   Press Return to continue:

Note:  The following licenses have been registered on this system

Active licenses on node :

------- Product ID --------    ---- Rating ----- -- Version --
Product           Producer    Units Avail Activ Version Release    Termination
OPENVMS-ALPHA     DEC             0  0     100    0.0  (none)      28-FEB-1995
VMSCLUSTER        DEC             0  0     100    0.0  (none)      28-FEB-1995
VOLSHAD           DEC             0  0     100    0.0  (none)      28-FEB-1995

   To register additional license PAKs, type the following command after
   the system restarts and you have logged in to the system manager's
   account:

   $ @SYS$UPDATE:VMSLICENSE

   Press Return to continue:

   The following information is essential to ensure that the system
   can be customized correctly:

   o  The current date and time
   o  The password for the system manager's account on the server
   o  Cluster Information
      -   The DECnet node name of the server
      -   The DECnet address of the server
      -   The cluster group number and password
      -   The value of the ALLOCLASS parameter for disk and tape
      -   Whether this server will serve tapes
      -   The name of a quorum disk if a quorum disk will be installed
          on the server.  (This disk must be attached to the server's
          device channel processor; that is, it must be a DUA disk and
          not the server's local system disk.)
      -   The number of Votes the server will contribute toward quorum
      -   The expected number of Votes in the cluster

   You cannot complete the Software Customization procedure if you do not
   have all the essential information.  If you do not have this information,
   you can obtain it from your system manager or network administrator.


   Are you ready to proceed with the Software Customization? [YES]: YES

   At the end of the Software Customization procedure, the system will
   restart itself.  Then it will be ready for use.

   +---------------------------------------------------------------------+
   |                       Set the system date and time                  |
   +---------------------------------------------------------------------+

   To set the server up correctly, you must set the system date and
   time.
```

(continued on next page)

Example 4–1 (Cont.) Software Customization Procedure Log File—Part 1

```
   To get help, type a question mark (?) and press Return.

   The current system date is 11-JAN-1994.

   Is this date correct? [YES]: NO

   Type the system date, using the following format: dd-mmm-yyyy
   For example, if the date is 19 July 1993, type the following:

   19-JUL-1993

   Type the current date []: 11-JAN-1995

   The current system time is 22:10:27

   Is this time correct? [YES]: NO

   Type the system time, using the following format: hh:mm
   For example, if the time is 3.05pm, type the following:

   15:05

   Type the current time []: 22:15

   The current system date and time is 11-JAN-1995 22:15:00.01

   Are the date and time correct? [YES]: YES

   +-------------------------------------------------------------------------+
   |                          Set the system password                        |
   +-------------------------------------------------------------------------+

   The system password gives you access to the system manager's account on
   the server.  When the Software Customization procedure completes, you will
   be able to log in to the system manager's account on the server by typing
   the user name SYSTEM, followed by the system password.

   The system password that you specify must contain at least 8 characters.
   For security reasons, the characters you type are not displayed on the
   screen.

   IMPORTANT:  Do not forget the system password, otherwise you will be
               unable to log in to the system manager's account.

   Type the password for the SYSTEM account:

   To verify that the password is correct, type it again now.

   Password verification:

   The password for the SYSTEM account has been verified.

   +-------------------------------------------------------------------------+
   |                          Configure Cluster                              |
   +-------------------------------------------------------------------------+

What is the server's DECnet node name? FSERV1
What is the server's DECnet node address? 63.849
Enter this cluster's group number: 3396
Enter this cluster's password:
Re-enter this cluster's password for verification:
```

(continued on next page)

Example 4–1 (Cont.) Software Customization Procedure Log File—Part 1

```
Do you wish FSERV1 to automatically serve disks? (Y/N): Y
Enter a value for FSERV1's ALLOCLASS parameter: 31
Do you wish FSERV1 to automatically serve tapes? (Y/N): Y
Enter a value for FSERV1's TAPE_ALLOCLASS parameter: 31
Will a quorum disk be installed on the server? (Y/N): Y
Enter the device name of the QUORUM disk: $31$DUA100:
Enter the number of votes that the server will contribute toward quorum: 1
Enter a value for the number of expected votes in the cluster: 3

      You have chosen the following cluster parameters:

Server name (SCSNODE)                  FSERV1
System id (SCSSYSTEMID)                65361 (63.849)
Cluster group number                   3396
Automatically serve disks              YES
Disk allocation class (ALLOCLASS)      31
Quorum disk                            $31$DUA100
Tape server (TMSCP_LOAD)               YES
Tape allocation class (TAPE_ALLOCLASS) 31
Number of Votes for this server        1
Number of Expected Votes in the cluster 3

Would you like to change any of these parameters? (Y/N) : NO
```

If you respond with an N to the last question shown in the previous example, the Software Customization Procedure continues as shown in Example 4–2. If you respond with a Y, then the previous script repeats.

Example 4–2 Software Customization Procedure Log File—Part 2

```
   ********************************************************************
   *                                                                  *
   *    The following parameters have been set for FSERV1:            *
   *                  VOTES = 1                                       *
   *                  QDSKVOTES = 1                                   *
   *                                                                  *
   *    After FSERV1 has booted into the cluster, you must            *
   *    adjust the value for EXPECTED_VOTES in every cluster          *
   *    member's  MODPARAMS.DAT.  You must then run AUTOGEN on        *
   *    each node and shutdown and reboot the cluster.  This          *
   *    will cause EXPECTED_VOTES to be correct on each               *
   *    node.                                                         *
   *                                                                  *
   *    For example, if the server includes two Alpha VMS             *
   *    systems and a quorum disk, then EXPECTED_VOTES for            *
   *    each cluster member will have to be increased by              *
   *    a total of 3.                                                 *
   *                                                                  *
   ********************************************************************

  Cluster set-up complete.  Continuing...

 +------------------------------------------------------------------------+
 |           Install Volume Shadowing for OpenVMS Alpha ECO kit           |
 +------------------------------------------------------------------------+

    You now have the option of installing the Volume Shadowing ECO kit.
```

(continued on next page)

Example 4–2 (Cont.) Software Customization Procedure Log File—Part 2

```
    The StorageWorks FDDI Server system disk and CD-ROM
    contain an optional patch kit that provides the following
    corrections to the OpenVMS Alpha Version 6.1 Volume Shadowing
    software:

    o  Enabling of minimerge features in a mixed-architecture
       VMScluster that includes OpenVMS Alpha Version 6.1 nodes

    o  Support for geometry-based shadowing device recognition
       that allows devices, such as the RZ28 and RZ28B, to work
       in the same shadow set

    o  Software enhancements that provide increased
       interoperability for a mixed-architecture VMScluster

Press Return to continue:

    Digital recommends that you install this kit if you are
    using the OpenVMS Volume Shadowing product on your system.
    This can be done at this time or after the system is
    rebooted by using VMSINSTAL to accomplish this task by
    entering the following command:

    $ @SYS$UPDATE:VMSINSTAL AXPSHAD04_061 sys$sysdevice:[000000]

    These corrections are also available in a TIMA remedial
    kit that is available through Digital Customer Services
    worldwide. For more information, please call your Customer
    Services representative.

Press Return to continue:
     ******************************************************
     WARNING WARNING WARNING WARNING WARNING WARNING WARNING
     ******************************************************

  If you have a mixed-architecture cluster, you must be absolutely
  sure to install this kit --- as well as the VAX version of this kit
  on VAX nodes --- on each node in the cluster BEFORE you bring up both
  types of systems in a cluster again.  If all nodes have not had the
  appropriate kit installed, you may not be able to create shadow sets.

  If this kit has been installed on each node of your existing cluster,
  then allow this procedure to install the kit on the server.

  If the kit has NOT been installed on each node of your cluster, then
  do not install the kit at this time.  You may install the kit later.
  However, please note the following restrictions:

   o  VMSclusters with shadowed  SCSI  disks  and  mixed-architecture
      VMSclusters  running OpenVMS Alpha Version 6.1 must apply the kit
      to each node in the cluster and reboot the entire cluster
      simultaneously.  In these cases, rolling upgrades are not supported.
   o  Working configurations that contain SCSI shadow sets on
      dissimilar controllers may no longer work.

Press Return to continue:

  For more information, please see the Problem Description section
  of the release notes supplied with the patch kits. The release notes
  are located on the server's system disk as follows:

  SYS$HELP:AXPSHAD04_061.RELEASE_NOTES   ! Alpha version
  SYS$HELP:VAXSHAD04_061.RELEASE_NOTES   ! VAX version

  The patch kits are located in the root directory of the server's
  system disk:
```

(continued on next page)

Example 4–2 (Cont.) Software Customization Procedure Log File—Part 2

```
     sys$sysdevice:[000000]AXPSHAD04_061.*     ! Alpha version
     sys$sysdevice:[000000]VAXSHAD04_061.*     ! VAX version

     Please re-read the above notes regarding the cluster impact
     of patch installation.

   Are you sure you are ready to install this patch kit? YES

    Installation of the patch kit will now take place.

    Since the server's CD-ROM provides a backup for the server's system disk,
    please answer YES to the next question regarding system disk backup.

   Press Return to continue:

        OpenVMS AXP Software Product Installation Procedure V6.1-1H2

It is 11-JAN-1995 at 22:16.

Enter a question mark (?) at any time for help.

* Are you satisfied with the backup of your system disk [YES]? YES

The following products will be processed:

  AXPSHAD04_ V6.1

        Beginning installation of AXPSHAD04_ T6.1 at 22:16

%VMSINSTAL-I-USEANS, The auto-answer file will be used.
%VMSINSTAL-I-RESTORE, Restoring product save set A ...
%VMSINSTAL-I-RELMOVED, Product's release notes have been moved to SYS$HELP.

               WARNING
               -------

        If you apply this remedial kit, please note the following
        restrictions:

   o    VMSclusters with shadowed SCSI disks and mixed-architecture
        VMSclusters running OpenVMS AXP Version 6.1 must apply the kit
        and reboot the entire cluster simultaneously. In these cases,
        rolling upgrades are not supported.

   o    Working configurations that contain SCSI shadow sets on dissimilar
        controllers may no longer work.

        For more information, please see the Problem Description section of
        the cover letter/release notes supplied with this kit.

* Do you want to continue [NO]? YES

        To complete the installation of this product, you must reboot
        the system.  If it is not convenient to reboot at this time, then
        enter NO to the following question. The installation of this kit
        will continue and the files moved to their appropriate locations
        without forcing the system to reboot upon completion of the
        installation. The system can than be rebooted at some more
        convenient time to actually have this update take effect.

        Entering YES will cause the system to automatically reboot
        upon the installation of this kit.

* Will you allow a system shutdown after this product is installed? [YES]: NO

    No more questions will be asked ...

    Now applying AXPSHAD04_061 ...
```

(continued on next page)

Example 4–2 (Cont.) Software Customization Procedure Log File—Part 2

```
%VMSINSTAL-I-RESTORE, Restoring product save set C ...
 0) MOUNTSHR (new image)
 1) MSCP (new image)
 2) SHADOW_SERVER (new image)
 3) SYS$DUDRIVER (new image)
 4) SYS$SHDRIVER (new image)
 5) SYS$TUDRIVER (new image)
 6) SYS$VCC (new image)
 7) SYS$VCC_MON (new image)
 8) TMSCP (new image)

%VMSINSTAL-I-MOVEFILES, Files will now be moved to their target directories...
        Installation of AXPSHAD04_ T6.1 completed at 22:16
    Adding history entry in VMI$ROOT:[SYSUPD]VMSINSTAL.HISTORY
    Creating installation data file: VMI$ROOT:[SYSUPD]AXPSHAD04_061.VMI_DATA
        VMSINSTAL procedure done at 22:17

   +------------------------------------------------------------------------+
   |                           End of Data Entry                            |
   +------------------------------------------------------------------------+

   This is the end of the data entry part of the Software Customization
   procedure -- you do not need to provide any more information. The server
   will now shut down and restart automatically. When it restarts, you can
   log in to the system manager's account using the user name SYSTEM and
   the password that you specified earlier.

   Press Return to complete the Customization procedure:

   The server is now being configured. This will take less than one
   minute.

   Please wait...

   Server is being configured...     22:17:20

   The server SYSGEN parameters have been set. Please wait for the
   server to restart.

   IMPORTANT:  If the server has not restarted within 3 minutes,
               power cycle the server processor by pressing the
               POWER button on the server control panel twice
               to turn power off and then back on.  Then type BOOT
               when the console prompt (>>>) is displayed.

   The server will now restart automatically. Please wait...

halted CPU 0

halt code = 5
HALT instruction executed
PC = ffffffff8004f91c

CPU 0 booting
```

(continued on next page)

Example 4–2 (Cont.) Software Customization Procedure Log File—Part 2

```
(boot dka0.0.0.6.0 -flags 0)
block 0 of dka0.0.0.6.0 is a valid boot block
reading 1007 blocks from dka0.0.0.6.0
bootstrap code read in
base = 12e000, image_start = 0, image_bytes = 7de00
initializing HWRPB at 2000
initializing page table at 120000
initializing machine state
setting affinity to the primary CPU
jumping to bootstrap code

    OpenVMS AXP (TM) Operating System, Version V6.1-1H2
```

After the server has shut down, it is necessary to cycle power for the server processor by pressing its power switch off and then on again. After the power-on self-tests complete and the system prompt >>> appears, type B to reboot the system. (See Example 3–1 for an example display of the power-on self-tests.)

Pressing the Return key at this point results in the system prompting for username and password.

Note

If you specified a quorum disk as part of the StorageWorks FDDI server software installation, the quorum disk must be mounted after the OpenVMS Alpha operating system is running. As soon as possible, add a command to mount the quorum disk on the server to the startup procedure for the server's processor. This allows you to maintain quorum in the event one of the server processors fail (see Section 7.1.1).

5

Configuring the StorageWorks FDDI Server Storage Devices

The StorageWorks FDDI server requires you to configure the StorageWorks storage system to meet your particular storage needs. Appendix A provides a detailed description of stripesets and RAIDsets that you can configure for the StorageWorks FDDI server.

This chapter begins by providing a few general guidelines to help you decide how you may want to configure the StorageWorks storage system, continues with the steps required to use the CFMENU utility from a terminal connected to the server processor, and concludes by providing a sample configuration.

5.1 General Considerations for Configuring a StorageWorks Storage System

Configuring the storage system is the process of defining, through parameters, the logical organization of the storage devices. The storage devices may be configured as one or more of the following:

- Stripeset—A virtual disk drive with its physical data spread across multiple physical disks.
- RAIDset—Three or more physical disks that are connected to present an array of disks as a single virtual unit to the host.
- Spareset—A pool of disk drives used by the device channel processor to replace failing members of a RAIDset.
- Failedset—A group of disk drives that have been removed from RAIDsets due to a failure or a manual removal.
- Container—An entity that is capable of storing data, whether it is one physical device or a group of physical devices. A disk, a stripeset, and a RAIDset are examples of a container.
- Passthrough Container—A virtual device used to pass SCSI commands that perform operations other than simple device reads or writes. Typically used with tape libraries that contain standard tape devices plus a robotic mechanism for selecting and loading cartridges from a pool of tapes. A passthrough container must exist to pass SCSI commands to the robotic mechanism.
- Unit—A logical entity composed of one or more devices and treated as a single addressable storage structure by the cluster members. Units can be single devices, stripesets, or RAIDsets.

5.2 Running CFMENU from the StorageWorks FDDI Server Terminal

Before CFMENU can be invoked from a terminal connected to the StorageWorks FDDI server, you must establish a logical connection to the device channel processor. To create a logical connection to the device channel processor, perform the following steps:

1. Log into the system manager's or other appropriate account on the StorageWorks FDDI server (the account must have DIAGNOSE privilege).
2. At the DCL prompt, enter the SHOW CLUSTER command. The system will respond with default cluster information similar to that shown in Figure 5–1 and return control to the DCL command level.

Figure 5–1 Example Cluster Information

View of Cluster from system ID 65315 Node: FUTBAL 1−JAN−1995 12:54:34

SYSTEMS		MEMBERS
NODE	**SOFTWARE**	**STATUS**
GNDBAL	**VMS V6.1**	**MEMBER**
FUTBAL	**VMS V6.1**	**MEMBER**
SOFBAL	**VMS V6.1**	**MEMBER**
RLRBAL	**VMS V6.1**	**MEMBER**
HS1CP1	**HSD E35D**	
HS1CP2	**HSD E35D**	

3. Identify the names of the device channel processors (note, there may be more than one). The device channel processor can be identified by searching for the nodes that are named HS1CP1 or HS1CP2.
4. Change the terminal format to accommodate 132 characters with this command:

   ```
   $ SET TERMINAL/WIDTH=132
   ```

5. Logically connect the StorageWorks FDDI server terminal to the device channel processor by entering the following command at the DCL prompt:

   ```
   $ SET HOST/DUP/SERVER=mscp$dup/TASK=CLI node-name
   ```

   Where:

   node-name is the name assigned to the device channel processor (for example, HS1CP1).

   The device channel processor responds with its prompt (for example, HS1CP>).

6. Enter the following command to start the CFMENU utility:

   ```
   HS1CP> RUN CFMENU
   ```

The CFMENU utility begins executing. After CFMENU's initial message, press the Return key to get to CFMENU's main menu as shown in Figure 5–2. Any recognized devices that have not been added to your configuration will appear, with their Port Target LUN (PTL), in the column to the right of the menu options. All the devices shown in Figure 5–2 are attached but not yet configured.

Figure 5–2 CFMENU Main Menu

```
MAIN MENU:                    |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In- Re-       W W
1. Add/delete devices         | Dev.PTLs |       PTLs     Name          ID            Name    Type    Size  sp. it'd duc  Unit P B
2. Add/delete stripesets      |----------|      ------ --------- ----------------  --------- ------ ------- --- ---- --- ----- - -
3. Add/delete raidsets/       | 100 (dsk)|
   sparesets/failedsets       | 110 (dsk)|
4. Add/delete passthrough     | 120 (dsk)|
5. Initialize devices         | 130 (dsk)|
   and/or storagesets         | 140 (dsk)|
6. Add/delete units           | 150 (dsk)|
7. Setup terminal             | 200 (dsk)|
8. Exit CFMENU                | 210 (dsk)|
                              | 220 (dsk)|
                              | 230 (dsk)|
                              | 240 (dsk)|
                              | 250 (dsk)|
  D=Scroll down  U=Scroll up  | 300 (dsk)|
Enter menu choice (1,5) [5] ?5
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
```

Note, if your terminal is capable of displaying more than 24 rows, you may want to enter option 7 from the main menu to set the number of rows CFMENU will display.

Table 5–1 describes the information headings of the main menu. To avoid confusion, the information headings are presented exactly as they appear on the main menu.

Table 5–1 CFMENU Information Headings

Information Heading	Description
Main Menu	Lists the major operations that are available
Unconfig'd Dev.PTLs	Unconfigured PTLS—The device PTLs that are recognized by the HS1CP but that have not yet been added to the configuration.
Config'd PTLs	Configured PTLS—The device PTLs that have been configured by the HS1CP.
Device Name	Device name—The name automatically assigned to the device when it is configured by the HS1CP. The name can be changed through use of the CLI commands. (See Appendix B for complete descriptions of the CLI commands.)
Product ID	Product ID—Information that identifies the device.
Stor.set Name	Storage set name—The name assigned by CFMENU to the storage set. The name is assigned using one of the following conventions: • S*x* for stripesets • R*x* for RAIDsets • P*x* for passthrough containers The number denoted by *x* is assigned sequentially beginning with 1 for each type of storage set. The name can be changed through use of the CLI commands. (See Appendix B for complete descriptions of the CLI commands.)
Stor.set Type	Storage Set—The type of storage set: STRP for stripesets, RAID for RAIDsets, PASS for passthrough containers.
Chnk Size	Chunksize—The size, in blocks, of data transfers with the device.
Trn sp.	Transportable—The letter, Y or N, to indicate if the device is transportable or not. A transportable device is one that does not have metadata written on it. A nontransportable device has a small amount of metadata written on it.
In- it'd	Initialized—The letter, Y or N, to indicate if the device has been initialized.
Re- duc	Reduced—Valid for RAIDsets only, indicates the RAIDset is missing one member.
Unit	Unit—The logical number of the unit as assigned by you and preceded by the letter D for disk or T for tape.
W P	Write Protected—The letter, Y or N, to indicate whether the device is write protected.
W B	Write-back—The letter, Y or N, to indicate if the device has write-back caching enabled.

5.3 Completing a Sample Configuration Using CFMENU

CFMENU allows you to quickly configure storage devices attached to the HS1CP. CFMENU presents, in a menu format, configuration commands normally entered at the Command Line Interpreter (CLI). Because CFMENU prompts you to choose options for devices, storagesets, and units based on the command qualifiers of the CLI, you should have a good understanding of the various CLI commands. A complete discussion of CLI commands can be found in Appendix B.

Note

The following configuration is meant as a sample only. It does not configure all the devices.

The following sample configuration assumes a HS121 model StorageWorks FDDI server with a complement of 36 storage devices. It is further assumed that you have decided to configure the following:

- Two stripesets, one consisting of six drives, the other consisting of five drives
- One RAIDset consisting of five drives
- One spareset with two drives

In order to create the configuration you will perform the following tasks:

- Add the attached storage devices to the configuration
- Add the two stripesets
- Add the RAIDset
- Add devices to the spareset
- Initialize the containers
- Add the units

Note

You must perform the same configuration on each HS1CP device channel processor.

5.3.1 Adding Devices

From the main menu, enter option 1 then press the Return key. The device menu appears as shown in Figure 5–3. The list of unconfigured devices appears to the right of the options. If the list is too long to be shown on one screen, enter D or U to scroll the information down or up.

Figure 5–3 CFMENU Device Menu (Before Adding Devices)

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
DEVICE MENU:                |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In- Re-       W W
1. Add a device from list   | Dev.PTLs |       PTLs     Name         ID            Name     Type     Size  sp. it'd duc  Unit P B
   of PTLs not configured   |----------|      ------ --------- ----------------  --------- ------ ------- --- ---- --- ----- - -
   (marked with ^)          |^100 (dsk)|
2. Delete an unbounded device |^110 (dsk)|
   (marked with *)          |^120 (dsk)|
3. Add all devices from list |^130 (dsk)|
   of PTLs not configured   |^140 (dsk)|
   (marked with ^)          |^150 (dsk)|
4. Delete all unbounded     |^200 (dsk)|
   devices (marked with *)  |^210 (dsk)|
5. Return to main menu      |^220 (dsk)|
                            |^230 (dsk)|
                            |^240 (dsk)|
                            |^250 (dsk)|
  D=Scroll down  U=Scroll up|^300 (dsk)|
```

In Figure 5–3, none of the devices have been added to the configuration yet. If you enter option 1, CFMENU asks you whether or not (y/n/q) to add each device on the list. If you enter option 3, CFMENU will add *all* the unknown devices. The y/n/q asks whether you want the option (yes), do not want the option (no), or want to stop what you are doing (quit).

After entering option 1 and adding all of the devices, the screen resembles Figure 5–4.

Figure 5–4 CFMENU Device Menu (After Adding Devices)

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
DEVICE MENU:                |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In- Re-       W W
1. Add a device from list   | Dev.PTLs |       PTLs     Name         ID            Name     Type     Size  sp. it'd duc  Unit P B
   of PTLs not configured   |----------|      ------ --------- ----------------  --------- ------ ------- --- ---- --- ----- - -
   (marked with ^)          |          |disks:  100 *  DISK100 RZ28     (C) DEC                          N    Y
2. Delete an unbounded device |          |        110 *  DISK110 RZ28     (C) DEC                          N    Y
   (marked with *)          |          |        120 *  DISK120 RZ28     (C) DEC                          N    Y
3. Add all devices from list |          |        130 *  DISK130 RZ28     (C) DEC                          N    Y
   of PTLs not configured   |          |        140 *  DISK140 RZ28     (C) DEC                          N    Y
   (marked with ^)          |          |        150 *  DISK150 RZ28     (C) DEC                          N    Y
4. Delete all unbounded     |          |        200 *  DISK200 RZ28     (C) DEC                          N    Y
   devices (marked with *)  |          |        210 *  DISK210 RZ28     (C) DEC                          N    Y
5. Return to main menu      |          |        220 *  DISK220 RZ28     (C) DEC                          N    Y
                            |          |        230 *  DISK230 RZ28     (C) DEC                          N    Y
                            |          |        240 *  DISK240 RZ28     (C) DEC                          N    Y
                            |          |        250 *  DISK250 RZ28     (C) DEC                          N    Y
  D=Scroll down  U=Scroll up|          |        300 *  DISK300 RZ28     (C) DEC                          N    Y
```

After adding devices, return to the main menu.

5.3.2 Adding Stripesets

Enter option 2 from the main menu to configure stripesets. From the stripeset menu (see Figure 5–5), enter option 1 to create a stripeset. CFMENU prompts you for how many and which devices from the configured PTLs list you wish to include in the stripeset (2–14 devices allowable). Figure 5–6 shows the result of creating two stripesets from the disks at the PTL's shown.

Figure 5–5 CFMENU Stripeset Menu

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
STRIPESET MENU:              |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Create a stripeset        | Dev.PTLs |       PTLs     Name          ID           Name     Type     Size  sp.  it'd duc  Unit P B
   (eligible devices marked  |----------|      ------ --------- ----------------  --------- ------ ------- ---  ---- --- ----- - -
   by ^)                     |          |disks:  100 ^  DISK100 RZ28     (C) DEC                          N    Y
2. Delete an unbounded       |          |        110 ^  DISK110 RZ28     (C) DEC                          N    Y
   stripeset (marked by *)   |          |        120 ^  DISK120 RZ28     (C) DEC                          N    Y
3. Delete all unbounded      |          |        130 ^  DISK130 RZ28     (C) DEC                          N    Y
   stripesets (marked by *)  |          |        140 ^  DISK140 RZ28     (C) DEC                          N    Y
4. Return to main menu       |          |        150 ^  DISK150 RZ28     (C) DEC                          N    Y
                             |          |        200 ^  DISK200 RZ28     (C) DEC                          N    Y
                             |          |        210 ^  DISK210 RZ28     (C) DEC                          N    Y
                             |          |        220 ^  DISK220 RZ28     (C) DEC                          N    Y
                             |          |        230 ^  DISK230 RZ28     (C) DEC                          N    Y
                             |          |        240 ^  DISK240 RZ28     (C) DEC                          N    Y
                             |          |        250 ^  DISK250 RZ28     (C) DEC                          N    Y
  D=Scroll down  U=Scroll up |          |        300 ^  DISK300 RZ28     (C) DEC                          N    Y

Enter menu choice (1,4) [4] ?1
```

Figure 5–6 CFMENU Showing Created Stripesets

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
STRIPESET MENU:              |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Create a stripeset        | Dev.PTLs |       PTLs     Name          ID           Name     Type     Size  sp.  it'd duc  Unit P B
   (eligible devices marked  |----------|      ------ --------- ----------------  --------- ------ ------- ---  ---- --- ----- - -
   by ^)                     |          |        640 ^  DISK640 RZ28     (C) DEC                          N    Y
2. Delete an unbounded       |          |        650 ^  DISK650 RZ28     (C) DEC                          N    Y
   stripeset (marked by *)   |          |strps:  100    DISK100 RZ28     (C) DEC *    S1     STRP    unk          N
3. Delete all unbounded      |          |        110    DISK110 RZ28     (C) DEC      "        "      "           "
   stripesets (marked by *)  |          |        120    DISK120 RZ28     (C) DEC      "        "      "           "
4. Return to main menu       |          |        130    DISK130 RZ28     (C) DEC      "        "      "           "
                             |          |        140    DISK140 RZ28     (C) DEC      "        "      "           "
                             |          |        150    DISK150 RZ28     (C) DEC      "        "      "           "
                             |          |        200    DISK200 RZ28     (C) DEC *    S2     STRP    unk          N
                             |          |        300    DISK300 RZ28     (C) DEC      "        "      "           "
                             |          |        400    DISK400 RZ28     (C) DEC      "        "      "           "
                             |          |        500    DISK500 RZ28     (C) DEC      "        "      "           "
  D=Scroll down  U=Scroll up |          |        600    DISK600 RZ28     (C) DEC      "        "      "           "
Enter menu choice (1,4) [4] ?
```

After adding stripesets, return to the main menu.

5.3.3 Adding RAIDsets

Enter option 3 from the main menu to configure RAIDsets. From the RAIDset menu (see Figure 5–7), enter option 1 to add a RAIDset. CFMENU prompts you for how many and which devices from the configured PTLs list you wish to include in the RAIDset (3–14 devices allowable). Figure 5–8 shows the result of creating a RAIDset from disks at the PTLs shown.

Note

RAIDsets require that write-back cache be licensed and enabled.

Figure 5–7 CFMENU RAIDset Menu

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
RAIDSET MENU:                |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Create a raidset (eligible | Dev.PTLs |       PTLs     Name         ID            Name     Type     Size  sp.  it'd duc  Unit P B
   devices marked by ^)      |----------|      ------ --------- ----------------  --------- ------ ------- ---  ---- --- ----- - -
2. Delete an unbounded       |          |disks:  210 ^  DISK210 RZ28     (C) DEC                          N    Y
   raidset (marked by *)     |          |        220 ^  DISK220 RZ28     (C) DEC                          N    Y
3. Delete all unbounded      |          |        230 ^  DISK230 RZ28     (C) DEC                          N    Y
   raidsets (marked by *)    |          |        240 ^  DISK240 RZ28     (C) DEC                          N    Y
4. Add/delete device in      |          |        250 ^  DISK250 RZ28     (C) DEC                          N    Y
   SPARESET or FAILEDSET     |          |        310 ^  DISK310 RZ28     (C) DEC                          N    Y
5. Replace member of a       |          |        320 ^  DISK320 RZ28     (C) DEC                          N    Y
   reduced raidset           |          |        330 ^  DISK330 RZ28     (C) DEC                          N    Y
6. Return to main menu       |          |        340 ^  DISK340 RZ28     (C) DEC                          N    Y
                             |          |        350 ^  DISK350 RZ28     (C) DEC                          N    Y
                             |          |        410 ^  DISK410 RZ28     (C) DEC                          N    Y
                             |          |        420 ^  DISK420 RZ28     (C) DEC                          N    Y
  D=Scroll down  U=Scroll up |          |        430 ^  DISK430 RZ28     (C) DEC                          N    Y
```

Figure 5–8 CFMENU Showing Created RAIDsets

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
RAIDSET MENU:                |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Create a raidset (eligible | Dev.PTLs |       PTLs     Name         ID            Name     Type     Size  sp.  it'd duc  Unit P B
   devices marked by ^)      |----------|      ------ --------- ----------------  --------- ------ ------- ---  ---- --- ----- - -
2. Delete an unbounded       |          |        130    DISK130 RZ28     (C) DEC      "        "      "          "
   raidset (marked by *)     |          |        140    DISK140 RZ28     (C) DEC      "        "      "          "
3. Delete all unbounded      |          |        150    DISK150 RZ28     (C) DEC      "        "      "          "
   raidsets (marked by *)    |          |        200    DISK200 RZ28     (C) DEC      S2       STRP   unk        N
4. Add/delete device in      |          |        300    DISK300 RZ28     (C) DEC      "        "      "          "
   SPARESET or FAILEDSET     |          |        400    DISK400 RZ28     (C) DEC      "        "      "          "
5. Replace member of a       |          |        500    DISK500 RZ28     (C) DEC      "        "      "          "
   reduced raidset           |          |        600    DISK600 RZ28     (C) DEC      "        "      "          "
6. Return to main menu       |          |raid5:  210    DISK210 RZ28     (C) DEC *    R1       RAID   unk        N   N
                             |          |        220    DISK220 RZ28     (C) DEC      "        "      "          "   "
                             |          |        230    DISK230 RZ28     (C) DEC      "        "      "          "   "
                             |          |        240    DISK240 RZ28     (C) DEC      "        "      "          "   "
  D=Scroll down  U=Scroll up |          |        250    DISK250 RZ28     (C) DEC      "        "      "          "   "
Enter menu choice (1,6) [6] ?
```

5.3.4 Adding to Sparesets

Enter option 4 from the RAIDset menu to configure sparesets and failedsets associated with RAIDsets. From the spareset/failedset menu shown in Figure 5–9, enter option 1 to add a device to the spareset. CFMENU prompts you for which devices from the configured PTLs list that you wish to include in the spareset. In Figure 5–10, two devices, PTL 310 and PTL 320, were added to the spareset.

Figure 5–9 CFMENU Spareset/Failedset Menu

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
SPARESET/FAILEDSET MENU:       |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Add a device to the         | Dev.PTLs |       PTLs     Name          ID            Name     Type     Size  sp.  it'd duc  Unit P B
   SPARESET (eligible          |----------|      ------ --------- ----------------  --------- ------ ------- --- ---- --- ----- - -
   devices marked by ^)        |          |disks:  310 ^  DISK310 RZ28     (C) DEC                          N    Y
2. Remove a device from the    |          |        320 ^  DISK320 RZ28     (C) DEC                          N    Y
   SPARESET                    |          |        330 ^  DISK330 RZ28     (C) DEC                          N    Y
3. Move a device from a        |          |        340 ^  DISK340 RZ28     (C) DEC                          N    Y
   RAIDSET to the FAILEDSET    |          |        350 ^  DISK350 RZ28     (C) DEC                          N    Y
   (eligible devices marked    |          |        410 ^  DISK410 RZ28     (C) DEC                          N    Y
   by *)                       |          |        420 ^  DISK420 RZ28     (C) DEC                          N    Y
4. Remove a device from the    |          |        430 ^  DISK430 RZ28     (C) DEC                          N    Y
   FAILEDSET                   |          |        440 ^  DISK440 RZ28     (C) DEC                          N    Y
5. Return to RAIDSET menu      |          |        450 ^  DISK450 RZ28     (C) DEC                          N    Y
                               |          |        510 ^  DISK510 RZ28     (C) DEC                          N    Y
                               |          |        520 ^  DISK520 RZ28     (C) DEC                          N    Y
  D=Scroll down  U=Scroll up   |          |        530 ^  DISK530 RZ28     (C) DEC                          N    Y
Enter menu choice (1,6) [6] ?1
```

Figure 5–10 CFMENU Showing Created Spareset

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
SPARESET/FAILEDSET MENU:       |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Add a device to the         | Dev.PTLs |       PTLs     Name          ID            Name     Type     Size  sp.  it'd duc  Unit P B
   SPARESET (eligible          |----------|      ------ --------- ----------------  --------- ------ ------- --- ---- --- ----- - -
   devices marked by ^)        |          |        150    DISK150 RZ28     (C) DEC      "        "      "          "
2. Remove a device from the    |          |        200    DISK200 RZ28     (C) DEC      S2      STRP    unk        N
   SPARESET                    |          |        300    DISK300 RZ28     (C) DEC      "        "      "          "
3. Move a device from a        |          |        400    DISK400 RZ28     (C) DEC      "        "      "          "
   RAIDSET to the FAILEDSET    |          |        500    DISK500 RZ28     (C) DEC      "        "      "          "
   (eligible devices marked    |          |        600    DISK600 RZ28     (C) DEC      "        "      "          "
   by *)                       |          |raid5:  210    DISK210 RZ28     (C) DEC      R1      RAID    unk        N   N
4. Remove a device from the    |          |        220    DISK220 RZ28     (C) DEC      "        "      "          "   "
   FAILEDSET                   |          |        230    DISK230 RZ28     (C) DEC      "        "      "          "   "
5. Return to RAIDSET menu      |          |        240    DISK240 RZ28     (C) DEC      "        "      "          "   "
                               |          |        250    DISK250 RZ28     (C) DEC      "        "      "          "   "
                               |          |spare:  310    DISK310 RZ28     (C) DEC
  D=Scroll down  U=Scroll up   |          |        320    DISK320 RZ28     (C) DEC
Enter menu choice (1,5) [5] ?
```

After adding sparesets, return to the main menu via the RAIDset menu.

5.3.5 Adding Passthrough Containers

Enter option 4 from the main menu to add passthrough containers. From the passthrough menu (see Figure 5–11), enter option 1. CFMENU prompts you (y/n) whether you want to create a passthrough container from the eligible devices. Figure 5–12 shows the device LDR120 has been created as a passthrough container.

Figure 5–11 CFMENU Passthrough Menu

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
PASSTHROUGH MENU:             |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Create a passthrough       | Dev.PTLs |       PTLs     Name          ID            Name     Type     Size  sp.  it'd duc  Unit P B
   (eligible devices marked   |----------|      ------ --------- ----------------  --------- ------ ------- ---  ---- --- ----- - -
   by ^)                      |          |        500    DISK500 RZ28     (C) DEC                          N    Y       D500 N Y
2. Delete an unbounded        |          |        510    DISK510 RZ26     (C) DEC                          N    Y       D510 N Y
   passthrough (marked        |          |        520    DISK520 RZ26     (C) DEC                          N    Y       D520 N Y
   by *)                      |          |        530    DISK530 RZ28B    (C) DEC                          N    Y       D530 N Y
3. Delete all unbounded       |          |        540    DISK540 RZ28B    (C) DEC                          N    Y       D540 Y N
   passthroughs (marked       |          |        550    DISK550 RZ26     (C) DEC                          N    Y       D550 N Y
   by *)                      |          |        600    DISK600 RZ26     (C) DEC                          N    Y       D600 N Y
4. Return to main menu        |          |        610    DISK610 RZ26     (C) DEC                          N    Y       D610 N Y
                              |          |        620    DISK620 RZ26     (C) DEC                          N    Y       D620 N Y
                              |          |        630    DISK630 RZ28     (C) DEC                          N    Y       D630 N Y
                              |          |        640 ^  DISK640 RZ28     (C) DEC                          N    Y
                              |          |        650    DISK650 RZ26     (C) DEC                          N    Y       D650 N Y
  D=Scroll down  U=Scroll up  |          |loadr:  120 ^   LDR120 TL820    (C) DEC
------------------------------------------------------------------------------------------------------------------------------------
Enter menu choice (1,4) [4] ?1
Create passthrough to device LDR120 (y/n/q) [n] ?y
```

Figure 5–12 CFMENU Passthrough Menu Showing a Passthrough Container

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
PASSTHROUGH MENU:             |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Create a passthrough       | Dev.PTLs |       PTLs     Name          ID            Name     Type     Size  sp.  it'd duc  Unit P B
   (eligible devices marked   |----------|      ------ --------- ----------------  --------- ------ ------- ---  ---- --- ----- - -
   by ^)                      |          |        500    DISK500 RZ28     (C) DEC                          N    Y       D500 N Y
2. Delete an unbounded        |          |        510    DISK510 RZ26     (C) DEC                          N    Y       D510 N Y
   passthrough (marked        |          |        520    DISK520 RZ26     (C) DEC                          N    Y       D520 N Y
   by *)                      |          |        530    DISK530 RZ28B    (C) DEC                          N    Y       D530 N Y
3. Delete all unbounded       |          |        540    DISK540 RZ28B    (C) DEC                          N    Y       D540 Y N
   passthroughs (marked       |          |        550    DISK550 RZ26     (C) DEC                          N    Y       D550 N Y
   by *)                      |          |        600    DISK600 RZ26     (C) DEC                          N    Y       D600 N Y
4. Return to main menu        |          |        610    DISK610 RZ26     (C) DEC                          N    Y       D610 N Y
                              |          |        620    DISK620 RZ26     (C) DEC                          N    Y       D620 N Y
                              |          |        630    DISK630 RZ28     (C) DEC                          N    Y       D630 N Y
                              |          |        640 ^  DISK640 RZ28     (C) DEC                          N    Y
                              |          |        650    DISK650 RZ26     (C) DEC                          N    Y       D650 N Y
  D=Scroll down  U=Scroll up  |          |pass:   120     LDR120 TL820    (C) DEC *    P1     PASS
------------------------------------------------------------------------------------------------------------------------------------

Enter menu choice (1,4) [4] ?
```

5.3.6 Initializing Containers

Enter option 5 from the main menu to initialize containers (devices or storagesets). From the initialization menu (see Figure 5–13), enter option 1. CFMENU prompts (y/n) whether you want to initialize each eligible container.

In addition, CFMENU prompts you to decide on other operating qualifiers, depending on whether the container is a device, stripeset, or RAIDset.

In Figure 5–14, the stripesets S1 and S2 are initialized. Notice that each now has a chunksize specified.

Figure 5–13 CFMENU Initialization Menu

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
INITIALIZATION MENU:          |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In- Re-       W W
1. Initialize a device or     | Dev.PTLs |       PTLs     Name          ID            Name    Type    Size  sp. it'd duc  Unit P B
   storageset (eligible       |----------|      ------ --------- ----------------  --------- ------ ------- --- ---- --- ----- - -
   entities marked with ^)    |          |        150    DISK150 RZ28     (C) DEC      "        "      "         "
2. Return to main menu        |          |        200    DISK200 RZ28     (C) DEC ^    S2     STRP    unk       N
                              |          |        300    DISK300 RZ28     (C) DEC      "        "      "         "
                              |          |        400    DISK400 RZ28     (C) DEC      "        "      "         "
                              |          |        500    DISK500 RZ28     (C) DEC      "        "      "         "
                              |          |        600    DISK600 RZ28     (C) DEC      "        "      "         "
                              |          |raid5:  210    DISK210 RZ28     (C) DEC ^    R1     RAID    unk       N   N
   WARNING!  Initialization of|          |        220    DISK220 RZ28     (C) DEC      "        "      "         "   "
   any device or storageset   |          |        230    DISK230 RZ28     (C) DEC      "        "      "         "   "
   will destroy all of its    |          |        240    DISK240 RZ28     (C) DEC      "        "      "         "   "
   current data.              |          |        250    DISK250 RZ28     (C) DEC      "        "      "         "   "
                              |          |spare:  310    DISK310 RZ28     (C) DEC
  D=Scroll down  U=Scroll up  |          |        320    DISK320 RZ28     (C) DEC
Enter menu choice (1,2) [2] ?
```

Figure 5–14 CFMENU Showing Initialization Results

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
INITIALIZATION MENU:          |Unconfig'd|     Config'd  Device       Product       Stor.set Stor.set Chnk  Trn  In- Re-       W W
1. Initialize a device or     | Dev.PTLs |       PTLs     Name          ID            Name    Type    Size  sp. it'd duc  Unit P B
   storageset (eligible       |----------|      ------ --------- ----------------  --------- ------ ------- --- ---- --- ----- - -
   entities marked with ^)    |          |        610 ^  DISK610 RZ28     (C) DEC                                N    Y
2. Return to main menu        |          |        620 ^  DISK620 RZ28     (C) DEC                                N    Y
                              |          |        630 ^  DISK630 RZ28     (C) DEC                                N    Y
                              |          |        640 ^  DISK640 RZ28     (C) DEC                                N    Y
                              |          |        650 ^  DISK650 RZ28     (C) DEC                                N    Y
                              |          |strps:  100    DISK100 RZ28     (C) DEC ^    S1     STRP     85        Y
                              |          |        110    DISK110 RZ28     (C) DEC      "        "      "         "
   WARNING!  Initialization of|          |        120    DISK120 RZ28     (C) DEC      "        "      "         "
   any device or storageset   |          |        130    DISK130 RZ28     (C) DEC      "        "      "         "
   will destroy all of its    |          |        140    DISK140 RZ28     (C) DEC      "        "      "         "
   current data.              |          |        150    DISK150 RZ28     (C) DEC      "        "      "         "
                              |          |        200    DISK200 RZ28     (C) DEC ^    S2     STRP     85        Y
  D=Scroll down  U=Scroll up  |          |        300    DISK300 RZ28     (C) DEC      "        "      "         "
Enter menu choice (1,2) [2] ?
```

After initializing containers, return to the main menu.

5.3.7 Adding Units

Enter option 6 from the main menu to configure units. From the unit menu, shown in Figure 5–15, enter option 1 to add a unit. CFMENU prompts you for which initialized containers you wish to create units from.

CFMENU also will prompt you to assign a unit number. (The program automatically assigns a "D" or "T" to the unit number when listing the unit, as shown in Figure 5–16.) In addition, CFMENU prompts you to decide on other unit qualifiers.

In Figure 5–16, three units were created from disks at PTLs 330, 340, and 350.

Figure 5–15 CFMENU Unit Menu

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
UNIT MENU:                    |Unconfig'd|     Config'd  Device       Product      Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Create a unit (eligible    | Dev.PTLs |       PTLs     Name         ID           Name     Type     Size  sp.  it'd duc  Unit P B
   entities marked by ^)      |----------|     ------ --------- ----------------  --------- ------ ------- ---  ---- --- ----- - -
2. Delete a unit (eligible    |          |disks:  330 ^  DISK330 RZ28     (C) DEC                          N    Y
   units marked by *)         |          |        340 ^  DISK340 RZ28     (C) DEC                          N    Y
3. Return to main menu        |          |        350 ^  DISK350 RZ28     (C) DEC                          N    Y
                              |          |        410 ^  DISK410 RZ28     (C) DEC                          N    Y
                              |          |        420 ^  DISK420 RZ28     (C) DEC                          N    Y
                              |          |        430 ^  DISK430 RZ28     (C) DEC                          N    Y
                              |          |        440 ^  DISK440 RZ28     (C) DEC                          N    Y
                              |          |        450 ^  DISK450 RZ28     (C) DEC                          N    Y
                              |          |        510 ^  DISK510 RZ28     (C) DEC                          N    Y
                              |          |        520 ^  DISK520 RZ28     (C) DEC                          N    Y
                              |          |        530 ^  DISK530 RZ28     (C) DEC                          N    Y
                              |          |        540 ^  DISK540 RZ28     (C) DEC                          N    Y
  D=Scroll down  U=Scroll up  |          |        550 ^  DISK550 RZ28     (C) DEC                          N    Y
Enter menu choice (1,3) [3] ?1
```

Figure 5–16 CFMENU Showing Created Units

```
------------------------------------------------ CFMENU  Configuration Menu Utility ------------------------------------------------
UNIT MENU:                    |Unconfig'd|     Config'd  Device       Product      Stor.set Stor.set Chnk  Trn  In-  Re-       W W
1. Create a unit (eligible    | Dev.PTLs |       PTLs     Name         ID           Name     Type     Size  sp.  it'd duc  Unit P B
   entities marked by ^)      |----------|     ------ --------- ----------------  --------- ------ ------- ---  ---- --- ----- - -
2. Delete a unit (eligible    |          |disks:  330    DISK330 RZ28     (C) DEC                          N    Y    * D330 N Y
   units marked by *)         |          |        340    DISK340 RZ28     (C) DEC                          N    Y    * D340 N N
3. Return to main menu        |          |        350    DISK350 RZ28     (C) DEC                          N    Y    * D350 N N
                              |          |        410 ^  DISK410 RZ28     (C) DEC                          N    Y
                              |          |        420 ^  DISK420 RZ28     (C) DEC                          N    Y
                              |          |        430 ^  DISK430 RZ28     (C) DEC                          N    Y
                              |          |        440 ^  DISK440 RZ28     (C) DEC                          N    Y
                              |          |        450 ^  DISK450 RZ28     (C) DEC                          N    Y
                              |          |        510 ^  DISK510 RZ28     (C) DEC                          N    Y
                              |          |        520 ^  DISK520 RZ28     (C) DEC                          N    Y
                              |          |        530 ^  DISK530 RZ28     (C) DEC                          N    Y
                              |          |        540 ^  DISK540 RZ28     (C) DEC                          N    Y
  D=Scroll down  U=Scroll up  |          |        550 ^  DISK550 RZ28     (C) DEC                          N    Y
```

5.3.8 Modifying Allocation Classes

During the Software Customization Procedure, you were prompted to type in the MSCP and TMSCP allocation classes of the HS1*xx* server. The HS1CP device channel processors must have the same MSCP and TMSCP allocation class as the server processors. Set the allocation classes with this procedure:

1. From the currently active menu, type in the number to select the option:

   ```
   Return to main menu
   ```

2. Exit from the CFMENU utility to return the the HS1CP> prompt. From the Main Menu type in 8 at the Enter Menu choice:

   ```
   (1,8) [8]?
   ```

3. Enter the following command to determine the MSCP allocation class and the TMSCP allocation class:

   ```
   HS1CP> SHOW THIS
   ```

4. If the display shows the correct allocation class information, then proceed with step 7. If the display shows the incorrect or incomplete allocation class information, then continue.

5. Set the MSCP allocation class by entering this command:

   ```
   HS1CP> SET THIS MSCP_ALLOCATION_CLASS = n
   ```

   Where:

   n is the number you entered for the server.

   The following will be displayed on the terminal screen:

   ```
   Warning 4020: A restart of both this and the other controller is
          required before all the parameters modified will
          take effect.
   Restart of this controller required
   Restart of the other controller required
   ```

6. Set the TMSCP allocation class by entering this command:

   ```
   HS1CP> SET THIS TMSCP_ALLOCATION_CLASS = n
   ```

   Where:

   n is the number you entered for the server.

   The following will be displayed on the terminal screen:

   ```
   Warning 4020: A restart of both this and the other controller is
          required before all the parameters modified will
          take effect.
   Restart of this controller required
   Restart of the other controller required
   ```

7. If you are configuring storage for an HS121 server, copy the configuration from this HS1CP to the other HS1CP with the following command:

   ```
   HS1CP> SET FAILOVER COPY = THIS
   ```

   These messages will appear on the screen:

   ```
   Restart of this controller required
   Restart of the other controller required
   ```

   If you are configuring storage for an HS111 server, proceed to step 9.

8. Restart the other HS1CP device channel processor with this command:

   ```
   HS1CP> RESTART OTHER
   ```

   These messages will appear on the screen:

   ```
   Received LAST GASP message from other controller, Fail Code: 08080000
   HS1CP>
   %PAx0, Software is Closing Virtual Circuit -- REMOTE NODE HS1CPx
   ```

9. Restart the HS1CP device channel processor to which you are currently connected with this command:

   ```
   HS1CP> RESTART THIS
   ```

   These messages will appear on the screen:

   ```
   %PAx0, Software is Closing Virtual Circuit -- REMOTE NODE HS1CPx

   %HSCPAD-F-NOLOCEXE, Local program not executing
   -SYSTEM-F-VCBROKEN, virtual circuit broken
   %HSCPAD-S-END, Control returned to node XXXXXX
   ```

   At this time you will return to the DCL prompt from which you connected to the HS1CP, and your configuration procedure is complete.

5.3.9 Configuring Storage for the Second HS1CP

If you are configuring storage for an HS121 StorageWorks FDDI server, both device channel processors must have the same configuration. This procedure describes how the copy the storage configuration from one device channel processor to another.

If you have just completed the procedure in Section 5.3.8, you have already performed the tasks to copy the storage configuration (steps 7 through 9), and it is not necessary to perform this procedure. You can continue with Section 5.4.

1. To begin, make sure the terminal is displaying the HS1CP> prompt.
   - If you have exited the CFMENU utility, then the terminal should be displaying the HS1CP> prompt.
   - If your terminal is currently displaying a CFMENU display, exit with this procedure:
     a. From the currently active menu, type in the number to select the option:

        ```
        Return to main menu
        ```

     b. Exit from the CFMENU utility to return the the HS1CP> prompt. From the Main Menu, type in 8 at the Enter Menu choice:

        ```
        (1,8) [8]?
        ```

   - If the terminal is displaying the DCL prompt, then you must connect to the device channel processor with this command:

     ```
     $ SET HOST/DUP/SERVER=MSCP$DUP/TASK=CLI node-name
     ```

     Where:

     node-name is the name assigned to the device channel processor (for example, HS1CP1).

2. Copy the configuration from this HS1CP to the other HS1CP with the following command:

   ```
   HS1CP> SET FAILOVER COPY = THIS
   ```

   These messages will appear on the screen:

   ```
   Restart of this controller required
   Restart of the other controller required
   ```

3. Restart the HS1CP device channel processor to which you copied the configuration with this command:

   ```
   HS1CP> RESTART OTHER
   ```

   These messages will appear on the screen:

   ```
   Received LAST GASP message from other controller, Fail Code: 08080000
   HS1CP>
   %PAx0, Software is Closing Virtual Circuit -- REMOTE NODE HS1CPx
   ```

4. Restart the HSC1CP device channel processor to which you are currently connected with this command:

   ```
   HS1CP> RESTART THIS
   ```

   These messages will appear on the screen:

   ```
   %PAx0, Software is Closing Virtual Circuit -- REMOTE NODE HS1CPx

   %HSCPAD-F-NOLOCEXE, Local program not executing
   -SYSTEM-F-VCBROKEN, virtual circuit broken
   %HSCPAD-S-END, Control returned to node XXXXXX
   ```

   At this time, you will return to the DCL prompt from which you connected to the HS1CP, and your configuration procedure is complete.

5.4 Saving the Configuration Setup

After defining the configuration setup, make sure it is printed and kept available to assist in servicing the subsystem in the future. Make a new printout each time you change your configuration.

To capture your device setup, you should perform the following steps:

1. Enter the following command at the DCL propmt to logically connect the StorageWorks FDDI server terminal to the device channel processor and record the results of the session in a file called CONFIG.LOG in the SYS$MANAGER directory on the server's system disk.

   ```
   $ SET HOST/DUP/SERVER=mscp$dup/TASK=CLI/LOG=SYS$MANAGER:CONFIG.LOG node-name
   ```

   Where:

   node-name is the name assigned to the device channel processor (for example, HS1CP1).

   The device channel processor responds with a brief display and its prompt.

   ```
   Copyright (C) Digital Equipment Corporation 1994
   HS1CP Firmware version E35D-0, Hardware version AX01

   Last fail code: 018000A0

   Press " ?" at any time for help.

   HS1CP>
   ```

2. At the HS1CP> prompt enter the following command:

```
HS1CP> SHO DEVICE FULL
```

The following information will be captured in the CONFIG.LOG file in the SYS$MANAGER directory on the server's system disk as it appears on the screen:

```
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

DISK100       disk             1    0    0                    D1100
          DEC      RZ28     (C) DEC D41C
        Switches:
          NOTRANSPORTABLE
DISK110       disk             1    1    0                    D1110
          DEC      RZ28B    (C) DEC 0003
        Switches:
          NOTRANSPORTABLE
DISK120       disk             1    2    0                    D1120
          DEC      RZ28     (C) DEC D41C
        Switches:
          NOTRANSPORTABLE
DISK130       disk             1    3    0                    D1130
          DEC      RZ28     (C) DEC D41C
        Switches:
          NOTRANSPORTABLE
```

3. Exit the display by typing this command at the prompt:

```
HS1CP> EXIT
```

At this point, the HS1CPs and their storage devices have been configured. You should leave the terminal attached to the server processor while you perform other tasks in the installation procedure.

6

Integrating the Server into the VMScluster System

After you have customized the OpenVMS Alpha operating system parameters and configured the storage system, you can complete the installation by connecting the server to the FDDI interconnect and by completing a few additional tasks. In addition, there are other optional postinstallation tasks that you may want to complete. This chapter discusses each task, required and optional, to complete the installation.

6.1 Connecting to the FDDI Interconnect

After the software customization and storage configuration have been completed, you can connect the StorageWorks FDDI server to the FDDI interconnect and boot the server so that it becomes a member of the cluster containing the client nodes. Perform the following steps to accomplish the task:

1. Shutdown the StorageWorks FDDI server OpenVMS Alpha operating system as follows:
   a. Log in to a privileged account such as the system manager account established earlier.
   b. Execute the system shutdown command file. For example, entering the following command invokes a system shutdown:

      ```
      $ @SYS$SYSTEM:SHUTDOWN
      ```

   c. Answer the questions asked during the server's execution of the shutdown command file and wait until the system has completed a logical shutdown.
2. Connect the StorageWorks FDDI server to the FDDI interconnect as follows:
   a. Remove the dust covers attached to the FDDI connections of the StorageWorks FDDI server.
   b. Place the dust covers in a protective container such as a resealable plastic bag. Place the container in a safe place for later use if you disconnect the StorageWorks FDDI server from the FDDI interconnect.
   c. Connect the StorageWorks FDDI server to the FDDI interconnect.
3. Cycle power for the server processor by pressing its power switch off and then on again. After the power-on self-tests complete and the system prompt >>> appears, type B to reboot the system. (See Example 3–1 for an example display of the power-on self-tests.)

As the StorageWorks FDDI server boots, it will automatically request membership in the cluster containing the client nodes. On completion of the boot process, the StorageWorks FDDI server will be able to serve its storage devices to the client nodes.

6.2 Preparing for Operations

At this point, you have set the StorageWorks FDDI server operating system parameters, configured your StorageWorks storage devices, and connected the StorageWorks FDDI server to your VMScluster FDDI interconnect.

6.2.1 Accessing Necessary Documentation

Before completing the StorageWorks FDDI server installation procedure, you may want to access and print some documentation provided on the server system disk. For example, if this is the first Alpha-based node joining your cluster, you should print the *OpenVMS Alpha Version 6.1 Upgrade and Installation Manual* (AA–PV6XB–TE), located on the StorageWorks FDDI server's system disk. This manual provides detailed information on system halt, shutdown, and boot procedures.

The .PS and .TXT documentation for OpenVMS Alpha software is usually found in the [DOCUMENTATION.V0*xx*] directory, where *xx* represents the OpenVMS operating system software version number. For example, for Version 6.1, the directory is [DOCUMENTATION.V061].

When needed, take these steps to print the documentation:

1. Copy the appropriate document from SYS$SYSDEVICE (the server's system disk) to one of the devices on the server that can be accessed by another VMScluster node.
2. From another VMScluster node, print the file on an available printer.

6.2.2 Ensuring VMScluster System Compatibility

After installing the StorageWorks FDDI server software, there remain some tasks that will ensure the system operates correctly in the cluster. The required tasks are as follows:

- Install the Volume Shadowing Interoperability Kit (I14Y) on VAX and/or Alpha nodes in your cluster as appropriate. See the release notes (*AXPSHADnn-061.RELEASE_NOTES* and *VAXSHADnn-061.RELEASE_NOTES* where *nn* refers to the version) found on the server's system disk for reference information.
- Update other nodes in the cluster as required by any changes to the quorum scheme caused by installation of the StorageWorks FDDI server. You must update the EXPECTED_VOTES value in each cluster member's MODPARAMS.DAT file, then run AUTOGEN on each node followed by shutting down and rebooting the cluster. For example, if you add an HS121 and a quorum disk, then each cluster member will need to have its EXPECTED_VOTES parameter increased by 3. See *VMScluster Systems for OpenVMS* for additional information on quorum issues.
- Run AUTOGEN on the server to update cluster parameters. Refer to Section 7.5 for details on running AUTOGEN.

6.2.3 Ensuring Viability of the Write-Back Cache

The write-back cache module installed in your StorageWorks FDDI server contains batteries that were completely charged at the factory. It is normal for these batteries to discharge slightly in shipment.

The server's write-back cache and RAID features require fully-charged batteries to maintain absolute data integrity. After installation, these advanced features may not be available until the batteries have had an opportunity to completely recharge. The charging process may take up to 4 hours to complete.

7

Supporting and Operating the StorageWorks FDDI Server

This chapter discusses system management issues for the StorageWorks FDDI server such as preparing for and installing system management software, setting certain operational parameters, and backing up the server's system disk.

Note

For optimal performance of StorageWorks FDDI servers, you should limit the applications you install on the server processors to storage management and system management applications. By installing and executing production software on the processors, you could degrade the performance of the StorageWorks FDDI server.

You also should consider the kinds of operations you plan to perform with the servers. For example, to realize optimal performance of your StorageWorks FDDI server, you should perform backup operations of online storage to devices controlled by the same server.

7.1 Customizing the System

Manage the StorageWorks FDDI server processors as you would other VMScluster nodes with respect to maintaining a functional, highly reliable computing environment. You can customize system files and install DECnet software and layered products as needed.

Before you make changes to the files on the server system disk, you should make a backup copy (see Section 7.3). If the changes you make cause configuration problems, you can restore your configuration to a known state.

7.1.1 Customizing System Files

You can edit system files that allow you to tailor your system management environment. If you created a quorum disk on the server, you need to add a MOUNT command in the system startup file to mount the quorum disk during server startup.

When customizing your environment, you may decide to set system parameters in the following command files:

- SYCONFIG.COM
- SYLOGICALS.COM

- SYLOGIN.COM
- SYSTARTUP_VMS.COM

The release notes, release notes addendum, and *A Comparison of System Management on OpenVMS Alpha and OpenVMS VAX* contain detailed information for notes and restrictions that might be relevant to your customization plans. Additional information about maintaining the server's performance can be found in the *Guide to OpenVMS Performance Management*.

Note

Immediately after changing your server processor system parameters, you should make a backup copy to protect your modifications.

7.1.2 Configuring and Starting DECnet for OpenVMS Alpha Software

If you plan to run DECnet™ for OpenVMS Alpha software, you must do the following:

1. After you have registered the license for the DECnet for OpenVMS Alpha software, execute the interactive command procedure SYS$MANAGER:NETCONFIG.COM to automatically configure your system for networking. See the *DECnet for OpenVMS Guide to Networking* for instructions on using NETCONFIG.COM.
2. After you start the queue manager (see the *OpenVMS System Manager's Manual*), edit the commands in SYS$COMMON:[SYSMGR]SYSTARTUP_VMS.COM that pertain to networking so that the DECnet for OpenVMS software starts automatically when your system is booted. Choose one of the following commands to start the network and remove the comment delimiter (!) from that command:

```
$! IF F$SEARCH("SYS$SYSTEM:NETACP.EXE") .NES. "" THEN @SYS$MANAGER:STARTNET
$! IF F$SEARCH("SYS$SYSTEM:NETACP.EXE") .NES. "" THEN SUBMIT SYS$MANAGER:STARTNET.COM
```

Both of the previous commands perform the same task. However, the first command executes STARTNET.COM and delays further processing until the procedure is completed; the second submits STARTNET.COM to a batch queue and continues executing SYSTARTUP_VMS.COM.

Note

Immediately after changing your server processor system parameters, you should make a backup copy to protect your modifications.

7.1.3 Installing Layered Products

The StorageWorks FDDI server ships from the factory with all the layered products required for its operation. If you decide to install additional layered products such as Storage Library System for OpenVMS or Host Based RAID, note the following:

- You must register the licenses for OpenVMS Alpha layered products such as Host Based RAID for OpenVMS Alpha software.

  You can invoke the OpenVMS License utility by entering the following command:

  ```
  $ @SYS$UPDATE:VMSLICENSE
  ```

  (You also can use the LICENSE REGISTER command.)

  For additional information about registering licenses, see the *OpenVMS License Management Utility Manual*.

- For general information about installing layered products, see the *OpenVMS System Manager's Manual*. For more specific information, see the documentation you received with each layered product.

The DECwindows components provided with the OpenVMS Alpha Version 6.1–1H2 operating system supply only DECwindows base support and workstation support files. To get full DECwindows support, you also must install the separate DECwindows™ Motif® for OpenVMS Alpha layered product, which supports both the Motif and XUI environments.

See the most recent version of the *DECwindows Motif for OpenVMS Installation Guide* for information about installing the separate DECwindows Motif for OpenVMS Alpha layered product. After the installation, follow the directions in that manual and in *Managing DECwindows Motif for OpenVMS Systems* to customize your DECwindows environment.

Note

Immediately after changing your server processor system parameters, you should make a backup copy to protect your modifications.

7.2 Booting a Server Processor

At various times, you may have to boot the server processor. You can boot a server processor from either the server's system disk or the CD–ROM device.

After the processor has been halted or after it has been powered up, the terminal connected to the processor displays the >>> prompt. To determine the name of the device from which you want to boot the server processor, then boot the device, follow this procedure:

1. Determine the names of the devices available to the server processor. Enter the following command at the >>> prompt:

   ```
   >>> SHOW DEVICE
   ```

A display similar to the following appears:

```
dka0.0.0.6.0      DKA0  ❶ RZ28       D41C
dka500.5.0.6.0    DKA500  ❷ RRD43     0064
dva0.0.0.1        DVA0
era0.0.0.3.1      ERA0       08-00-2B-BC-06-99
fra0.0.0.3.1      FRA0       08-00-2B-A6-0B-D4
pka0.7.0.6.0      PKA0       SCSI BUS ID 7
```

Where:

❶ DKA0 is the name of the server processor system disk, an RZ28 drive

❷ DKA500 is the name of the server processor CD–ROM drive, an RRD43 drive.

2. To boot the server processor, specify the device by name in the boot command. For example, the following command shows how to boot from the CD–ROM shown in the previous example:

```
>>>boot DKA500
```

7.3 Backing Up the Server System Disk

After you create a backup copy of the files on the system disk, you should keep the copy in a safe location. For complete information about backup operations, see the *OpenVMS Alpha Upgrade and Installation Manual*.

Perform the following steps to back up the StorageWorks FDDI server OpenVMS Alpha operating system:

- Shut down the server.

Note

Any node in the VMScluster system whose system disk is served through an HS111 also should be shut down. Shut down that system before shutting down the HS111 server.

- Boot the server from the CD–ROM (see Section 7.2). This starts the menu-driven command procedure shown in the following example.

```
        OpenVMS Alpha (TM) Operating System, Version 6.1-1H2
        Copyright (c) 1994 Digital Equipment Corporation.  All rights reserved.
        Installing required known files...
        Configuring devices...
        ****************************************************************

        You can install or upgrade the OpenVMS Alpha operating system.
        You can also execute DCL commands and procedures to perform
        "standalone" tasks, such as backing up the system disk.

         Please choose one of the following:

         1)  Install or upgrade OpenVMS Alpha Version V6.1-1H2
         2)  Execute DCL commands and procedures
         3)  Shut down this system

      Enter CHOICE or "?" to repeat menu: (1/2/3/?)
```

- Enter 2, then press the Return key.
- At the triple dollar sign prompt ($$$), enter the SHOW DEVICES command.
- Examine the list of devices so you can determine which device is the source drive (the system disk you want to back up) and which device is your target drive (the supported disk or tape device that will hold the backed up files).
- When you have determined which devices will be the source drive and target drive, mount those devices (and any other output devices you plan to use) before you perform any backup operations. Enter the MOUNT commands in the following format:

  ```
  $$$ MOUNT/OVERRIDE=IDENTIFICATION source-drive
  $$$ MOUNT/FOREIGN target-drive
  ```

  Note the following conventions:

  - *source-drive* is the name of the drive holding the system disk.
  - *target-drive* is the name of the drive that will hold the backup files.
- When the system disk and output devices are mounted, back up the system disk by entering the BACKUP command in the following format:

  ```
  $$$ BACKUP/IMAGE/VERIFY source-drive: target-drive:
  ```

  You also must include the save set name and the /SAVE_SET qualifier if the target drive is a tape device.

In this example, the system disk and a target disk are mounted so the BACKUP command can create a backup disk. (You can use a backup disk as a system disk.)

```
$$$ MOUNT/OVERRIDE=IDENTIFICATION DKA200
$$$ MOUNT/FOREIGN DKA300
$$$ BACKUP/IMAGE/VERIFY DKA200: DKA300:
```

In this example the system disk and a target tape device are mounted so the BACKUP command can create a backup tape.

```
$$$ MOUNT/OVERRIDE=IDENTIFICATION DKA200
$$$ MOUNT/FOREIGN MKA300
$$$ BACKUP/IMAGE/VERIFY DKA200: MKA300:APR_06_BACKUP.BCK/SAVE_SET
```

The BACKUP command creates a system disk that includes a set of volume parameters provided by Digital, including a CLUSTER_SIZE (disk access scheme) that is appropriate for your system. (The CLUSTER_SIZE refers to the way files are stored on the disk, not to VMScluster environments.) You can change most volume parameters later with the SET VOLUME command.

However, to change the CLUSTER_SIZE, you must back up the system disk to a disk that has been previously initialized with the CLUSTER_SIZE that you want. For more information about initializing a disk and using the BACKUP command, see the *OpenVMS System Manager's Manual* and the *OpenVMS System Management Utilities Reference Manual*, and see the description of the INITIALIZE and BACKUP commands in the OpenVMS DCL Dictionary.

After you complete the backup operation, do the following:

- Enter the LOGOUT command to exit from the DCL environment and return to the menu.
- Choose the shutdown option (3).
- After the shutdown completes, boot from the system disk.

7.4 Defining Packet Sizes

The StorageWorks FDDI server operating system is OpenVMS Alpha operating system Version 6.1. The SYSGEN parameter governing packet size has been preset to 4468. All other nodes in the same cluster and on the FDDI interconnect should be set to the same packet size according to the following rules:

IF THE OPERATING SYSTEM OF THE NODE IS . . .	THEN SET THE SYSGEN PARAMETER . . .
OpenVMS VAX Version 6.1	NISCS_MAX_PKTSZ to 4468
OpenVMS Alpha Version 6.1	NISCS_MAX_PKTSZ to 4468
OpenVMS Version 5.5–2	LRPSIZE to 4474

More information on packet sizes and related configuration considerations can be found in the manual titled *VMScluster Systems for OpenVMS*.

7.5 Running AUTOGEN

When you installed the operating system, the system executed the AUTOGEN.COM procedure to set the values of system parameters and the sizes of the page, swap, and dump files according to the system configuration. As a postinstallation procedure, you need to run the AUTOGEN.COM procedure again to properly tune the system.

Run AUTOGEN as follows:

1. After 24 hours of operation, run AUTOGEN in feedback mode and reboot the system.
2. Run AUTOGEN again in feedback mode two workdays later, and then reboot the system.
3. Digital recommends that you run AUTOGEN from SAVPARAMS through TESTFILES on a weekly basis thereafter, and examine AGEN$PARAMS.REPORT to determine the need for additional changes.

Based on your examination of AGEN$PARAMS.REPORT, you might need to modify parameter values in MODPARAMS.DAT. If so, note the following:

- Hardcoded values in MODPARAMS.DAT should not hinder AUTOGEN's ability to calculate feedback parameters. AUTOGEN generally does not reduce the value of parameters that allocate resources; it considers current parameter values to be minimum values, which means that you do not have to add MIN_* symbols to MODPARAMS.DAT.
- AUTOGEN does increase parameter values according to its calculations unless you have specified explicit or maximum values (by adding MAX_* symbols) in MODPARAMS.DAT.

For more information about the MODPARAMS.DAT file and about using AUTOGEN in general, see the *OpenVMS System Manager's Manual*. Also, see the *OpenVMS System Manager's Manual: Tuning, Monitoring and Complex Systems* for information on running AUTOGEN and for additional information on system tuning.

7.6 Balancing I/O Load on the HS121

When a device is accessed on a server from multiple client nodes, performance can be enhanced if the access pathways used in device mounting are different. In the HS121 FDDI Server, the two processors and two HS1CP controllers typically offer four possible paths to a given disk device. By distributing mount connections to a device, bandwidth contention along a given path is minimized. When the load is balanced across the two server processors, the server provides better performance than if one server processor was heavily loaded while the other was mostly idle.

Load balancing can be implemented in one or both of two domains: at the device bus level, and at the server/network (mount path) level.

7.6.1 Device-Level Balancing

To ensure that all resources connected to the device bus are utilized fairly equally in dual-controller HS121 systems, the HS1CP provides a means of tuning their usage via the PREFER qualifier. This can be done on one of the HS1CPs:

```
HS1CP> SET unit PREFERRED_PATH=HS1CP_NAME
```

where *unit* specifies a device and *HS1CP_NAME* is THIS (for the HS1CP currently used) or OTHER (for the alternate HS1CP).

7.6.2 Mount Path-Level Balancing via the DCL PREFER Command

A DCL-level command, PREFER, can be incorporated into the DCL command table to provide a means of explicitly specifying a preferred disk path at mount time, or to change an existing mount path. You can use the PREFER command to set up device paths for each VMScluster node to take maximum advantage of the pathways available.

The PREFER command files should be available in the SYS$EXAMPLES directory of your system disk (as PREFER.*). To set up the PREFER command:

1. Place PREFER.EXE in the SYS$SYSTEM directory. (Another directory may be used; edit the 'image' line in file PREFER.CLD to specify the directory.)
2. To define PREFER in the local process DCL command table, issue the command:

   ```
   $ SET COMMAND PREFER
   ```

   To define PREFER in the systemwide DCL command table, issue the following command from an account with CMKRNL privilege:

   ```
   $ SET COMMAND /TABLE=SYS$LIBRARY:DCLTABLES -
   _$ /OUTPUT=SYS$COMMON:[SYSLIB]DCLTABLES PREFER
   ```

   VMS Install is needed to make PREFER available to other users, again requiring CMKRNL privilege:

   ```
   $ INSTALL
   INSTALL> REPLACE SYS$LIBRARY:DCLTABLES
   INSTALL> EXIT
   ```

Note

The Install utility commands must be executed from all nodes in a VMScluster system. If your system configuration deviates from the standard usage of DCLTABLES, modify the above commands to suit your site. For further information on defining command verbs or using the Install utility, consult the VMS documentation set.

To use the preferred path functionality, enter the command:

```
$ PREFER unit/HOST=host[/FORCE]
```

where *unit* specifies a device and *host* is the name of the server node that constitutes the preferred path. /FORCE is an optional qualifier described below.

In the example below, device 10DUA10: has a primary path through node FSERV1 and a secondary path through node FSERV2. To select FSERV1 as the primary path, enter the following command:

```
$ PREFER $10$DUA10:/HOST=FSERV1
```

This command sets the preferred path on the local node so that the next mount will use the selected path. If the /CLUSTER qualifier is used on the MOUNT command, all nodes in the cluster will use the selected path. Note that no change is made to the SHOW DEVICE output.

The /FORCE qualifier is used to select a preferred path for mounted disks. If the disk is mounted (not /FOREIGN), then the /FORCE qualifier will force the drive into mount verification and move it to the new controller. For example:

```
$ PREFER $10$DUA10:/HOST=FSERV1/FORCE
```

The path used to remount the disk will be the preferred path of the node that performs the mount verification. No other nodes in the VMScluster system will alter their paths to the device. For proper operation, the /FORCE qualifier should be issued only from a host on which the device is mounted.

To select a preferred path for all nodes in a VMScluster system, use the VMS SYSMAN Utility to set the preferred path on all nodes in the cluster before executing the PREFER command with the /FORCE qualifier.

Note

The PREFER command must be defined as a DCL command verb on all nodes in the VMScluster system before using it within SYSMAN.

In the example below, the SYSMAN PREFER command sets the preferred path on all nodes so that the next mount command from any node will use the selected path. The DCL PREFER command using the /FORCE qualifier causes the device to enter mount verification. After mount verification completes, the device will be remounted by all nodes through the preferred path, FSERV1.

```
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO PREFER $10$DUA10:/HOST=FSERV1
SYSMAN> EXIT
$ PREFER $10$DUA10:/HOST=FSERV1/FORCE
```

8

System Software Maintenance

Much of the information contained in this chapter is found in the current *OpenVMS Alpha Upgrade and Installation Manual*, available in hardcopy as well as documentation files on the server system disk and CD–ROM. These documents also include other information that can be used in maintaining server and client nodes in a VMScluster environment.

8.1 Upgrading the OpenVMS Alpha Software Version

Note

Digital recommends that you first make a backup copy of your system disk prior to upgrading it. If there is any problem during the upgrade that might affect the integrity of the disk, you will have the backup copy as a safeguard.

Perform the following steps to upgrade the StorageWorks FDDI server OpenVMS Alpha operating system:

1. Shut down the server.

Note

Any node in the VMScluster system whose system disk is served through an HS111 also should be shut down. Shut down that system before shutting down the HS111 server.

2. Boot the server from the CD–ROM (see Section 7.2). This starts the menu-driven command procedure shown in the following example:

```
        OpenVMS Alpha (TM) Operating System, Version 6.1-1H2
        Copyright (c) 1994 Digital Equipment Corporation.  All rights reserved.
        Installing required known files...
        Configuring devices...
        ************************************************************

        You can install or upgrade the OpenVMS Alpha operating system.
        You can also execute DCL commands and procedures to perform
        "standalone" tasks, such as backing up the system disk.

         Please choose one of the following:

         1)  Install or upgrade OpenVMS Alpha Version V6.1-1H2
         2)  Execute DCL commands and procedures
         3)  Shut down this system

 Enter CHOICE or "?" to repeat menu: (1/2/3/?)
```

3. Enter 1, then press the Return key. The install or upgrade script (the beginning of which is shown in the following example) starts by providing brief instructions that help you provide answers to questions asked during the remainder of the script.

```
        ************************************************************

        The installation procedure will ask a series of questions.

              () - encloses acceptable answers
              [] - encloses default answers

         Type your response and press the <Return> key.  Type:

              ? - to repeat an explanation
              ^ - to change prior input (not always possible)

            There are two choices for Installation/Upgrade:

            INITIALIZE - removes all software and data files that were
                previously on the target disk and installs
    OpenVMS Alpha.

            PRESERVE -- installs or upgrades OpenVMS Alpha on the target disk
                and retains all other contents of the target disk.

           * NOTE: You cannot use PRESERVE to install OpenVMS Alpha on a disk on
                which OpenVMS VAX or any other operating system is installed.

         Do you want to INITIALIZE or to PRESERVE? [PRESERVE]
```

Because you are upgrading the OpenVMS Alpha operating system, choose the default reply, PRESERVE. When you specify the PRESERVE option, the following operations take place:

IF . . .	THEN . . .
The OpenVMS Alpha operating system is not already installed on the target disk,	The following operations take place: • The OpenVMS Alpha operating system is installed. • All other contents of the target disk are retained, or preserved.
The OpenVMS Alpha operating system is installed on the target disk,	The OpenVMS Alpha operating system is upgraded, as follows: • Old operating system files and new files are merged or replaced. • All other contents of the target disk are retained, or preserved.

4. Follow the instructions provided by the remainder of the upgrade script.

Refer to the *OpenVMS Alpha Upgrade and Installation Manual* for detailed information on upgrading the OpenVMS Alpha operating system.

8.2 Restoring the System Disk from a Backup

Perform the following steps to restore the StorageWorks FDDI server OpenVMS Alpha operating system:

1. Shut down the server.

   Note

   Any node in the VMScluster system whose system disk is served through an HS111 also should be shut down. Shut down that system before shutting down the HS111 server.

2. Boot the server from the CD–ROM (see Section 7.2). This starts the menu-driven command procedure shown in the following example.

```
        OpenVMS Alpha (TM) Operating System, Version 6.1-1H2
        Copyright (c) 1994 Digital Equipment Corporation.  All rights reserved.
        Installing required known files...
        Configuring devices...
        ************************************************************

        You can install or upgrade the OpenVMS Alpha operating system.
        You can also execute DCL commands and procedures to perform
        "standalone" tasks, such as backing up the system disk.

         Please choose one of the following:

         1)  Install or upgrade OpenVMS Alpha Version V6.1-1H2
         2)  Execute DCL commands and procedures
         3)  Shut down this system

 Enter CHOICE or "?" to repeat menu: (1/2/3/?)
```

3. Enter 2, then press the Return key to begin a limited DCL environment that is identified by a triple dollar sign prompt ($$$).
4. At the triple dollar sign prompt ($$$), enter the SHOW DEVICES command.
5. Examine the list of devices so you can determine which device is the source drive (the drive holding the backed up files you want to restore) and which device is your target drive (the disk on which you want the files restored).
6. When you have determined which devices will be the source drive and target drive, mount those devices (and any other output devices you plan to use) before you perform any restore operations. Enter the MOUNT commands in the following format:

```
$$$ MOUNT/OVERRIDE=IDENTIFICATION source-drive
$$$ MOUNT/FOREIGN target-drive
```

   where:

   source-drive is the device holding the files you want to restore.
   target-drive is the destination.

   (Note, however, that you must use the MOUNT/FOREIGN command if the source drive is a tape device.)

7. Enter the BACKUP command in the following format:

   ```
   $$$ BACKUP/IMAGE/VERIFY source-drive: target-drive:
   ```

   (You also must include the save set name and the /SAVE_SET qualifier if the source drive is a tape device.)

In this example, a backup disk and a target disk are mounted so the BACKUP command can restore the system disk from the backup disk:

```
$$$ MOUNT/OVERRIDE=IDENTIFICATION DKA300
$$$ MOUNT/FOREIGN DKA200
$$$ BACKUP/IMAGE/VERIFY DKA300: DKA200:
```

In this example, a backup tape and a target disk are mounted so the BACKUP command can restore the system disk from the backup tape:

```
$$$ MOUNT/FOREIGN MKA300
$$$ MOUNT/FOREIGN DKA200
$$$ BACKUP/IMAGE/VERIFY MKA300:APR_06_BACKUP.BCK/SAVE_SET DKA200:
```

After you complete the restore operation, do the following:

- Enter the LOGOUT command to exit from the DCL environment and return to the menu.
- Choose the shutdown option (3).
- After the shutdown completes, boot from the system disk.

8.3 Rebuilding the System Disk from the Server CD–ROM

Perform the following steps to rebuild the StorageWorks FDDI server OpenVMS Alpha operating system:

1. Shut down the server.

Note

Any node in the VMScluster system whose system disk is served through an HS111 also should be shut down. Shut down that system before shutting down the HS111 server.

2. Boot the server from the CD–ROM (see Section 7.2). This starts the menu-driven command procedure shown in the following example:

```
OpenVMS Alpha (TM) Operating System, Version 6.1-1H2
Copyright (c) 1994 Digital Equipment Corporation.  All rights reserved.
Installing required known files...
Configuring devices...
************************************************************

You can install or upgrade the OpenVMS Alpha operating system.
You can also execute DCL commands and procedures to perform
"standalone" tasks, such as backing up the system disk.

 Please choose one of the following:

 1)  Install or upgrade OpenVMS Alpha Version V6.1-1H2
 2)  Execute DCL commands and procedures
 3)  Shut down this system

Enter CHOICE or "?" to repeat menu: (1/2/3/?)
```

3. Enter 1, then press the Return key. The install or upgrade script (the beginning of which is shown in the following example) starts by providing brief instructions that help you provide acceptable answers to questions asked during the remainder of the script.

```
************************************************************

The installation procedure will ask a series of questions.

      () - encloses acceptable answers
      [] - encloses default answers

 Type your response and press the <Return> key.  Type:

       ? - to repeat an explanation
       ^ - to change prior input (not always possible)

    There are two choices for Installation/Upgrade:

    INITIALIZE - removes all software and data files that were
        previously on the target disk and installs OpenVMS Alpha.

    PRESERVE -- installs or upgrades OpenVMS Alpha on the target disk
        and retains all other contents of the target disk.

   * NOTE: You cannot use PRESERVE to install OpenVMS Alpha on a disk on
        which OpenVMS VAX or any other operating system is installed.

 Do you want to INITIALIZE or to PRESERVE? [PRESERVE]
```

4. Because you are rebuilding the OpenVMS Alpha operating system, choose INITIALIZE as your reply. When you specify INITIALIZE, the following operations take place:

   - All software and data files that were previously on the target disk are removed.
   - The OpenVMS Alpha operating system is installed.

For further information on the operating system installation refer to the *OpenVMS Alpha Upgrade and Installation Manual*.

A

Working with RAIDsets

Note

For discussions in this manual, the term RAIDsets refers to RAID level 5, and the term stripesets refers to RAID level 0.

The disk striping and RAID level 5 facilities of the HS1CP provide you with a variety of options for controlling the cost, performance, and data availability characteristics of disk storage attached to the controllers. The disk striping facility included with the basic firmware, provides high I/O performance for applications requiring either high I/O request rates or high data transfer rates. The RAID facility available as an optional feature in conjunction with the write-back cache module combines elements of RAID level 5 and RAID level 3 technology to provide affordable data availability (compared to RAID level 0) without the I/O performance penalties usually associated with RAID level 5.

To use these RAID facilities effectively, you need to make some configuration decisions. Study the controller's CLI commands listed in this appendix and in Appendix B to configure your RAIDsets and stripesets.

In the HS1CP firmware, the RAID level 5 facility uses a distributed data mapping technique just like that used for disk striping. A powerful measure of protection against hardware component failure is added by reserving some of the blocks in each RAIDset's disk containers for the storage of redundant information. This redundant information allows the contents of any block of application data stored in the RAIDset to be *regenerated* in the case of a disk container failure on which the data is stored (as long as the remaining RAIDset members are functioning properly).

RAID levels 3 and 5 are sometimes called parity RAID levels, because the redundant information they store is in the form of a parity block that corresponds to data blocks in each of the RAIDset's disk containers. *Parity* is any kind of checksum that allows the regeneration of unretrievable data. Parity is typically combined with data stored in positionally corresponding blocks of other disk containers in the RAIDset to regenerate the missing data.

For detailed information about RAID technology, refer to *The RAIDBOOK—A Source for RAID Technology* published by The RAID Advisory Board, St. Peter, MN.

A.1 RAID Level 0

RAID level 0 is known as striping. Striping spreads data across multiple disks, breaking the user data into segments designated as "chunks." For example, in a four disk stripeset, A, B, C, and D, the first chunk is written on disk A, the second on disk B, the third on disk C, the fourth on disk D, the fifth on disk A, and so on.

CAUTION

If any member of a RAID level 0 stripeset fails, all data is lost from the entire set.

The system administrator sets the chunksize based upon application requirements. If the chunksize is set to be relatively large related to the average input/output (I/O) size, all of the disks may be able to execute different read/write requests simultaneously. If there are large numbers of frequently accessed files, this may be especially beneficial.

If the chunksize is set significantly smaller than the average I/O size, then most or all of the disks in the stripeset will be able to transfer data for a single request in parallel. This method increases the data transfer rate for large I/Os.

RAID level 0 provides high performance for a wide variety of I/O intensive applications. Depending on the hardware configuration and the chunksize set, RAID level 0 improves either data transfer rate or I/O request rate.

A.2 RAID Level 5

RAID level 5 stripes data and rotates parity across all disks in the RAIDset. The controller combines incoming data with existing parity data.

RAID level 5 is suited for applications whose I/O loads consist predominantly of a large number of asynchronous read requests. Transaction processing and office automation applications often fall into this category. It also is good for data transfer intensive applications, such as image analysis, which make mostly read requests. It is not as well suited for write intensive applications (such as data entry, or scientific or engineering data collection).

Note

If using RAID level 5, all data in the RAIDset will be lost if a second drive fails in the same set before the first failed drive is repaired.

A.3 RAID Level 3

Industry standard RAID level 3 achieves higher bandwidths by transferring data to and from the disks, as a result of transferring a part of each I/O's data from each RAIDset member in parallel. To achieve high bandwidths with conventional fixed-block disks (typically 512 data bytes in size), all I/O requests must specify an amount of data equal to the member block size, multiplied by the number of members in the RAIDset, minus one. Also, the requests' starting addresses must be aligned so that correspondingly located data from each member is transferred. To permit this data transfer to take place in parallel, industry standard RAID

level 3 often requires special disks or configurations to ensure that all disks in the RAIDset are rotating in perfect synchronization.

Industry standard RAID level 3 performs as though the RAID level 3 RAIDset is a single disk with a specific large (virtual) sector size. This results in substantial performance penalties for I/Os that are not perfectly aligned multiples of the 2048 (or larger) data byte size. Few applications use extremely large I/O sizes (and these may not easily be modified to use a multiple of the RAID level 3 virtual sector size). In any event, many operating systems cannot easily accommodate virtual disks with unconventional sector sizes.

Digital's implementation of RAID level 3 for the HS1CP achieves higher bandwidth levels without the virtual sector size or special device/configuration disadvantages. This is achieved with special algorithms related to RAID level 5 technology, but without the write performance penalty associated with conventional RAID level 5 (not occurring with conventional RAID level 3) implementation.

For convenience, this capability is controlled by setting the RAIDset's chunksize to a lower value and performing many sequential write operations (in write-back mode). This permits higher bandwidth performance results approaching industry standard RAID level 3 operation. With the capability of setting chunksize, you can conveniently choose between more bandwidth-oriented or more throughput-oriented performance using the same configuration and CLI commands. When you specify intermediate chunksizes, you realize very large I/O benefits from RAID level 3 technology while getting small I/O benefits from RAID level 5 technology.

Note

References to RAID level 5 in this document relate to higher level RAID operation for the HS1CP.

Chunksize is set with the CLI INITIALIZE CHUNKSIZE= command (refer to Appendix B).

A.4 RAIDset Terminology

The following terms are used with HS1CP RAIDset technology:

- Chunksize—The number of blocks written to one RAIDset member before data is written to the next RAIDset member.
- Container—Any entity capable of storing data, whether it is a physical device or a group of physical devices. A disk, a stripeset, and a RAIDset are examples of containers.
- Failedset—A group of disk drives that have been removed from RAIDsets due to a failure or manual removal. Disk drives in the failedset should be considered defective and should be tested and repaired before being placed in the spareset pool or back in their original locations.
- Parity—Any kind of checksum that allows the regeneration of unretrievable data. Parity is typically combined with data stored in positionally corresponding blocks of other disk containers in the RAIDset to regenerate the missing data.

- RAIDset—A virtual disk drive with its physical data spread across multiple physical disks. A RAIDset contains parity data to be used to regenerate data in the event that one member fails.
- RAIDset states
  - Normal state—All members are present and all data is redundant.
  - Reduced state—A failed RAIDset member has been detected and removed from the RAIDset.
  - Reconstructing state—All members are present and redundancy is being restored.
  - Reconstruct types (process of restoring redundancy to the RAIDset)

    There are two different types of reconstruct: one that takes place when a unit is created from a RAIDset; the other takes place during the replace operation of a failed RAIDset member. Each kind indicates a different error recovery operation/choice that the controller will make.

    * Initial reconstruct operation—Establishes initial redundancy following an ADD RAIDset . . . ADD UNIT command sequence. Note that all data written by the host is immediately fully redundant.

      When a RAIDset is initialized using the INITIALIZE container-name command, the controller does not take the lengthy period of time to make all the parity blocks consistent with the data. Instead, the controller marks all the parity blocks as bad and starts a reconstruct operation. The reconstruct operation recalculates and rewrites the parity blocks and marks them as good. This process allows the RAIDset to be used immediately. All new data written to the RAIDset is immediately fully redundant.

    * Reconstructing a replaced member—Regenerates the data for that member and restores redundancy.

      To reconstruct means to restore redundancy. This could be either recalculating the parity or recalculating a user data block using the remaining blocks.

      When a reduced RAIDset has a member added back to it, all the blocks on the replacement member are marked as bad (parity and user data blocks) and a reconstruct scan is started. The reconstruct operation recalculates the parity blocks on the new member, recalculates the user data blocks on the new member, writes the blocks, and marks the blocks good.

Note

If a second RAIDset member fails during a reconstruct operation, the RAIDset goes inoperative.

To determine which type of reconstruct operation is taking place, enter the SHOW RAIDset command. If a reconstructing member is not identified, the controller is performing an initial reconstruct.

- Redundancy—A RAIDset is considered to be redundant when user data is recorded directly to one member, and all of the other members and associated parity also are recorded. If a member is missing from the RAIDset, its data can be regenerated as needed, but the RAIDset is no longer redundant until the missing member is replaced and reconstructed.
- Regenerate—The process of calculating missing data from the redundant data.
- Replacement policy—The firmware controlled method by which a spare disk is selected to replace a disk that has failed in a RAIDset. Your replacement policy choices are BEST_FIT, BEST_PERFORMANCE, or NOPOLICY. Refer to the ADD RAIDSET and SET *raidset-container-name* CLI commands in Appendix B.
- Spareset—A pool of disk devices available to the controller to replace failed RAID level 5 RAIDset members.
- Stripe—The data and parity from the associated chunks of each member of the RAIDset.
- Stripe size—The capacity determined by *n*–1 times the chunksize. (*n* is the number of RAIDset members.)
- Stripeset—A virtual disk drive with its physical data spread across multiple physical disks.

A.5 RAIDset Rules and Important Information

The following list gives rules to remember about RAIDsets:

- You must always have a write-back cache module when creating RAID level 5 RAIDsets.
- Do not attempt to use any RAIDset commands with mismatched cache modules. Both cache modules must be write-back cache, and both must have the same number of megabytes.
- You must purchase licenses for write-back caching and RAID level 5 functionality. If you turn on write-back caching without entering a valid license key, you will receive an hourly error message at the terminal, and an hourly error in the host error log. Once you enter your key, the error message stops.

Note

You may activate write-back caching and RAID level 5 after turning them on via the firmware licensing system (FLS) utility.

- NOWRITEBACK_CACHE (write-through caching) is automatically set for units created from RAIDsets. To increase the unit's performance, switch to WRITEBACK_CACHE.
- RAIDsets may be made up of 3 to 14 members.
- For HS1CP, RAIDsets in a dual-redundant configuration flush cache and failover to the companion controller if the write-back cache module battery for one controller in the pair has a low charge.

- RAIDsets will go inoperative (and write-protected) if *both* write-back cache modules' batteries fail, or if a RAIDset belongs to a nonredundant controller. Any other storagesets (stripesets) will flush cache data and become write-through (cache module acts as read cache) instead of write-back. RAIDsets become operational when the batteries become fully charged again.

  Note

  When the power for the write-back cache module batteries is too low, a console message is displayed. You can check the status of the batteries at any time by entering the CLI SHOW THIS_CONTROLLER command (or SHOW OTHER_CONTROLLER as appropriate).

- Your RAIDsets can achieve the high performance characteristics of RAID level 3 provided you set your RAIDset chunksize to the minimum value (16) and your application calls for large sequential I/O operations.
- RAIDsets can contain disks of different sizes, but the disk space used by the RAIDset is limited to the maximum size of the smallest disk in the RAIDset.
- Place RAIDset members on different ports. This keeps your RAIDset from going inoperative in the event that a single port bus failure occurs.

A.6 Planning Your RAIDsets

The following items should be considered before creating your RAIDsets:

- RAIDset size (3 to 14 members)
- RAIDset chunksize
- RAIDset replacement policy
- RAIDset reconstruction rate
- RAIDset spares
- RAIDset configurations for availability and performance
- RAIDset hardware requirement

A.6.1 Creating a RAIDset

Enter the following commands to create a RAIDset:

1. Use the ADD DISK *container-name SCSI-location* command to add new disk drives to your configuration and name them.
2. Use the ADD RAIDSET *container-name container-name1 container-name2 [container-nameN]* command and set the appropriate replacement policy and reconstruct qualifiers. You are not required to set a replacement policy.

   The following is an example for using the replacement policy and reconstruct qualifiers on the same command line as the ADD RAIDSET command:

```
HS1CP1> ADD RAIDSET R3 DISK100 DISK200 DISK300 POLICY=BEST_FIT
    RECONSTRUCTION=NORMAL
```

Where:

R3 is the name assigned to the RAIDset.
DISK100, *DISK200*, and *DISK300* are the names assigned to the RAIDset members of a three member RAIDset.

POLICY=BEST_FIT is the replacement policy qualifier that will be used when a RAIDset member fails. (You can choose between three different replacement policy qualifiers: POLICY=BEST_FIT, POLICY=BEST_PERFORMANCE, or NOPOLICY.)

RECONSTRUCTION= *rate or no rate* is the reconstruction rate qualifier that is used when a RAIDset member fails and a new member is taken from the spareset as a replacement for the failed device. You can choose between three reconstruct qualifiers: RECONSTRUCT=NORMAL (the default), RECONSTRUCT=FAST, or NORECONSTRUCT.

3. Enter the INITIALIZE command for your RAIDset. This is the time to specify your chunksize. The metadata on the container (in this case, the RAIDset) must be initialized before a unit may be created from it. If the container's metadata cannot be found, or is incorrect, an error will be displayed and the unit will not be created.

   ```
   HS1CP1> INITIALIZE R3 CHUNKSIZE=n
   ```

   where *n* is the chunksize in blocks

   or

   ```
   HS1CP1> INITIALIZE R3 CHUNKSIZE=DEFAULT
   ```

   where the controller determines the optimal chunksize.

4. Enter the ADD UNIT *unit-number container-name* command to create a logical unit from the RAIDset for the host to access, followed by the appropriate qualifier for cache transfer sizes, preferred path, cache access, write protection, and so forth (as described in Appendix B).

   ```
   HS1CP1> ADD UNIT D170 R3
   ```

   where D170 is the unit name, and R3 is the RAIDset name.

5. Enter the SHOW RAIDSETS command to display all of the RAIDsets known to the controller. By adding the FULL qualifier, more information concerning all of the known RAIDsets is displayed.

   To show information about a particular RAIDset, enter the SHOW *raidset-container-name* command. Where *raidset-container-name* is the name assigned to the particular RAIDset.

6. Using the ADD SPARESET command, populate the spareset pool with disk drives that closely match the geometry of the other disk drives in your subsystem.

Refer to Appendix B for descriptions and examples for choosing and using the appropriate ADD UNIT qualifiers.

A.6.2 Storageset SHOW Commands

The following CLI SHOW commands are used to display the status of your RAIDsets, stripesets, and storagesets.

Enter the following CLI command to display all RAIDsets known to the controller:

```
HS1CP1> SHOW RAIDSETS
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

RAID0         raidset                       DISK100          D0
                                            DISK300
                                            DISK400
                                            DISK500
```

Enter the following CLI command to display additional information about all RAIDsets known to the controller:

```
HS1CP1> SHOW RAIDSETS FULL
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

RAID0         raidset                       DISK100          D0
                                            DISK300
                                            DISK400
                                            DISK500
        Switches:
          CHUNKSIZE = 64 blocks
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
        State:
          RECONSTRUCT 3% complete on member DISK500
```

Enter the following CLI command to display information about a particular RAIDset:

```
HS1CP1> SHOW RAID0
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

RAID0         raidset                       DISK100          D0
                                            DISK300
                                            DISK400
                                            DISK500
        Switches:
          CHUNKSIZE = 64 blocks
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
        State:
          RECONSTRUCT 3% complete on member DISK500
```

Enter the following CLI command to display all information about all stripesets known to the controller:

```
HS1CP1> SHOW STRIPESETS FULL
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

STRIPE0       stripeset                     DISK110          D1
                                            DISK210
        Switches:
          CHUNKSIZE = 64 blocks
```

Enter the following CLI command to display all information about all storagesets known to the controller:

```
HS1CP1> SHOW STORAGESETS FULL
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

STRIPE0       stripeset                     DISK110          D1
                                            DISK210
        Switches:
          CHUNKSIZE = 64 blocks
```

```
RAID0         raidset                        DISK100          D0
                                             DISK300
                                             DISK400
                                             DISK500
        Switches:
          CHUNKSIZE = 64 blocks
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
        State:
          RECONSTRUCT 4% complete on member DISK500
SPARESET      spareset                       DISK310
                                             DISK600
FAILEDSET     failedset                      DISK200
```

A.6.3 Adding and Deleting Spareset Members

The spareset is a pool of disk drives available to the controller to replace failing members of a RAIDset. The ADD SPARESET command adds disk drives to the spareset pool and initializes the metadata on the drives so they may be used for replacements into RAIDsets. The DELETE SPARESET command removes disk drives from the spareset.

Note

The spareset cannot be deleted, it is always available. It may contain no spares.

Enter the following CLI commands to add one or more disks to the spareset:

```
HS1CP1> ADD SPARESET disk-container-name0 [disk-container-nameN]
Example: HS1CP1> ADD SPARESET DISK100
         HS1CP1> ADD SPARESET DISK100 DISK200 DISK300
```

Enter the following CLI commands to remove one or more disks from the spareset:

```
HS1CP1> DELETE SPARESET disk-container-name0 [disk-containter-nameN]
Example: HS1CP1> DELETE SPARESET DISK100
         HS1CP1> DELETE SPARESET DISK100 DISK200 DISK300
```

Enter the following command to show the spareset:

```
HS1CP1> SHOW SPARESET
Name          Storageset                     Uses             Used by
------------------------------------------------------------------------------
SPARESET      spareset                       DISK310
                                             DISK600
```

A.6.4 Showing and Deleting Failedset Members

A failedset is a group of disk drives that were removed from RAIDsets because they failed or were manually removed (via the SET *RAIDset-container-name REMOVE=disk-container-name command*). Drives in the failedset should be considered defective. These drives must be tested and repaired before placing them back in operation.

The DELETE FAILEDSET command removes drives from the failedset so that they can be physically removed from the device shelves for testing and repair. Enter the following commands to show, and then remove, one or more disk drives from the failedset:

```
HS1CP1> DELETE FAILEDSET DISK99
HS1CP1> DELETE FAILEDSET DISK99 DISK88 DISK77
```

Enter the following CLI command to show a failedset:

```
HS1CP1> SHOW FAILEDSET
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

FAILEDSET     failedset                     DISK200
```

Note

A failedset cannot be deleted; it is always available. It may contain no devices.

A.6.5 Changing RAIDset Characteristics

To change certain characteristics of a RAIDset, use the SET *RAIDset-container-name* command.

When a RAIDset looses a member, a new member is automatically added to the RAIDset from the spareset pool (providing you have a replacement policy set, and an appropriate spare is in the spareset). If you specified NOPOLICY, when you created your RAIDset or you want to change your replacement policy, enter one of the following commands:

```
HS1CP1> SET RAIDset-container-name POLICY=BEST_FIT
HS1CP1> SET RAIDset-container-name POLICY=BEST_PERFORMANCE
HS1CP1> SET RAIDset-container-name NOPOLICY
```

To change the speed at which a RAIDset will be reconstructed when a new member is added to the RAIDset or immediately after the RAIDset is initialized, enter one of the following commands:

```
HS1CP1> SET RAIDset-container-name RECONSTRUCT=NORMAL
HS1CP1> SET RAIDset-container-name RECONSTRUCT=FAST
```

If you do not want your RAIDset to be reconstructed, enter the following command:

```
HS1CP1> SET RAIDset-container-name NORECONSTRUCT
```

If you need to remove a disk member from a RAIDset, enter the following command:

```
HS1CP1> SET RAIDset-container-name REMOVE=disk-container-name
```

For example:

```
HS1CP1> SET R3 REMOVE=DISK100
```

If the RAIDset is already in a reduced state when the REMOVE= qualifier is used, an error is printed and the command is rejected. If a replacement policy is specified, the replacement drive is automatically taken from the spareset to replace the removed member using the specified policy.

If NOPOLICY is specified, the RAIDset continues to operate in a reduced state until a replacement is manually specified or a policy is specified. The disk drive removed via the REMOVE= qualifier is automatically added to the failedset.

To manually place a disk member into a reduced RAIDset when NOPOLICY was specified, enter the following command:

```
HS1CP1> SET RAIDset-container-name REPLACE=disk-container-name
```

For example:

```
HS1CP1> SET R3 REPLACE=DISK550
```

Where *R3* is the RAIDset name, and *DISK550* is the replacement disk name.

The disk called DISK550 is added to the reduced RAIDset (R3). A reconstruct operation begins immediately on the newly added disk (as long as the reconstruct is not disabled).

Note

No other qualifiers can be used with the SET *RAIDset-container-name* command when either the REPLACE or REMOVE qualifiers are specified.

A.6.6 Deleting a RAIDset

Use the DELETE *container-name* command to delete a RAIDset. This command determines whether the container (RAIDset) is used by a unit. If the container is in use, an error is printed and the container is not deleted. If the container is not is use, it is deleted.

Enter the following command to delete a RAIDset:

```
HS1CP1> DELETE container-name
```

For example:

```
HS1CP1> DELETE R3
```

Where *R3* is the name of the RAIDset being deleted.

A.6.7 Moving a RAIDset

You may physically relocate some or all of a RAIDset's member devices according to the following procedure:

CAUTION

If you lose track of the RAIDset members at any point during this procedure, you will have to attempt to restore the RAIDset by guessing where its members are installed. There is currently no way to retrace your steps using the controller or HS operating firmware.

To move a RAIDset you must do the following:

1. Make note of all devices comprising the RAIDset. Digital recommends marking them after using the HS1CP1> LOCATE command to find all RAIDset members.
2. Delete the UNIT that uses the RAIDset with the DELETE *unit-number* command.
3. Delete the RAIDset with the DELETE *container-name* command.
4. Delete each disk from that RAIDset with the DELETE *container-name* command.
5. Physically remove the disks from the storage shelf.
6. Move the disks to the new port/target/LUN (PTL) location.
7. Add each disk with the ADD DISK *container-name SCSI-location* command using the new **PTL** location.
8. Re-add the RAIDset with the ADD RAIDSET *container-name container-name1 container-name2 [container-nameN]* command. Make sure you create it from the exact, original set of drives.

   CAUTION

   Do *not* initialize the RAIDset or you will destroy its data.

9. Recreate the logical unit from the RAIDset with the ADD UNIT *unit-number container-name* command.

The following example shows the unit "D100" made of RAIDset "RAID99." "RAID99" has member disks at PTLs 200, 210, and 400. The member at PTL 210 can be relocated to PTL 300 as follows:

```
HS1CP1> DELETE D100
HS1CP1> DELETE RAID99
HS1CP1> DELETE DISK210
(Move the disk to PTL 300.)
HS1CP1> ADD DISK DISK300 3 0 0
HS1CP1> ADD RAIDSET RAID99 DISK200 DISK300 DISK400
HS1CP1> ADD UNIT D100 RAID99
```

If you move a RAIDset from one controller to another and you damage one member, you must specify all of that RAIDset's members when you re-add the RAIDset to the new controller. The controller will automatically reduce the RAIDset when it discovers that one member is inoperative.

Using the REDUCED Qualifier with the ADD RAIDset Command

Only use the REDUCED qualifier (with the ADD RAIDSET command) when you want to move a RAIDset that is already reduced. For example, you have a four member RAIDset that has been reduced to a three member RAIDset on Controller A and you wish to move the RAIDset to Controller B. Physically move the three members to Controller B and enter the following command:

```
HS1CP1> ADD RAIDSET container-name container-name1 container-name2
    container-name3 REDUCED
```

For example:

```
HS1CP1> ADD RAIDSET R3 DISK100 DISK300 DISK400 REDUCED
```

A.7 Adding a Stripeset (RAID Level 0)

Use the ADD STRIPESET *container-name container-name1 container-name2* command to add a stripeset and to name that stripeset. This command must be used when a new stripeset is added to a controller's configuration. A stripeset may contain from 2 to 14 members. To create a stripeset, add the individual disks, add the stripeset and name it, initialize the stripeset, and then create and name a unit from the stripeset as shown in the following example:

```
HS1CP1> ADD DISK DISK99 1 0 0
HS1CP1> ADD DISK DISK88 2 0 0
HS1CP1> ADD DISK DISK77 3 0 0
HS1CP1> ADD STRIPESET STRIPE0 DISK99 DISK88 DISK77
HS1CP1> INITIALIZE STRIPE0
HS1CP1> ADD UNIT D0 STRIPE0
```

A.7.1 Moving a Stripeset or Stripeset Member

You may physically relocate some or all of a stripeset's member devices. However, if you lose track of the stripeset members at any point during the relocation, you will have to attempt to restore the stripeset by guessing where its members are installed. There is currently no way to retrace your steps using the controller or HS operating firmware.

You the same procedure as described in Section A.6.7 to move a stripeset or stripeset member.

A.7.2 Showing Stripesets

The SHOW STRIPESET command displays all the stripesets known by the controller. The SHOW STRIPESET FULL gives more information about all stripesets known to the controller. By entering the SHOW *stripeset-container-name* command, you are given specific information about a particular stripeset:

```
HS1CP1> SHOW STRIPESET
HS1CP1> SHOW STRIPESET FULL
HS1CP1> SHOW stripeset-container-name
HS1CP1> SHOW UNITS FULL
```

A.8 Configuring RAIDsets for Availability, Performance, and Cost

RAIDset size (and RAID level) recommendations depend on whether availability, performance, or cost is the priority for creating RAIDsets. Tradeoffs must be made, because no single RAID level provides the perfect balance of availability, performance, and cost. You need to determine what your priorities are before creating your RAIDsets or stripesets.

For availability and performance, it is important to put each RAIDset member on a different port (bus). This keeps the RAIDset from going inoperative in the event of a single port failure, and also provides better performance.

RAID level 5 is more economical for large RAIDsets than smaller RAIDsets because the cost of the parity blocks is amortized across a larger number of devices. However, large RAIDsets statistically have higher failure rates.

Stripesets provide high performance and a lower cost (no parity disk to buy), but do not provide redundancy for availability.

B

Device Channel Processor Command Line Interpreter

B.1 Overview

The command line interpreter (CLI) is the user interface to the HS1CP. The CLI allows you to add to or modify the HS1CP procssor's configuration by using CLI commands. The following sections explain how to define a configuration and modify it when needed. A detailed description of the CLI commands is given in Section B.2.

B.1.1 CLI Access

After initial configuration, you can access the CLI using a maintenance terminal or a virtual terminal. The actual method of establishing the virtual terminal connection varies depending on your operating system and interface type.

B.1.2 CLI Command Sets

CLI consists of the following six command sets:

1. CLI failover commands—These commands are used to support dual-redundant HS1CP configurations.
2. CLI HS1CP commands—These commands are used for the following:
   a. Sets and shows the basic HS1CP parameters
   b. Sets the HS1CP ID
   c. Sets the resident terminal characteristics
   d. Restarts the HS1CP
   e. Runs resident diagnostics and utility programs
3. CLI device commands—These commands allow you to specify and show the location of the *physical* SCSI–2 devices attached to each HS1CP. Locations of devices are specified using their SCSI port/target/LUN (PTL) designation.

   Note

   Only devices that have been defined by the ADD command are recognized by the HS1CP. Devices that have been placed in a storage shelf, but have not been defined using the ADD command are NOT recognized by the HS1CP. Locations of devices are specified using a Port-Target-LUN (**PTL**) format. (Other restrictions may apply when adding removable media devices.)

4. CLI storageset commands—These commands are used for the following:
   a. Adding
   b. Modifying (setting)
   c. Renaming
   d. Showing storagesets (stripesets, RAIDsets, and so forth)
5. CLI logical unit commands—These commands are used to add, modify, or show logical units that have been built from devices or storagesets.
6. CLI diagnostic and utility commands—These commands are used for general support functions on the HS1CP. They are used to invoke exercisers that test the data transfer capabilities of disk and tape drives. The two exercisers are DILX (a disk inline exerciser) and TILX (a tape inline exerciser).

Remember these two rules when using the CLI:

- Not all configuration parameters need to be specified on one line. They can be entered by using multiple SET commands.
- Only enough of each command needs to be entered to make the command unique (usually three characters). For example, SHO is equivalent to SHOW.

B.1.3 How to Exit the CLI

If you are using a maintenance terminal, you cannot exit the CLI. Entering the EXIT command merely restarts the CLI and redisplays the HS1CP type, and the last fail error information as shown in Example B–1.

Example B–1 CLI EXIT Command Message

```
HS1CP1> EXIT

Copyright  Digital Equipment Corporation 1994
HS1CP Firmware version E35D-0, Hardware Version AX01

Last fail code: 018700A0

Press " ?" at any time for help.

HS1CP1>
```

If you are using a DUP connection for the CLI, type EXIT after the CLI prompt to return control to the host.

If you are connected to a virtual terminal (via DUP for OpenVMS), and you specified the /LOG=CONFIGURATION.INFO qualifier on your command line, a log file of your session in CLI is created. You must use the EXIT command to exit the CLI in order to print the CONFIGURATION.INFO file.

B.1.4 Setting Configuration Parameters for an HS1CP Configuration

CAUTION

Unless you have a mating guide installed on the HS1CP processor's host port connector, you must turn off all power to all other devices, including the host CPU, on a DSSI bus before connecting or disconnecting a DSSI host port cable.

Without the mating guide, if you accidentally short DSSI connector pins during aligning and connecting/disconnecting the connector, you risk blowing the fuses of *all* members of the DSSI bus.

With the mating guide installed, you may disconnect or connect the DSSI host port cable connector with power applied to the HS1CP and other members on the host bus. However, Digital recommends carefully connecting or disconnecting the DSSI connector *at all times*, even if you have the mating guide installed.

B.1.5 Setting Configuration Parameters for an HS1CP Configuration

Perform the following steps to identify the devices, storagesets, and logical units to the host. If these steps are not completed, the host does not recognize the devices. These steps can be done from a virtual (host) terminal.

CAUTION

Unless you have a mating guide installed on the HS1CP processor's host port connector, you must turn off all power to the HS1CP processor's and all other devices, including the host CPU, on a DSSI bus before connecting or disconnecting a DSSI host port cable.

Without the mating guide, if you accidentally short DSSI connector pins during aligning and connecting/disconnecting the connector, you risk blowing the fuses of *all* members of the DSSI bus.

With the mating guide installed, you may disconnect or connect the DSSI host port cable connector with power applied to the HS1CP and other members on the host bus. However, Digital recommends carefully connecting or disconnecting the DSSI connector *at all times*, even if you have the mating guide installed.

Some DSSI host port cable connectors are too large to allow access to the trilink connector screws. You must remove the trilink with the host port cables attached to warm swap your HS1CP. You cannot perform HS1CP warm swap if you have this situation.

Newer host port cable connectors have a slot in the side of the host port connector to allow access to the trilink with a very small flathead screwdriver.

To automatically configure devices on the HS1CP, use either the Configure (CONFIG) or CFMENU utility. For manual configuration, use the following steps to add devices, storagesets, and logical units.

1. Add the physical devices by using the following command:

   ```
   HS1CP1> ADD device-type device-name scsi-location
   ```

   For example:

   ```
   HS1CP1> ADD DISK DISK0 1 0 0
   HS1CP1> ADD TAPE TAPE0 5 1 0
   HS1CP1> ADD CDROM CDROM0 6 0 0
   ```

   where:

   device-type is the type of device to be added. This can be DISK, TAPE, or CDROM.

   device-name is the name to refer to that device. The name is referenced when creating units or storagesets.

   SCSI-location is the HS1CP processor's port, target, and LUN for the device. (When typing the port, target, and LUN designations, at least one space must separate the port, target, and LUN numbers.)

   You can use the ADD *device-type* command anytime you need to add a device to the configuration.

   Note

   If you add a removable media device to an HS1CP, it is not known to the host until one of the following occurs: the media is loaded into the device, the host is rebooted, or the virtual circuit is broken and reestablished. This behavior is a feature of MSCP/TMSCP.

---

2. Add the storagesets for the devices, for example:

```
HS1CP1> ADD STRIPESET STRIPE0 DISK0 DISK1 DISK2 DISK3
 or
HS1CP1> ADD RAIDSET RAID99 DISK0 DISK1 DISK2
```

   Refer to Section B.2 for examples for adding storagesets (for example, ADD STRIPESET or ADD RAIDSET). If you do not need storagesets for your configuration, skip this step.

3. Enter the INITIALIZE command to initialize *containers* (devices or storagesets) prior to adding them to the configuration.

```
HS1CP1> INITIALIZE container-name
```

   For example:

```
HS1CP1> INITIALIZE DISK0
HS1CP1> INITIALIZE STRIPE0
HS1CP1> INITIALIZE RAID99
```

   CAUTION

   The INITIALIZE command destroys all customer data on the container. Refer to Section B.2 for specific requirements for using the INITIALIZE command.

   When initializing a single disk drive container, if the NOTRANSPORTABLE qualifier was specified or allowed to default on the ADD DISK or SET *disk-name* commands, a small amount of disk space is made inaccessible to the host and used for metadata. The metadata will be initialized. However, if the TRANSPORTABLE qualifier was specified, any metadata is destroyed on the device and the full device is accessible to the host. Refer to Section B.2 for details of when initializing is required and when it is not required.

4. Add the units that use either the devices directly or the storagesets built from the devices by entering the following command at the HS1CP1> prompt:

```
HS1CP1> ADD UNIT logical-unit-number container-name
```

   ```
For example:  HS1CP1>ADD UNIT D300 DISK0
              HS1CP1>ADD UNIT D200 RAID99
   ```

   where:

   logical-unit-number is the unit number the host uses to access the device.

   container-name identifies the device or the storageset.

B.2 CLI Commands

The following sections describe each of the valid commands in the CLI, along with their required parameters and qualifiers. Examples are given after the command format, parameters, description, and qualifiers.

ADD CDROM

Adds a CDROM drive to the list of known CDROM drives.

Format

ADD CDROM *container-name SCSI-location*

Parameters

container-name
Specifies the name that is used to refer to this CDROM drive. This name is referred to when creating units. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

SCSI-location
The location of the CDROM drive to be added in the form PTL where **P** designates the port (1–6), **T** designates the target ID of the device, (0–6, in a nonfailover configuration, or 0–5 if the HS1CP is in a failover configuration), and **L** designates the LUN of the device (must be 0).

When entering the PTL, at least one space must separate the the port, target, and LUN numbers.

Description

Adds a CDROM drive to the list of known CDROM drives and names the drive. This command must be used when a new SCSI–2 CDROM drive is to be added to the configuration.

Examples

1.
```
HS1CP1> ADD CDROM CD_PLAYER 1 0 0
```
Adds a CDROM drive to port 1, target 0, LUN 0, and named CD_PLAYER.

ADD DISK

Adds a disk drive to the list of known disk drives.

Format

ADD DISK *container-name SCSI-location*

Parameters

container-name
Specifies the name that is used to refer to this disk drive. This name is referred to when creating units and stripesets. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

SCSI-location
The location of the disk drive to be added in the form PTL where **P** designates the port (1–6), **T** designates the target ID of the device, (0–6, in a nonfailover configuration, or 0–5 if the HS1CP is in a failover configuration), and **L** designates the LUN of the device (must be 0).

When entering the PTL, at least one space must separate the the port, target, and LUN numbers.

Description

Adds a disk drive to the list of known disk drives and names the drive. This command must be used when a new SCSI–2 disk drive is to be added to the configuration.

Qualifiers

TRANSPORTABLE
NOTRANSPORTABLE (Default)
In normal operations, the HS1CP makes a small portion of the disk inaccessible to the host and uses this area to store metadata, which improves data reliability, error detection, and recovery. This vast improvement comes at the expense of transportability.

If NOTRANSPORTABLE is specified and there is no valid metadata on the unit, the unit must be initialized.

Note

Digital recommends that you avoid specifying TRANSPORTABLE unless transportability of disk drive or media is imperative and there is no other way to accomplish moving the data.

Examples

1.
```
HS1CP1> ADD DISK RZ26_100 1 0 0
```

   Adds a non transportable disk to port 1, target 0, LUN 0 and names it RZ26_100.

2.
```
HS1CP1> ADD DISK DISK0 2 3 0 NOTRANSPORTABLE
```

   Adds a non transportable disk to port 2, target 3, LUN 0 and names it DISK0.

3.
```
HS1CP1> ADD DISK TDISK0 3 2 0 TRANSPORTABLE
```

   Adds a transportable disk to port 3, target 2, LUN 0 and names it TDISK0.

ADD LOADER

Adds a loader to the list of known loaders.

Format

ADD LOADER *container-name SCSI-location*

Parameters

container-name
Specifies the name that is used to refer to this loader. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

SCSI-location
The location of the loader to be added in the form PTL where **P** designates the port (1–6), **T** designates the target ID of the device, (0–6, in a nonfailover configuration, or 0–5 if the HS1CP is in a failover configuration), and **L** designates the LUN of the device (must be 0).

When entering the PTL, at least one space must separate the the port, target, and LUN numbers.

Description

Adds a loader to the list of known loaders and names the loader. This command must be used when a new SCSI–2 loader is to be added to the configuration.

Examples

1.
```
HS1CP1> ADD LOADER LOAD0 2 0 0
```

   Adds a loader to port 2, target 0, LUN 0 and names it LOAD0.

ADD OPTICAL

Adds an optical drive to the list of known optical drives.

Format

ADD OPTICAL *container-name SCSI-location*

Parameters

container-name
Specifies the name that is used to refer to this optical drive. This name is referred to when creating units and stripesets. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

SCSI-location
The location of the optical drive to be added in the form PTL where **P** designates the port (1–6), **T** designates the target ID of the device, (0–6, in a nonfailover configuration, or 0–5 if the HS1CP is in a failover configuration), and **L** designates the LUN of the device (must be 0).

When entering the PTL, at least one space must separate the the port, target, and LUN numbers.

Description

Adds an optical drive to the list of known optical drives and names the drive. This command must be used when a new SCSI–2 optical drive is to be added to the configuration.

Qualifiers

TRANSPORTABLE
NOTRANSPORTABLE (Default)
In normal operations, the HS1CP makes a small portion of the optical disk inaccessible to the host and uses this area to store metadata, which improves data reliability, error detection, and recovery. This vast improvement comes at the expense of transportability.

If NOTRANSPORTABLE is specified and there is no valid metadata on the unit, the unit must be initialized.

Note

Digital recommends that you avoid specifying TRANSPORTABLE unless transportability of optical drive or media is imperative and there is no other way to accomplish moving the data.

Examples

1. `HS1CP1> ADD OPTICAL OPT0 2 3 0 NOTRANSPORTABLE`

   Adds a non transportable optical drive to port 2, target 3, LUN 0 and names it OPT0.

2. `HS1CP1> ADD OPTICAL TOPT0 3 2 0 TRANSPORTABLE`

   Adds a transportable optical drive to port 3, target 2, LUN 0 and names it TOPT0.

ADD PASSTHROUGH

Creates a command disk (passthrough) container to allow direct access to a device.

Format

ADD PASSTHROUGH *container-name device-name*

Parameters

container-name
Specifies the name that is used to refer to this passthrough container. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

device-name
The device that receives passthrough commands. Only one device may be specified.

Description

Adds a passthrough container to the list of known passthrough containers and names the container. This command must be used when you want to communicate directly to a device using SCSI commands, such as a loader or a disk that is about to have new microcode downline loaded.

Examples

1. `HS1CP1> ADD PASSTHROUGH PASS0 DISK0`

   Creates a passthrough container to disk0 and names it PASS0.

ADD RAIDSET

Creates a RAIDset from a number of containers.

Format

ADD RAIDSET *container-name container-name1 container-name2 [container-nameN]*

Parameters

container-name
Specifies the name that is used to refer to this RAIDset. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

container-name1 container-name2 container-nameN
The containers that will make up this RAIDset. A RAIDset may be made up of from 3 to 14 containers.

Description

Adds a RAIDset to the list of known RAIDsets and names the RAIDset. This command must be used when a new RAIDset is to be added to the configuration.

Qualifiers

POLICY=BEST_FIT
POLICY=BEST_PERFORMANCE (Default)
NOPOLICY
Specifies the replacement policy to use when a member within the RAIDset fails.

BEST_FIT gives highest priority to finding a replacement device within the spareset that most closely matches the sizes of the remaining members of the RAIDset. After finding the most closely matching devices, the device that gives the best performance is selected.

BEST_PERFORMANCE (default) gives highest priority to finding a replacement device within the spareset that results in the best performance of the RAIDset. After finding the best performing devices, the device that most closely matches the size of the remaining members of the RAIDset is selected.

NOPOLICY retires a failing device from the RAIDset without selecting a replacement. This causes the RAIDset to run in a reduced state until a BEST_ FIT or BEST_PERFORMANCE policy is selected, or a member is manually replaced in the RAIDset (see SET *raidset-container-name*).

RECONSTRUCT=NORMAL (Default)
RECONSTRUCT=FAST
NORECONSTRUCT
Specifies the speed at which a RAIDset will be reconstructed when a new member is added to the RAIDset or immediately after the RAIDset is initialized.

RECONSTRUCT=NORMAL (default) balances overall performance of the HS1CP against the demand of reconstructing the RAIDset.

RECONSTRUCT=FAST reconstructs the RAIDset at the fastest rate possible resulting in some loss of performance of the HS1CP overall.

REDUCED
NOREDUCED (Default)
REDUCED specifies that the RAIDset being added is already missing one member. Use the REDUCED keyword when moving an already reduced RAIDset from one HS1CP to another. NOREDUCED (default) identifies that all RAIDset members that make up the RAIDset are being specified.

Examples

1. `HS1CP1> ADD RAIDSET RAID9 DISK0 DISK1 DISK2 DISK3`

   Creates a RAIDset with four disks (DISK0, DISK1, DISK2, and DISK3). The replacement policy is BEST_PERFORMANCE.

2. `HS1CP1> ADD RAIDSET RAID9 DISK0 DISK1 DISK2 DISK3 POLICY=BEST_FIT`

   Creates a RAIDset with four disks (DISK0, DISK1, DISK2, and DISK3). The replacement policy is BEST_FIT, as specified.

3. `HS1CP1> ADD RAIDSET RAID9 DISK0 DISK1 DISK2 DISK3 NOPOLICY`

   Creates a RAIDset with four disks (DISK0, DISK1, DISK2, and DISK3). If a member within the RAIDset fails, a replacement will *not* be selected.

4. `HS1CP1> ADD RAIDSET RAID9 DISK0 DISK1 DISK3 REDUCED`

   Creates a four member RAIDset with a raidset that was already reduced.

ADD SPARESET

Adds a disk drive to the spareset.

Format

ADD SPARESET *disk-container-name0 [disk-container-nameN]*

Parameters

disk-container-name0 disk-container-nameN
The disk drive container names to add to the spareset. Any number of disks may be added to the spareset using only one command.

Description

The **SPARESET** is a pool of drives available to the HS1CP to replace failing members of a RAIDset. The ADD SPARESET command adds disk drives to the spareset and initializes the metadata on the drives so they may be used for replacements by RAIDsets.

Examples

1. `HS1CP1> ADD SPARESET DISK0`

   Adds one disk to the spareset.

2. `HS1CP1> ADD SPARESET DISK0 DISK1 DISK2 DISK3 DISK4`

   Adds five disks to the spareset.

ADD STRIPESET

Creates a stripeset from a number of containers.

Format

ADD STRIPESET *container-name container-name1 container-name2 [container-nameN]*

Parameters

container-name
Specifies the name that is used to refer to this stripeset. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

container-name1 container-name2 container-nameN
The containers that will make up this stripeset. A stripeset may be made up of from 2 to 14 containers.

Description

Adds a stripeset to the list of known stripesets and names the stripeset. This command must be used when a new stripeset is added to the configuration.

Examples

1.
```
HS1CP1> ADD STRIPESET STRIPE0 DISK0 DISK1 DISK2 DISK3
```

   Creates a STRIPESET with four disks (DISK0, DISK1, DISK2, and DISK3).

ADD TAPE

Adds a tape drive to the list of known tape drives.

Format

ADD TAPE *device-name SCSI-location*

Parameters

device-name
Specifies the name that is used to refer to this tape drive. This name is referred to when creating units. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

SCSI-location
The location of the tape drive to be added in the form PTL where **P** designates the port (1–6), **T** designates the target ID of the device, (0–6, in a nonfailover configuration, or 0–5 if the HS1CP is in a failover configuration), and **L** designates the LUN of the device (must be 0).

When entering the PTL, at least one space must separate the the port, target, and LUN numbers.

Description

Adds a tape drive to the list of known tape drives and names the drive. This command must be used when a new SCSI–2 tape drive is to be added to the configuration.

Examples

1.
```
HS1CP1> ADD TAPE TAPE0 1 0 0
```

   Adds a tape drive to port 1, target 0, LUN 0 and names it TAPE0.

ADD UNIT

Adds a logical unit to the HS1CP.

Format

ADD UNIT *unit-number container-name*

Parameters

unit-number
The device type letter followed by the logical unit number (0–4094) that the host uses to access the unit. The device type letter is either "D" for disk devices (including CDROMs) or "T" for tape devices. Using this format, logical unit 3, which is made up of a disk or disks (such as a stripeset), would be specified as D3, and logical unit 7, which is made up of a tape device would be T7.

container-name
The name of the container that is used to create the unit.

Description

The ADD UNIT command adds a logical unit for the host to access. All requests by the host to the logical unit number are mapped as requests to the container specified in the ADD UNIT command.

For disk devices (and stripesets and RAIDsets built from disk devices), the metadata on the container must be initialized before a unit may be created from it. If the container's metadata cannot be found, or is incorrect, an error is displayed and the unit is not created.

Qualifiers for a Unit Created from a CDROM Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

Qualifiers for a Unit Created from a TRANSPORTABLE Disk Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

Qualifiers for a Unit Created from a NOTRANSPORTABLE Disk Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the HS1CP's write-back cache on this unit.

Note

It may take up to 5 minutes to flush unwritten data from the write-back cache once you disable write-back caching.

Note

When initially added, NOWRITEBACK_CACHE is the default.

Qualifiers for a Unit Created from a TRANSPORTABLE Optical Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be

made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

Qualifiers for a Unit Created from a NOTRANSPORTABLE Optical Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the HS1CP's write-back cache on this unit.

Note

It may take up to 5 minutes to flush unwritten data from the write-back cache once you disable write-back caching.

Note

When initially added, NOWRITEBACK_CACHE is the default.

Qualifiers for a Unit Created from a RAIDset

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

Note

Writes may still be performed to a write-protected RAIDset to satisfy a reconstruct pass or to reconstruct a newly replaced member. However, write protect will disable the writing of any new data.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the HS1CP's write-back cache on this unit.

Note

It may take up to 5 minutes to flush unwritten data from the write-back cache once you disable write-back caching.

Qualifiers for a Unit Created from a Stripeset

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the HS1CP's write-back cache on this unit.

Note

It may take up to 5 minutes to flush unwritten data from the write-back cache once you disable write-back caching.

Qualifiers for a Unit Created from a Tape Drive

DEFAULT_FORMAT=*format*
DEFAULT_FORMAT=DEVICE_DEFAULT (Default)
Specifies the tape format to be used unless overridden by the host. Note that not all devices support all formats. The easiest way to determine what formats are supported by a specific device is to enter the "SHOW <tape unit number> DEFAULT_FORMAT= ?" command—the valid options will be displayed.

Supported tape formats are as follows:

- DEVICE_DEFAULT (default)

  The default tape format is the default that the device uses, or, in the case of devices that can be set via switches on the front panel, the settings of those switches.
- 800BPI_9TRACK
- 1600BPI_9TRACK
- 6250BPI_9TRACK
- TZ85
- TZ86
- TZ87_NOCOMPRESSION
- TZ87_COMPRESSION
- DAT_NOCOMPRESSION
- DAT_COMPRESSION
- 3480_NOCOMPRESSION

- 3480_COMPRESSION

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

Examples

1. `HS1CP1> ADD UNIT D0 DISK0`

   Creates disk unit number 0 from container DISK0.

2. `HS1CP1> ADD UNIT T0 TAPE12`

   Creates tape unit number 0 from container TAPE12.

3. `HS1CP1> ADD UNIT D170 RAID9 WRITE_PROTECT`

   Creates disk unit number 170 from container RAID9 and write protects it.

CLEAR_ERRORS CLI

Stops displaying errors at the CLI prompt.

Format

CLEAR_ERRORS CLI

Description

Errors detected by HS1CP firmware are displayed before the CLI prompt. These errors are displayed even after the error condition is rectified, until the HS1CP is restarted or the CLEAR_ERRORS CLI command is issued.

Note

This command does not clear the error conditions, it only clears displaying the errors at the CLI prompt.

Examples

1.
```
HS1CP1>

All NVPM components initialized to their default settings.
HS1CP1> CLEAR_ERRORS CLI
HS1CP1>
```

Clears the message "All NVPM components initialized to their default settings." that was displayed at the CLI prompt.

CLEAR_ERRORS INVALID_CACHE

Clears all data from the cache and makes it usable by the specified HS1CP.

Format

CLEAR_ERRORS INVALID_CACHE *controller*

Parameters

controller
Specifies which HS1CP will clear the INVALID_CACHE condition. Either THIS_CONTROLLER or OTHER_CONTROLLER must be specified.

Description

CAUTION

This command causes loss of customer data.

Note

Because this command causes loss of customer data, "INVALID_CACHE" must be completely spelled out, not abbreviated.

If a write-back cache module with unwritten cache data from another HS1CP is installed on this HS1CP, or if the write-back cache module with unwritten cache data is removed from this HS1CP, an INVALID_CACHE error results. CLEAR_ERRORS INVALID_CACHE clears the invalid cache error, however *all customer data that was in cache is lost*.

For this reason, use great caution when considering using this command.

Examples

1. `HS1CP1> CLEAR_ERRORS INVALID_CACHE THIS_CONTROLLER`

   Clears all cache information from this HS1CP's cache and clears the invalid cache error.

CLEAR_ERRORS LOST_DATA

Clears the lost data error on a unit.

Format

CLEAR_ERRORS LOST_DATA *unit-number*

Parameters

unit-number
Specifies the logical unit number (D0–D4094 or T0–T4094) that will have the lost data error cleared. The *unit-number* is the name given the unit when it was created using the ADD UNIT command.

Description

CAUTION

This command causes loss of customer data.

Note

Because this command causes loss of customer data, "LOST_DATA" must be completely spelled out, not abbreviated.

It may take up to 5 minutes to clear lost data.

If customer data has been lost due to the removal or failure of the write-back cache, the lost data error is reported on the unit. CLEAR_ERRORS LOST_DATA clears the lost data error, however, *all customer data that had not been written to disk is lost.*

For this reason, use great caution when considering using this command.

Examples

1.
```
HS1CP1> CLEAR_ERRORS LOST_DATA D13
```

Clears the lost data error on disk unit D13.

CLEAR_ERRORS UNKNOWN

Clears the UNKNOWN error from a device.

Format

CLEAR_ERRORS UNKNOWN *device-name*

Parameters

device-name
Specifies the device name of the device with the UNKNOWN error.

Description

Note

"UNKNOWN" must be completely spelled out, not abbreviated.

If a device has a failure such that the HS1CP marks the device as UNKNOWN, the device is never automatically checked again to see if it has been repaired or if the failure condition was rectified. When you rectifiy a condition that caused a device to be marked UNKNOWN, this command must be issued for the HS1CP to recognize the device.

Examples

1.
```
HS1CP1> CLEAR_ERRORS UNKNOWN DISK300
```
Causes the HS1CP to recognize DISK300, a previously UNKNOWN device.

CLEAR_ERRORS UNWRITEABLE_DATA

Clears the unwriteable data error on a unit.

Format

CLEAR_ERRORS UNWRITEABLE_DATA *unit-number*

Parameters

unit-number
Specifies the logical unit number (D0–D4094 or T0–T4094) that will have the unwriteable data error cleared. The *unit-number* is the name given the unit when it was created using the ADD UNIT command.

Description

CAUTION

This command causes loss of customer data.

Note

Because this command causes loss of customer data, "UNWRITEABLE_ DATA" must be completely spelled out, not abbreviated.

If a container fails in a way that customer data in the write-back cache cannot be written to the container, the unwriteable data error is reported. CLEAR_ ERRORS UNWRITEABLE_DATA clears the unwriteable data error, however, *all customer data that has not been written to disk is lost.*

For this reason, use great caution when considering using this command.

Examples

1. `HS1CP1> CLEAR_ERRORS UNWRITEABLE_DATA D13`

   Clears the unwriteable data error on disk unit D13.

DELETE *container-name*

Deletes a container from the list of known containers.

Format

DELETE *container-name*

Parameters

container-name
Specifies the name that identifies the container. This is the name given the container when it was created using the ADD command (ADD DEVICE, ADD STRIPESET, and so forth).

Description

Checks to see if the container is used by any other containers or a unit. If the container is in use, an error is displayed and the container is not deleted.

If the container is not in use, it is deleted.

Note

The spareset and failedset containers cannot be deleted. See DELETE SPARESET and DELETE FAILEDSET commands.

Examples

1. `HS1CP1> DELETE DISK0`

   Deletes DISK0 from the list of known containers.

2. `HS1CP1> DELETE STRIPE0`

   Deletes STRIPE0 from the list of known containers.

3. `HS1CP1> DELETE RAID9`

   Deletes RAID9 from the list of known containers.

DELETE FAILEDSET

Deletes a disk drive from the failedset.

Format

DELETE FAILEDSET *disk-container-name0 [disk-container-nameN]*

Parameters

disk-container-name0 disk-container-nameN
The disk drive container names to delete from the failedset. Any number of disks may be deleted from the failedset using only one command.

Description

The **FAILEDSET** is a group of drives that were removed from RAIDsets because they failed or were manually removed (via the SET *raidset-container-name* REMOVE=*disk-container-name* command). Drives in the failedset should be considered defective and should be tested, then repaired or replaced. The DELETE FAILEDSET command removes drives from the failedset, typically before you remove them physically from the shelf for testing, repair, or replacement.

Examples

1. `HS1CP1> DELETE FAILEDSET DISK0`

   Deletes one disk from the failedset.

2. `HS1CP1> DELETE FAILEDSET DISK0 DISK1 DISK2 DISK3 DISK4`

   Deletes five disks from the failedset.

DELETE SPARESET

Deletes a disk drive from the spareset.

Format

DELETE SPARESET *disk-container-name0 [disk-container-nameN]*

Parameters

disk-container-name0 disk-container-nameN
The disk drive container names to delete from the spareset. Any number of disks may be deleted from the spareset using only one command.

Description

The **SPARESET** is a pool of drives available to the HS1CP to replace failing members of a RAIDset. The DELETE SPARESET command removes disk drives from the spareset.

Examples

1. `HS1CP1> DELETE SPARESET DISK0`

   Deletes one disk from the spareset.

2. `HS1CP1> DELETE SPARESET DISK0 DISK1 DISK2 DISK3 DISK4`

   Deletes five disks from the spareset.

DELETE *unit-number*

Deletes a unit from the list of known units.

Format

DELETE *unit-number*

Parameters

unit-number
Specifies the logical unit number (D0–D4094 or T0–T4094) that is to be deleted. The *unit-number* is the name given the unit when it was created using the ADD UNIT command.

Description

The DELETE command flushes any user data from the write-back cache to the disk and deletes the logical unit. If the logical unit specified is online to a host, the unit is not deleted unless the OVERRIDE_ONLINE qualifier is specified. If any errors occur when trying to flush the user data, the logical unit is not deleted.

In order to delete a unit that has cache errors, you must clear all cache errors associated with the unit via a CLEAR_ERRORS command.

Qualifiers

OVERRIDE_ONLINE
NOOVERRIDE_ONLINE (Default)
If the logical unit is online to the HS1CP, it is not deleted unless the OVERRIDE_ONLINE qualifier is specified.

If the OVERRIDE_ONLINE qualifier is specified, the unit is run down, the user data is flushed to disk, and the logical unit is deleted.

CAUTION

Customer data may be lost or corrupted if the OVERRIDE_ONLINE qualifier is specified.

Examples

1. `HS1CP1> DELETE D12`

   Deletes disk unit number 12 from the list of known units.

2. `HS1CP1> DELETE T3 OVERRIDE_ONLINE`

   Deletes tape unit number 3 from the list of known units even if it is currently online to a host.

DIRECTORY

Lists the diagnostics and utilities available on THIS_CONTROLLER.

Format

DIRECTORY

Description

The DIRECTORY command lists the various diagnostics and utilities that are available on THIS_CONTROLLER. A directory of diagnostics and utilities available on this HS1CP is displayed.

For specific information about the diagnostics and utilities available, refer to

Examples

1.
```
HS1CP1> DIRECTORY
TILX  X067   D
DILX  X067   D
VTDPY X067   D
FLS   X067   D
ECHO  X067   D
DIRECTX067   D
CLI   X067   D
```

Displays directory listing.

EXIT

Exits the CLI and breaks the virtual terminal connection.

Format

EXIT

Description

When entering the EXIT command from a host using a virtual terminal connection, the connection is broken and control is returned to the host. If entered from a maintenance terminal, the EXIT command restarts the CLI, displaying the copyright notice, the HS1CP type, and the last fail packet.

Examples

1.
```
HS1CP1> EXIT

Copyright  Digital Equipment Corporation 1994
HS1CP Firmware version E35D-0, Hardware Version AX01

Last fail code: 018700A0

Press " ?" at any time for help.

HS1CP1>
```

An EXIT command issued on a maintenance terminal.

2.
```
HS1CP1> EXIT
Control returned to host
$
```

An EXIT command issued on a terminal that was connected to the CLI via a DUP connection.

HELP

Displays an overview for getting help.

Format

HELP

Description

The HELP command displays a brief description for using the question mark "?" to obtain help on any command or CLI function.

Examples

1.
```
HS1CP1> HELP
Help may be requested by typing a question mark (?) at the CLI prompt.  This
will print a list of all available commands

For further information you may enter a partial command and type a space
followed by a "?" to print a list of all available options at that point in
the command.  For example:

          SET THIS_CONTROLLER ?

Prints a list of all legal SET THIS_CONTROLLER commands
```

   Displaying help using the HELP command.

2.
```
HS1CP1> SET ?
Your options are:
   FAILOVER
   OTHER_CONTROLLER
   NOFAILOVER
   THIS_CONTROLLER
   Unit number or container name
```

   Getting help on the SET command, using the "?" facility.

INITIALIZE

Initializes the metadata on the container specified.

Format

INITIALIZE *container-name*

Parameters

container-name
Specifies the container name to initialize.

Description

The INITIALIZE command initializes a container so a logical unit may be created from it. When initializing a single disk drive container, if NOTRANSPORTABLE was specified or allowed to default on the ADD DISK or SET *disk-name* commands, a small amount of disk space is made inaccessible to the host and used for metadata. The metadata is initialized. If TRANSPORTABLE was specified, any metadata is destroyed on the device and the full device is accessible to the host.

CAUTION

The INITIALIZE command destroys all customer data on the container.

Note

It may take up to 2 minutes to initialize a RAIDset or stripeset.

The INITIALIZE command is required when:

- A unit is going to be created from a newly installed disk
- A unit is going to be created from a newly created storageset, (RAIDset or stripeset)

The INITIALIZE command specifically is *not* required when:

- A unit has been deleted, and a new unit is going to be created from the same container
- A storageset that was initialized in the past is deleted, then added again using the same members that were in the original storageset

Qualifiers

CHUNKSIZE=*n*
CHUNKSIZE=DEFAULT (Default)
Specifies the chunksize to be used. The chunksize may be specified in blocks (CHUNKSIZE=*n*), or you can let the controller determine the optimal chunksize (CHUNKSIZE=DEFAULT).

Examples

1. `HS1CP1> INITIALIZE DISK0`

   Initializes container DISK0. If NOTRANSPORTABLE was specified (or allowed to default), metadata is written on the disk.

2. `HS1CP1> INITIALIZE STRIPE0 CHUNKSIZE=20`

   Initializes container STRIPE0 and writes metadata on it.

3. `HS1CP1> INITIALIZE RAID9 CHUNKSIZE=20`

   Initializes container RAID9 with a chunksize of 20 and writes metadata on it.

LOCATE

Locates units, storagesets, and devices by lighting the amber device fault LED on the front of the StorageWorks building block (SBB).

Format

LOCATE

Description

The LOCATE command illuminates the amber device fault LEDs (the lower LED on the front of an SBB) of the containers specified. The LOCATE command also can be used as a lamp test.

Qualifiers

ALL

The LOCATE ALL command turns on the amber device fault LEDs of all configured devices. This qualifier also can be used as a lamp test. See LOCATE CANCEL to turn off the LEDs.

An error is displayed if no devices have been configured.

CANCEL

The LOCATE CANCEL command turns off all amber device fault LEDs on all configured devices.

An error is displayed if no devices have been configured.

DISKS

The LOCATE DISKS command turns on the amber device fault LEDs of all configured disks. See LOCATE CANCEL to turn off the LEDs.

An error is displayed if no disks have been configured.

UNITS

The LOCATE UNITS command turns on the amber device fault LEDs of all devices used by units. This command is useful to determine which devices are not currently configured into logical units. See LOCATE CANCEL to turn off device the LEDs.

An error is displayed if no units have been configured.

PTL *SCSI-location*

The LOCATE PTL *SCSI-location* command turns on the amber device fault LEDs at the given SCSI location. *SCSI-location* is specified in the form PTL where **P** designates the port (1–6), **T** designates the target ID of the device (0–6 in a nonfailover configuration or 0–5 if the controller is in a failover configuration), and **L** designates the LUN of the device (0–7).

When entering the PTL, at least one space must separate the port, target, and LUN numbers. See LOCATE CANCEL to turn off the LEDs.

An error is displayed if the port, target, or LUN is invalid, or if no device is configured at that location.

device or storageset name or unit number (*entity*)
The LOCATE *entity* command turns on the amber device fault LEDs that make up the entity supplied. If a device name is given, the device's LED is lit. If a storageset name is given, all device LEDs that make up the storageset are lit. If a unit number is given, all device LEDs that make up the unit are lit. See LOCATE CANCEL to turn off the LEDs.

An error is displayed if no entity by that name or number has been configured.

Examples

1. `HS1CP1> LOCATE DISK0`

   Turns on the device fault LED on device DISK0.

2. `HS1CP1> LOCATE D12`

   Turns on the device fault LEDs on all devices that make up disk unit number 12.

3. `HS1CP1> LOCATE DISKS`

   Turns on the device fault LEDs on all configured disk devices.

RENAME

Renames a container.

Format

RENAME *old-container-name new-container-name*

Parameters

old-container-name
Specifies the existing name that identifies the container.

new-container-name
Specifies the new name to identify the container. This name is referred to when creating units and storagesets. The name must start with a letter (A–Z) and can then consist of up to eight more characters made up of letters A–Z, numbers 0–9, periods (.), dashes (-), or underscores (_), for a total of nine characters.

Description

Gives a known container a new name by which to be referred.

Examples

1.
```
HS1CP1> RENAME DISK0 DISK100
```

   Renames container DISK0 to DISK100.

RESTART OTHER_CONTROLLER

Restarts the other HS1CP.

Format

RESTART OTHER_CONTROLLER

Description

The RESTART OTHER_CONTROLLER command flushes all user data from the other HS1CP's write-back cache (if present), then restarts the other HS1CP.

If any disks are online to the other HS1CP, the HS1CP does not restart unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the HS1CP does not restart unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes the other HS1CP to restart immediately without flushing any user data to the disks, even if drives are online to the host.

The RESTART OTHER_CONTROLLER command does not cause a failover to this HS1CP in a dual-redundant configuration. The other HS1CP restarts and resumes operations where it was interrupted.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the HS1CP is not restarted unless IGNORE_ERROR is specified.

If the IGNORE_ERRORS qualifier is specified, the HS1CP restarts even if all customer data cannot be written to disk from the write-back cache.

CAUTION

Customer data may be lost or corrupted if the IGNORE_ERRORS qualifier is specified.

IMMEDIATE
NOIMMEDIATE (Default)
If IMMEDIATE is specified, the HS1CP is immediately restarted without checking for online devices or flushing user data from write-back cache to disk.

CAUTION

Customer data may be lost or corrupted if the IMMEDIATE qualifier is specified.

OVERRIDE_ONLINE
NOOVERRIDE_ONLINE (Default)
If any units are online to the HS1CP, the HS1CP is not restarted unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the HS1CP restarts after all customer data is written to disk.

CAUTION

Customer data may be lost or corrupted if the OVERRIDE_ONLINE qualifier is specified.

Examples

1. `HS1CP1> RESTART OTHER_CONTROLLER`

   Restarts the other HS1CP as long as the other HS1CP does not have any units online.

2. `HS1CP1> RESTART OTHER_CONTROLLER OVERRIDE_ONLINE`

   Restarts the other HS1CP even if there are units online to the other HS1CP.

RESTART THIS_CONTROLLER

Restarts this HS1CP.

Format

RESTART THIS_CONTROLLER

Description

The RESTART THIS_CONTROLLER command flushes all user data from this HS1CP's write-back cache (if present), then restarts this HS1CP.

If any disks are online to this HS1CP, the HS1CP does not restart unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the HS1CP does not restart unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes this HS1CP to restart immediately without flushing any user data to the disks, even if drives are online to a host.

The RESTART THIS_CONTROLLER command does not cause a failover to the other HS1CP in a dual-redundant configuration. This HS1CP restarts and resumes operations where it was interrupted.

Note

If you enter the RESTART THIS_CONTROLLER command and you are using a virtual terminal to communicate with the HS1CP, the connection is lost when this HS1CP restarts.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the HS1CP is not restarted unless IGNORE_ERROR is specified.

If the IGNORE_ERRORS qualifier is specified, the HS1CP restarts even if all customer data cannot be written to disk from the write-back cache.

CAUTION

Customer data may be lost or corrupted if the IGNORE_ERRORS qualifier is specified.

IMMEDIATE
NOIMMEDIATE (Default)
If IMMEDIATE is specified, the HS1CP is immediately restarted without checking for online devices or flushing user data from write-back cache to disk.

CAUTION

Customer data may be lost or corrupted if the IMMEDIATE qualifier is specified.

OVERRIDE_ONLINE
NOOVERRIDE_ONLINE (Default)
If any units are online to the HS1CP, the HS1CP is not restarted unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the HS1CP restarts after all customer data is written to disk.

CAUTION

Customer data may be lost or corrupted if the OVERRIDE_ONLINE qualifier is specified.

Examples

1. `HS1CP1> RESTART THIS_CONTROLLER`

   Restarts this HS1CP as long as this HS1CP does not have any units that are online.

2. `HS1CP1> RESTART THIS_CONTROLLER OVERRIDE_ONLINE`

   Restarts this HS1CP even if there are units online to this HS1CP.

RETRY_ERRORS UNWRITEABLE_DATA

Tries to write the unwriteable data on a unit.

Format

RETRY_ERRORS UNWRITEABLE_DATA *unit-number*

Parameters

unit-number
Specifies the logical unit number (D0–D4094 or T0–T4094) which the write operation of the unwriteable data is attempted. The *unit-number* is the name given the unit when it was created using the ADD UNIT command.

Description

If a container fails in a way that customer data in the write-back cache cannot be written to the container, the unwriteable data error is reported. If possible the condition that is causing the unwriteable data should be corrected and the write operation should be attempted again. RETRY_ERRORS UNWRITEABLE_DATA attempts to write the unwriteable data error. No data is lost if the retry fails.

Examples

1.
```
HS1CP1> RETRY_ERRORS UNWRITEABLE_DATA D13
```

   Attempts to write the cached data on disk unit D13 that was previously marked unwriteable.

RUN

Runs a diagnostic or utility on THIS_CONTROLLER.

Format

RUN *program-name*

Parameters

program-name
The name of the diagnostic or utility to be run. DILX and TILX are examples of utilities and diagnostics that can be run from the CLI.

Description

The RUN command enables various diagnostics and utilities on THIS_CONTROLLER. Diagnostics and utilities can be run *only* on the HS1CP where the terminal or DUP connection is connected.

For specific information about available diagnostics and utilities, refer to

Examples

1.
```
HS1CP1> RUN DILX
Disk Inline Exerciser - version 2.0
            .
            .
            .
```
Runs the DILX diagnostic.

SELFTEST OTHER_CONTROLLER

Runs a self-test on the other HS1CP.

Format

SELFTEST OTHER_CONTROLLER

Description

The SELFTEST OTHER_CONTROLLER command flushes all user data from the other HS1CP's write-back cache (if present), shuts down the other HS1CP, then restarts it in DAEMON loop-on-self-test mode. The OCP reset (//) button must be pressed to take the other HS1CP out of loop-on-self-test mode.

If any disks are online to the other HS1CP, the HS1CP does not self-test unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the HS1CP does not self-test unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes the other HS1CP to self-test immediately without flushing any user data to the disks, even if drives are online to the host.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the HS1CP does not start the self-test unless IGNORE_ERRORS is specified.

If the IGNORE_ERRORS qualifier is specified, the HS1CP starts the self-test even if all customer data cannot be written to disk from the write-back cache.

CAUTION

Customer data may be lost or corrupted if the IGNORE_ERRORS qualifier is specified.

IMMEDIATE
NOIMMEDIATE (Default)
If IMMEDIATE is specified, the HS1CP will immediately start self-test without checking for online devices or flushing user data from write cache to disk.

CAUTION

Customer data may be lost or corrupted if the IMMEDIATE qualifier is specified.

OVERRIDE_ONLINE
NOOVERRIDE_ONLINE (Default)
If any units are online to the HS1CP, the HS1CP does not self-test unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the HS1CP starts the self-test after all customer data is written to disk from the write-back cache.

CAUTION

Customer data may be lost or corrupted if the OVERRIDE_ONLINE qualifier is specified.

Examples

1. `HS1CP1> SELFTEST OTHER_CONTROLLER`

   Starts the self-test on the other HS1CP, as long as the other HS1CP does not have any units online.

2. `HS1CP1> SELFTEST OTHER_CONTROLLER OVERRIDE_ONLINE`

   Starts the self-test on the other HS1CP even if there are units online to the other HS1CP.

SELFTEST THIS_CONTROLLER

Runs a self-test on this HS1CP.

Format

SELFTEST THIS_CONTROLLER

Description

The SELFTEST THIS_CONTROLLER command flushes all user data from this HS1CP's write-back cache (if present), shuts down this HS1CP, then restarts it in DAEMON loop-on-self-test mode. The OCP reset (//) button must be pressed to take this HS1CP out of loop-on-self-test mode.

If any disks are online to this HS1CP, the HS1CP does not self-test unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the HS1CP does not self-test unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes this HS1CP to self-test immediately without flushing any user data to the disks, even if drives are online to a host.

Note

If you enter a SELFTEST THIS_CONTROLLER command, and you are using a virtual terminal to communicate with the HS1CP, the connection is lost when this HS1CP starts the self-test.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the HS1CP does not start the self-test unless IGNORE_ERRORS is specified.

If the IGNORE_ERRORS qualifier is specified, the HS1CP starts the self-test even if all customer data cannot be written to disk from the write-back cache.

CAUTION

Customer data may be lost or corrupted if the IGNORE_ERRORS qualifier is specified.

IMMEDIATE
NOIMMEDIATE (Default)
If IMMEDIATE is specified, the HS1CP will immediately start self-test without checking for online devices or flushing user data from write cache to disk.

CAUTION

Customer data may be lost or corrupted if the IMMEDIATE qualifier is specified.

OVERRIDE_ONLINE
NOOVERRIDE_ONLINE (Default)
If any units are online to the HS1CP, the HS1CP does not self-test unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the HS1CP starts the self-test after all customer data is written to disk from the write-back cache.

CAUTION

Customer data may be lost or corrupted if the OVERRIDE_ONLINE qualifier is specified.

Examples

1. `HS1CP1> SELFTEST THIS_CONTROLLER`

   Starts the self-test on this HS1CP as long as this HS1CP does not have any units online.

2. `HS1CP1> SELFTEST THIS_CONTROLLER OVERRIDE_ONLINE`

   Starts the self-test on this HS1CP even if there are units on line to this HS1CP.

SET *disk-container-name*

Changes the characteristics of a disk drive.

Format

SET *disk-container-name*

Parameters

disk-container-name
The name of the disk drive that will have its characteristics changed.

Description

Changes the characteristics of a disk drive.

Qualifiers

TRANSPORTABLE
NOTRANSPORTABLE (Default)
In normal operations, the HS1CP makes a small portion of the disk inaccessible to the host and uses this area to store metadata, which improves data reliability, error detection, and recovery. This vast improvement comes at the expense of transportability.

If NOTRANSPORTABLE is specified and there is no valid metadata on the unit, the unit must be initialized.

Note

Digital recommends that you avoid specifying TRANSPORTABLE unless transportability of disk drive or media is imperative and there is no other way to accomplish moving the data.

Examples

1. `HS1CP1> SET DISK130 TRANSPORTABLE`

   Sets DISK130 to transportable.

SET FAILOVER

Places THIS_CONTROLLER and OTHER_CONTROLLER into a dual-redundant configuration.

Format

SET FAILOVER *COPY=configuration-source*

Parameters

COPY=configuration-source
Specifies where the "good" copy of the device configuration resides.

If THIS_CONTROLLER is specified for *configuration-source*, all the device configuration information on THIS_CONTROLLER (the HS1CP that either the maintenance terminal is connected to or the virtual terminal is connected to) is copied to the other HS1CP.

If OTHER_CONTROLLER is specified for *configuration-source*, all the device configuration information on the OTHER_CONTROLLER (the HS1CP that either the maintenance terminal or the virtual terminal connection is *not* connected to) is copied to this HS1CP.

Description

The SET FAILOVER command places THIS_CONTROLLER and the OTHER_CONTROLLER in a dual-redundant configuration. After entering this command, if one of the two HS1CPs fail, the devices and cache (if any) attached to the failed HS1CP become available to and accessible through the operating HS1CP.

CAUTION

All device configuration information on the HS1CP *not* specified by the COPY= parameter is destroyed and overwritten by the configuration information found in the HS1CP specified by the COPY= parameter. **Make sure you know where your good configuration information is stored, or you have a complete copy of the device configuration, BEFORE entering this command**.

A considerable amount of work and effort is lost by overwriting a good configuration with incorrect information if the wrong HS1CP is specified by the COPY= parameter.

Also note that due to the amount of information that must be passed between the two HS1CPs, this command may take up to 1 minute to complete.

Examples

1. `HS1CP1> SET FAILOVER COPY=THIS_CONTROLLER`

   Places two HS1CPs into a dual-redundant configuration, where the "good" data was on the HS1CP that the maintenance terminal or virtual terminal connection was connected to.

2. `HS1CP1> SET FAILOVER COPY=OTHER_CONTROLLER`

   Places two HS1CPs into a dual-redundant configuration, where the "good" data was on the HS1CP that the maintenance terminal or virtual terminal connection was *not* connected to.

SET NOFAILOVER

Removes THIS_CONTROLLER and OTHER_CONTROLLER (if reachable) from a dual-redundant configuration.

Format

SET NOFAILOVER

Description

The SET NOFAILOVER command removes THIS_CONTROLLER and the OTHER_CONTROLLER (if currently reachable) from a dual-redundant configuration. Before or immediately after entering this command, one HS1CP should be physically removed because the sharing of devices is not supported by single HS1CP configurations.

The HS1CP on which the command was entered is always removed from a dual-redundant state, even if the other HS1CP is not currently reachable. No configuration information is lost when leaving a dual-redundant state.

Examples

1. `HS1CP1> SET NOFAILOVER`

   Removes the two HS1CPs from a dual-redundant configuration.

SET OTHER_CONTROLLER

Changes the other HS1CP's parameters (in a dual-redundant configuration the HS1CP that the maintenance terminal is *not* connected to or the HS1CP that is *not* the target of the DUP connection.

Format

SET OTHER_CONTROLLER

Description

The SET OTHER_CONTROLLER command allows you to modify the HS1CP parameters of the other HS1CP in a dual-redundant configuration.

Qualifiers

CACHE_FLUSH_TIMER=*n*
CACHE_FLUSH_TIMER=DEFAULT
Specifies how many seconds (1–65535) of idle time may elapse before the write-back cache flushes its entire contents to disk. After the specified time, the write-back cache flushes its contents to disk to ensure data integrity.

ID=*n*
Specifies the DSSI node number (0–7).

MSCP_ALLOCATION_CLASS=*n*
Specifies the allocation class (0–255 in a single HS1CP configuration or 1–255 in a dual-redundant configuration).

When first installed, the HS1CP's MSCP_ALLOCATION_CLASS is set to 0.

PATH
NOPATH
Enables or disables the DSSI port.

When first installed, NOPATH is set.

PROMPT="*new prompt*"
Specifies a 1- to 16-character prompt enclosed in quotes that will be displayed when the HS1CP's CLI prompts for input. Only printable ASCII characters are valid.

When first installed, the CLI prompt is set to HS1CP1>.

SCS_NODENAME="*xxxxxx*"
Specifies a one- to six-character name for node.

TERMINAL_PARITY=ODD
TERMINAL_PARITY=EVEN
NOTERMINAL_PARITY
Specifies the parity transmitted and expected. Parity options are ODD or EVEN. NOTERMINAL_PARITY causes the HS1CP to not check for, or transmit any parity on the terminal lines.

When first installed, the HS1CP's terminal parity is set to NOTERMINAL_PARITY.

TERMINAL_SPEED=*baud_rate*
Sets the terminal speed to 300, 600, 1200, 2400, 4800, 9600 or 19200 baud. The transmit speed is always equal to the receive speed.

When first installed, the HS1CP's terminal speed is set to 9600 baud.

TMSCP_ALLOCATION_CLASS=*n*
Specifies the allocation class (0–255 in a single HS1CP configuration or 1–255 in a dual-redundant configuration).

When first installed, the HS1CP's TMSCP_ALLOCATION_CLASS is set to 0.

Examples

1.
```
HS1CP1> SET OTHER_CONTROLLER PATH_A PATH_B SPEED=1200
```

   Modifies the other HS1CP's two CI paths and sets the terminal speed to 1200 baud.

SET *RAIDset-container-name*

Changes the characteristics of a RAIDset.

Format

SET *RAIDset-container-name*

Parameters

RAIDset-container-name
The name of the RAIDset that will have its characteristics modified.

Description

Changes the characteristics of a RAIDset.

Qualifiers

POLICY=BEST_FIT
POLICY=BEST_PERFORMANCE (Default)
NOPOLICY
Specifies the replacement policy to use when a member within the RAIDset fails.

BEST_FIT gives highest priority to finding a replacement device within the spareset that most closely matches the sizes of the remaining members of the RAIDset. After finding the most closely matching devices, the device that gives the best performance is selected.

BEST_PERFORMANCE (default) gives highest priority to finding a replacement device within the spareset that results in the best performance of the RAIDset. After finding the best performing devices, the device that most closely matches the size of the remaining members of the RAIDset is selected.

NOPOLICY retires a failing device from the RAIDset without selecting a replacement. This causes the RAIDset to run in a reduced state until a BEST_FIT or BEST_PERFORMANCE policy is selected, or a member is manually replaced in the RAIDset (see SET *raidset-container-name*).

RECONSTRUCT=NORMAL (Default)
RECONSTRUCT=FAST
NORECONSTRUCT
Specifies the speed at which a RAIDset will be reconstructed when a new member is added to the RAIDset or immediately after the RAIDset is initialized.

RECONSTRUCT=NORMAL (default) balances overall performance of the HS1CP against the demand of reconstructing the RAIDset.

RECONSTRUCT=FAST reconstructs the RAIDset at the fastest rate possible resulting in some loss of performance of the HS1CP overall.

REMOVE=*disk-container-name*
Specifies the removal of a disk member from a RAIDset. If the RAIDset is already in a reduced state, an error is displayed and the command is rejected. If a replacement policy is specified, the replacement is taken from the spareset to replace the removed member using the specified policy. If NOPOLICY is

specified, the RAIDset continues to operate in a reduced state until a replacement is manually specified (see SET *RAIDset-container-name*) REPLACE=) or a policy is specified (see SET *RAIDset-container-name* POLICY=).

The disk removed via the REMOVE= command is added to the failedset.

Note

No other qualifiers to the SET *RAIDset-container-name* command may be specified if REMOVE is specified.

REPLACE=*disk-container-name*

Specifies the replacement of a disk member into a reduced RAIDset. If the RAIDset is not in a reduced state, an error is displayed and the command is rejected. If a replacement policy is already specified, an error is displayed and the command is rejected. If the disk specified is already being used by a configuration (including a spareset), an error is displayed and the command is rejected. Otherwise, the disk specified is added as a member to the specified RAIDset and a reconstruct operation begins immediately.

Note

No other qualifiers to the SET *RAIDset-container-name* command may be specified if REPLACE is specified.

Examples

1.
```
HS1CP1> SET RAID9 POLICY=BEST_FIT
```

   Changes RAIDset RAID9's policy to BEST_FIT.

2.
```
HS1CP1> SET RAID9 REMOVE=DISK0
```

   Removes RAIDset RAID9's member DISK0 from the RAIDset. If there is a replacement policy, a new disk is taken from the spareset and placed in the RAIDset automatically.

3.
```
HS1CP1> SET RAID9 REPLACE=SPAREDISK
```

   Adds disk SPAREDISK to the reduced RAIDset, RAID9. A reconstruct operation begins immediately on SPAREDISK.

SET THIS_CONTROLLER

Changes this HS1CP's parameters (the HS1CP that the maintenance terminal is connected to or the target of the DUP connection).

Format

SET THIS_CONTROLLER

Description

The SET THIS_CONTROLLER command allows you to modify HS1CP parameters on THIS_CONTROLLER in single and dual-redundant configurations.

Qualifiers

CACHE_FLUSH_TIMER=*n*
CACHE_FLUSH_TIMER=DEFAULT
Specifies how many seconds (1–65535) of idle time may elapse before the write-back cache flushes its entire contents to disk. After the specified time, the write-back cache flushes its contents to disk to ensure data integrity.

ID=*n*
Specifies the DSSI node number (0–7).

MSCP_ALLOCATION_CLASS=*n*
Specifies the allocation class (0–255 in a single HS1CP configuration or 1–255 in a dual-redundant configuration).

When first installed, the HS1CP's MSCP_ALLOCATION_CLASS is set to 0.

PATH
NOPATH
Enables or disables the DSSI port.

When first installed, NOPATH is set.

PROMPT="*new prompt*"
Specifies a 1- to 16-character prompt enclosed in quotes that will be displayed when the HS1CP's CLI prompts for input. Only printable ASCII characters are valid.

When first installed, the CLI prompt is set to HS1CP1>.

SCS_NODENAME="*xxxxxx*"
Specifies a one- to six-character name for node.

TERMINAL_PARITY=ODD
TERMINAL_PARITY=EVEN
NOTERMINAL_PARITY
Specifies the parity transmitted and expected. Parity options are ODD or EVEN. NOTERMINAL_PARITY causes the HS1CP to not check for, or transmit any parity on the terminal lines.

When first installed, the HS1CP's terminal parity is set to NOTERMINAL_PARITY.

TERMINAL_SPEED=*baud_rate*
Sets the terminal speed to 300, 600, 1200, 2400, 4800, 9600 or 19200 baud. The transmit speed is always equal to the receive speed.

When first installed, the HS1CP's terminal speed is set to 9600 baud.

TMSCP_ALLOCATION_CLASS=*n*
Specifies the allocation class (0–255 in a single HS1CP configuration or 1–255 in a dual-redundant configuration).

When first installed, the HS1CP's TMSCP_ALLOCATION_CLASS is set to 0.

Examples

1. `HS1CP1> SET THIS_CONTROLLER PATH_A PATH_B SPEED=1200`

   Modifies this HS1CP's two CI paths and sets the terminal speed to 1200 baud.

2. `HS1CP1> SET THIS_CONTROLLER ID=5`

   Sets this HS1CP so it responds to requests for target 5.

3. `HS1CP1> SET THIS_CONTROLLER ID=(2,5)`

   Sets this HS1CP so it responds to requests for targets 2 and 5.

SET *unit-number*

Changes the unit parameters.

Format

SET *unit-number*

Parameters

unit-number
Specifies the logical unit number (D0–D4094 or T0–T4094) to modify the software switches. The *unit-number* is the name given the unit when it was created using the ADD UNIT command.

Description

The SET command is used to change logical unit parameters.

Qualifiers for a Unit Created from a CDROM Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

Qualifiers for a Unit Created from a TRANSPORTABLE Disk Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

Qualifiers for a Unit Created from a NOTRANSPORTABLE Disk Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the HS1CP's write-back cache on this unit.

Note

It may take up to 5 minutes to flush unwritten data from the write-back cache once you disable write-back caching.

Qualifiers for a Unit Created from a TRANSPORTABLE Optical Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

Qualifiers for a Unit Created from a NOTRANSPORTABLE Optical Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the HS1CP's write-back cache on this unit.

Note

It may take up to 5 minutes to flush unwritten data from the write-back cache once you disable write-back caching.

Qualifiers for a Unit Created from a RAIDset

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

Note

Writes may still be performed to a write-protected RAIDset to satisfy a reconstruct pass or to reconstruct a newly replaced member. However, write protect will disable the writing of any new data.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the HS1CP's write-back cache on this unit.

Note

It may take up to 5 minutes to flush unwritten data from the write-back cache once you disable write-back caching.

Qualifiers for a Unit Created from a Stripeset

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the HS1CP. Any transfers over this size are not cached. Valid values are 1–1024.

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

READ_CACHE (Default)
NOREAD_CACHE
Enables and disables the HS1CP's read cache on this unit.

RUN (Default)
NORUN
Enables and disables a unit's availibility to the host. When RUN (default) is specified, the devices that make up the unit will be spun up and the unit will be made available to the host. If NORUN is specified, the devices that make up the unit will still be spun up, but the unit will not be made available to the host.

WRITE_PROTECT
NOWRITE_PROTECT (Default)
Enables and disables write protection of the unit.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the HS1CP's write-back cache on this unit.

Note

It may take up to 5 minutes to flush unwritten data from the write-back cache once you disable write-back caching.

Qualifiers for a Unit Created from a Tape Drive

DEFAULT_FORMAT=*format*
DEFAULT_FORMAT=DEVICE_DEFAULT (Default)
Specifies the tape format to be used unless overridden by the host. Note that not all devices support all formats. The easiest way to determine what formats are supported by a specific device is to enter the "SHOW <tape unit number> DEFAULT_FORMAT= ?" command—the valid options will be displayed.

Supported tape formats are as follows:

- DEVICE_DEFAULT (default)

  The default tape format is the default that the device uses, or, in the case of devices that can be set via switches on the front panel, the settings of those switches.
- 800BPI_9TRACK
- 1600BPI_9TRACK
- 6250BPI_9TRACK
- TZ85
- TZ86
- TZ87_NOCOMPRESSION
- TZ87_COMPRESSION
- DAT_NOCOMPRESSION
- DAT_COMPRESSION
- 3480_NOCOMPRESSION
- 3480_COMPRESSION

PREFERRED_PATH=THIS_CONTROLLER
PREFERRED_PATH=OTHER_CONTROLLER
NOPREFERRED_PATH (Default)

Specifies the preferred HS1CP that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either HS1CP (NOPREFERRED_PATH).

The preferred path qualifier is used only if both HS1CPs are running in a dual-redundant configuration. If one HS1CP fails, then all the devices will be made accessible through the remaining HS1CP, ignoring the preferred path setting.

When the failed HS1CP is restarted, the drives automatically return to the HS1CP specified by the preferred path qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single HS1CP, however, the qualifier will not take effect until a second HS1CP is added and the two HS1CPs are configured for dual-redundancy. The second HS1CP will inherit any PREFERRED_PATH settings, and the two HS1CPs will operate using the preset PREFERRED_PATH options.

Examples

1. `HS1CP1> SET D1 WRITE_PROTECT NOREAD_CACHE`

   Sets the write protect and turns off the read cache on unit D1.

2. `HS1CP1> SET T47 DEFAULT_FORMAT=1600BPI_9TRACK`

   Sets unit T47 to 1600 bpi.

SHOW CDROMS

Shows all CDROM drives and drive information.

Format

SHOW CDROMS

Description

The SHOW CDROMS command displays all the CDROM drives known to the HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
HS1CP1> SHO CDROM
Name          Type          Port Targ  Lun                    Used by
------------------------------------------------------------------------------

CDROM230      cdrom            2    3    0                    D623
CDROM240      cdrom            2    4    0                    D624
```

Shows a basic listing of CDROMs.

2.
```
HS1CP1> SHO CDROM FULL
Name          Type          Port Targ  Lun                    Used by
------------------------------------------------------------------------------

CDROM230      cdrom            2    3    0                    D623
          DEC      RRD44   (C) DEC  3593
CDROM240      cdrom            2    4    0                    D624
          DEC      RRD44   (C) DEC  3593
```

Shows a full listing of CDROMs.

SHOW *cdrom-container-name*

Shows information about a CDROM.

Format

SHOW *cdrom-container-name*

Parameters

cdrom-container-name
The name of the CDROM drive to be displayed.

Description

The SHOW *cdrom-container-name* command is used to show specific information about a particular CDROM drive.

Examples

1.
```
HS1CP1> SHO CDROM230
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

CDROM230      cdrom            2    3    0                    D623
          DEC      RRD44   (C) DEC  3593
```

A listing of CDROM CDROM230.

SHOW DEVICES

Shows physical devices and physical device information.

Format

SHOW DEVICES

Description

The SHOW DEVICES command displays all the devices known to the HS1CP. First disks are shown, then tapes, then CDROMs and finally opticals.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Information contained in the additional information is dependent on the device type.

Examples

1.
```
HS1CP1> SHOW DEVICES
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

DI0           disk             1    0    0                    D100
DI1           disk             1    1    0                    D110
TAPE110       tape             3    1    0                    T110
TAPE130       tape             3    3    0                    T130
CDROM230      cdrom            2    3    0                    D623
CDROM240      cdrom            2    4    0                    D624
```

Shows a basic listing of devices attached to the HS1CP.

2.
```
HS1CP1> SHOW DEVICES FULL
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

DI0           disk             1    0    0                    D100
          DEC      RZ35     (C) DEC X388
DI1           disk             1    1    0                    D110
          DEC      RZ26     (C) DEC T386
TAPE110       tape             3    1    0                    T110
          DEC      TZ877    (C) DEC 930A
TAPE130       tape             3    3    0                    T130
          DEC      TZ877    (C) DEC 930A
CDROM230      cdrom            2    3    0                    D623
          DEC      RRD44   (C) DEC  3593
CDROM240      cdrom            2    4    0                    D624
          DEC      RRD44   (C) DEC  3593
```

Shows a full listing of devices attached to the HS1CP.

SHOW DISKS

Shows all disk drives and drive information.

Format

SHOW DISKS

Description

The SHOW DISKS command displays all the disk drives known to the HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
HS1CP1> SHOW DISKS
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

DI0           disk             1    0    0                    D100
DI1           disk             1    1    0                    D110
```

Shows a basic listing of disks attached to the HS1CP.

2.
```
HS1CP1> SHOW DISKS FULL
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

DI0           disk             1    0    0                    D100
          DEC      RZ35     (C) DEC X388
DI1           disk             1    1    0                    D110
          DEC      RZ26     (C) DEC T386
```

Shows a full listing of disks attached to the HS1CP.

SHOW *disk-container-name*

Shows information about a disk drive.

Format

SHOW *disk-container-name*

Parameters

disk-container-name
The name of the disk drive to be displayed.

Description

The SHOW *disk-container-name* command is used to show specific information about a particular disk.

Examples

1.
```
HS1CP1> SHOW DI3
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

DI3           disk             1    3    0                    D130
          DEC      RZ26     (C) DEC X388
```

Shows a listing of disk DI3.

SHOW FAILEDSET

Shows the members of the failedset.

Format

SHOW FAILEDSET

Description

The SHOW FAILEDSET command displays all the disk drives that are members of the failedset.

Examples

1.
```
HS1CP1> SHOW FAILEDSET
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------------

FAILEDSET     failedset                     DISK310
                                            DISK410
```

Shows a listing of the members of the failedset.

SHOW LOADERS

Shows all loaders and loader information.

Format

SHOW LOADERS

Description

The SHOW LOADERS command displays all the loaders known to the HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
HS1CP1> SHOW LOADERS
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------
LDR511        loader           5    1    1                    CMD511
```

Basic listing of loaders.

2.
```
HS1CP1> SHOW LOADERS FULL
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

LDR511        loader           5    1    1                    CMD511
          DEC      TZ Media Changer 930A
```

Full listing of loaders.

SHOW *loader-container-name*

Shows information about a loader.

Format

SHOW *loader-container-name*

Parameters

loader-container-name
The name of the loader to be displayed.

Description

The SHOW *loader-container-name* command is used to show specific information about a particular loader.

Examples

HS1CP1> SHOW LDR511

1.
```
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

LDR511        loader           5    1    1                   CMD511
         DEC      TZ Media Changer 930A
```

Listing of loader LDR511

SHOW OPTICALS

Shows all optical drives and drive information.

Format

SHOW OPTICALS

Description

The SHOW OPTICALS command displays all the optical drives known to the HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
HS1CP1> SHOW OPTICALS
Name          Type          Port Targ  Lun                    Used by
------------------------------------------------------------------------------

OP0           optical          1    0    0                    D100
OP1           optical          1    1    0                    D110
```

Shows a basic listing of optical drives attached to the HS1CP.

2.
```
HS1CP1> SHOW OPTICALS FULL
Name          Type          Port Targ  Lun                    Used by
------------------------------------------------------------------------------

OP0           optical          1    0    0                    D100
          DEC      RWZ52     (C)DEC 3404
OP1           optical          1    1    0                    D110
          DEC      RWZ52     (C)DEC 3404
```

Shows a full listing of optical drives attached to the HS1CP.

SHOW *optical-container-name*

Shows information about an optical drive.

Format

SHOW *optical-container-name*

Parameters

optical-container-name
The name of the optical drive to be displayed.

Description

The SHOW *optical-container-name* command is used to show specific information about a particular optical drive.

Examples

1.
```
HS1CP1> SHOW OP3
Name          Type          Port Targ  Lun                    Used by
------------------------------------------------------------------------------

OP3           optical         1    1    0                    D110
          DEC      RWZ52     (C)DEC 3404
```

Shows a listing of optical drive DI3.

SHOW OTHER_CONTROLLER

Shows information for the other HS1CP.

Format

SHOW OTHER_CONTROLLER

Description

Shows all HS1CP, port, and terminal information for the other HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information is displayed after the basic HS1CP information.

Example

```
HS1CP1> SHOW OTHER_CONTROLLER
Controller:
        HS1CP    (C) DEC ZG33400022 Firmware E35D-0, Hardware AX01
        Configured for dual-redundancy with ZG33400026
            In dual-redundant configuration
        SCSI address 6
        Time: 18-AUG-1994 18:23:27
Host port:
        Node name: HS1CP1, valid DSSI node 0
        Host path is ON
        MSCP allocation class    9
        TMSCP allocation class   9
Cache:
        32 megabyte write cache, version 2
        Cache is GOOD
        Battery is GOOD
        No unflushed data in cache
        CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
```

Shows the basic HS1CP information.

SHOW PASSTHROUGH

Shows passthrough containers and container information.

Format

SHOW PASSTHROUGH

Description

The SHOW PASSTHROUGH command displays all the passthrough containers known by the HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each storageset.

Examples

1.
```
HS1CP1> SHOW PASSTHROUGH
MASS> sho pass
Name          Storageset                    Uses            Used by
-------------------------------------------------------------------------------

CMD100        passthrough                   DISK100         D610
CMD240        passthrough                   DISK240         D624
CMD310        passthrough                   CD310           D631
CMD320        passthrough                   CD320           D632
CMD640        passthrough                   LDR400          D640
```

Shows a listing of all passthrough containers.

SHOW *passthrough-container-name*

Shows information about a passthrough container.

Format

SHOW *passthrough-container-name*

Parameters

passthrough-container-name
The name of the passthrough container to be displayed.

Description

The SHOW *passthrough-container-name* command is used to show specific information about a passthrough container.

Examples

1.
```
HS1CP1> SHOW CMD100
MASS> sho pass
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

CMD100        passthrough                   DISK100          D610
```

Shows a listing of passthrough container CMD100.

SHOW RAIDSETS

Shows RAIDsets and RAIDset information.

Format

SHOW RAIDSETS

Description

The SHOW RAIDSETS command displays all the RAIDsets known by the HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each storageset.

Examples

1.
```
HS1CP1> SHOW RAIDSETS
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

R0            raidset                       DISK110          D401
                                            DISK220
                                            DISK310
                                            DISK400

R1            raidset                       DISK130
                                            DISK240
                                            DISK330
                                            DISK420
```

Shows a basic listing of all RAIDsets.

2.
```
HS1CP1> SHOW RAIDSETS FULL
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

R0            raidset                       DISK110          D401
                                            DISK220
                                            DISK310
                                            DISK400
        Switches:
          CHUNKSIZE = 63 blocks
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
        State:
          RECONSTRUCT 3% complete
```

```
R1            raidset                          DISK130
                                               DISK240
                                               DISK330
                                               DISK420
        Switches:
          CHUNKSIZE = 63 blocks
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
        State:
          RECONSTRUCT 0% complete
```

Shows a full listing of all RAIDsets.

SHOW *raidset-container-name*

Shows information about a RAIDset.

Format

SHOW *raidset-container-name*

Parameters

raidset-container-name
The name of the RAIDset to be displayed.

Description

The SHOW *raidset-container-name* command is used to show specific information about a particular RAIDset.

Examples

1.
```
HS1CP1> SHOW RAID9
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

RAID9         raidset                       DISK130
                                            DISK240
                                            DISK330
                                            DISK420
        Switches:
          CHUNKSIZE = 63 blocks
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
        State:
          RECONSTRUCT 0% complete
```

Shows a listing of RAIDset RAID9.

SHOW SPARESET

Shows the members of the spareset.

Format

SHOW SPARESET

Description

The SHOW SPARESET command displays all the disk drives that are members of the spareset.

Examples

1.
```
HS1CP1> SHOW SPARESET
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------------

SPARESET      spareset                      DISK150
                                            DISK350
                                            DISK440
```

Shows a list of the members of the spareset.

SHOW STORAGESETS

Shows storagesets and storageset information.

Format

SHOW STORAGESETS

Description

The SHOW STORAGESETS command displays all the storagesets known by the HS1CP. A storageset is any collection of containers, such as stripesets, RAIDsets, the spareset and the failedset.

Stripesets are displayed first, followed by RAIDsets, sparesets, failedsets, and then passthrough containers.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each storageset.

Examples

1.
```
HS1CP1> SHOW STORAGESETS
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------------

S0            stripeset                     DISK500          D1
                                            DISK510
                                            DISK520
```

Shows a basic listing of all storagesets.

2.
```
HS1CP1> SHOW STORAGESETS FULL
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------------

S0            stripeset                     DISK530
                                            DISK550
                                            DISK600
        Switches:
          CHUNKSIZE = 24 blocks
S1            stripeset                     DISK620
                                            DISK640
        Switches:
          CHUNKSIZE = 24 blocks
```

```
R0            raidset                       DISK110          D401
                                            DISK220
                                            DISK310
                                            DISK400
        Switches:
          CHUNKSIZE = 63 blocks
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
        State:
          RECONSTRUCT 3% complete

R1            raidset                       DISK130
                                            DISK240
                                            DISK330
                                            DISK420
        Switches:
          CHUNKSIZE = 63 blocks
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
        State:
          RECONSTRUCT 0% complete

SPARESET      spareset                      DISK150
                                            DISK350
                                            DISK440

FAILEDSET     failedset

CMD100        passthrough                   DISK100          D610

CMD240        passthrough                   DISK250          D624

CMD310        passthrough                   CD310            D631
```

Shows a full listing of all storagesets.

SHOW STRIPESETS

Shows stripesets and related stripeset information.

Format

SHOW STRIPESETS

Description

The SHOW STRIPESET command displays all the stripesets known by the HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each storageset.

Examples

1.
```
HS1CP1> SHOW STRIPESETS
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

S0            stripeset                     DISK500          D1
                                            DISK510
                                            DISK520

S1            stripeset                     DISK400          D17
                                            DISK410
                                            DISK420
```

Shows a basic listing of all stripesets.

2.
```
HS1CP1> SHOW STRIPESETS FULL
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

S0            stripeset                     DISK530
                                            DISK550
                                            DISK600
        Switches:
          CHUNKSIZE = 24 blocks

S1            stripeset                     DISK620
                                            DISK640
        Switches:
          CHUNKSIZE = 24 blocks
```

Shows a full listing of all stripesets.

SHOW *stripeset-container-name*

Shows information about a specific stripeset.

Format

SHOW *stripeset-container-name*

Parameters

stripeset-container-name
The name of the stripeset to be displayed.

Description

The SHOW *stripeset-container-name* command is used to show specific information about a particular stripeset.

Examples

1.
```
HS1CP1> SHOW STRIPE0
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

STRIPE0       stripeset                     DISK530
                                            DISK550
                                            DISK600

        Switches:
          CHUNKSIZE = 24 blocks
```

Shows a listing of stripeset STRIPE0.

SHOW TAPES

Shows all tape drives and tape drive information.

Format

SHOW TAPES

Description

The SHOW TAPES command displays all the tape drives known to the HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
HS1CP1> SHOW TAPES
Name          Type          Port Targ  Lun                    Used by
------------------------------------------------------------------------------

TAPE200       tape             2    0    0
```

Shows a basic listing of tape drives.

2.
```
HS1CP1> SHOW TAPES FULL
Name          Type          Port Targ  Lun                    Used by
------------------------------------------------------------------------------

TAPE200       tape             2    0    0
          DEC      TLZ06     (C)DEC 0491
```

Shows a full listing of tape drives.

SHOW *tape-container-name*

Shows information about a specific tape drive.

Format

SHOW *tape-container-name*

Parameters

tape-container-name
The name of the tape drive to be displayed.

Description

The SHOW *tape-container-name* command is used to show specific information about a particular tape drive.

Examples

1.
```
HS1CP1> SHOW TAPE200
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

TAPE200       tape             2    0    0
          DEC      TLZ06     (C)DEC 0491
```

Shows a listing of TAPE200.

SHOW THIS_CONTROLLER

Shows information for this HS1CP.

Format

SHOW THIS_CONTROLLER

Description

Shows all HS1CP, port, and terminal information for this HS1CP.

Qualifiers

FULL
If the FULL qualifier is specified, additional information is displayed after the basic HS1CP information.

Example

```
HS1CP1> SHOW THIS_CONTROLLER
Controller:
        HS1CP    (C) DEC ZG33400026 Firmware E35D-0, Hardware AX01
        Configured for dual-redundancy with ZG33400022
            In dual-redundant configuration
        SCSI address 7
        Time: 18-AUG-1994 18:23:55
Host port:
        Node name: HS1CP1, valid DSSI node 1
        Host path is ON
        MSCP allocation class    9
        TMSCP allocation class   9
Cache:
        32 megabyte write cache, version 2
        Cache is GOOD
        Battery is GOOD
        No unflushed data in cache
        CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
```

SHOW UNITS

Shows all units and unit information.

Format

SHOW UNITS

Description

The SHOW UNITS command displays all the units known by the HS1CP. First disks (including CDROMs) are listed, then tapes.

Qualifiers

FULL
If the FULL qualifier is specified after UNITS, additional information may be displayed after each unit-number, such as the switch settings.

Examples

1.
```
HS1CP1> SHOW UNITS
MSCP unit                                       Uses
-------------------------------------------------------------

  D401                                          R0
  T41                                           TAPE200
```

Shows a basic listing of units available on the HS1CP.

2.
```
HS1CP1> SHOW UNITS FULL
MSCP unit                                       Uses
-------------------------------------------------------------

  D401                                          R0
        Switches:
          RUN                    NOWRITE_PROTECT        READ_CACHE
          NOWRITEBACK_CACHE
          MAXIMUM_CACHED_TRANSFER_SIZE = 32
        State:
          AVAILABLE
          No exclusive access
          NOPREFERRED_PATH
  T41                                           TAPE200
        Switches:
          DEFAULT_FORMAT = DEVICE_DEFAULT
        State:
          AVAILABLE
          No exclusive access
          Current Format = DAT_DDS_NOCOMPRESSION
          NOPREFERRED_PATH
```

Shows a full listing of units available on the HS1CP.

SHOW *unit-number*

Shows information about a specific unit.

Format

SHOW *unit-number*

Parameters

unit-number
The unit number of the unit that is to be displayed.

Description

The SHOW *unit-number* command is used to show specific information about a particular unit.

Examples

1.
```
HS1CP1> SHOW D150
MSCP unit                                       Uses
-------------------------------------------------------------

  D150                                          R0
        Switches:
          RUN                    NOWRITE_PROTECT        READ_CACHE
          NOWRITEBACK_CACHE
          MAXIMUM_CACHED_TRANSFER_SIZE = 32
        State:
          AVAILABLE
          No exclusive access
          NOPREFERRED_PATH
```

Shows a listing of a specific disk unit.

2.
```
HS1CP1> sho T110
MSCP unit                                       Uses
-------------------------------------------------------------

  T110                                          TAPE200
        Switches:
          DEFAULT_FORMAT = DEVICE_DEFAULT
        State:
          AVAILABLE
          No exclusive access
          Current Format = DAT_DDS_NOCOMPRESSION
          NOPREFERRED_PATH
```

Shows a listing of a specific tape unit.

SHUTDOWN OTHER_CONTROLLER

Shuts down and does not restart the other HS1CP.

Format

SHUTDOWN OTHER_CONTROLLER

Description

The SHUTDOWN OTHER_CONTROLLER command flushes all user data from the other HS1CP's write-back cache (if present), then shuts down the other HS1CP.

If any disks are online to the other HS1CP, the HS1CP does not shut down unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the HS1CP does not shut down unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes the other HS1CP to shut down immediately without flushing any user data to the disks, even if drives are online to the host.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the HS1CP is not shut down unless IGNORE_ERROR is specified.

If the IGNORE_ERRORS qualifier is specified, the HS1CP shuts down even if all customer data cannot be written to disk from the write-back cache.

CAUTION

Customer data may be lost or corrupted if the IGNORE_ERRORS qualifier is specified.

IMMEDIATE
NOIMMEDIATE (Default)
If IMMEDIATE is specified, the HS1CP shuts down immediately without checking for online devices or flushing user data from write-back cache to disk.

CAUTION

Customer data may be lost or corrupted if the IMMEDIATE qualifier is specified.

OVERRIDE_ONLINE
NOOVERRIDE_ONLINE (Default)
If any units are online to the HS1CP, the HS1CP is not shutdown unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the HS1CP shuts down after all customer data is written to disk.

CAUTION

Customer data may be lost or corrupted if the OVERRIDE_ONLINE qualifier is specified.

Examples

1. `HS1CP1> SHUTDOWN OTHER_CONTROLLER`

   Shuts down the other HS1CP as long as the other HS1CP does not have any units online.

2. `HS1CP1> SHUTDOWN OTHER_CONTROLLER OVERRIDE_ONLINE`

   Shuts down the other HS1CP even if there are units online to the other HS1CP.

SHUTDOWN THIS_CONTROLLER

Shuts down and does not restart this HS1CP.

Format

SHUTDOWN THIS_CONTROLLER

Description

The SHUTDOWN THIS_CONTROLLER command flushes all user data from this HS1CP's write-back cache (if present), then shuts down this HS1CP.

If any disks are online to this HS1CP, the HS1CP does not shut down unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the HS1CP does not shut down unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes this HS1CP to shut down immediately without flushing any user data to the disks, even if drives are online to a host.

Note

If you issue a SHUTDOWN THIS_CONTROLLER command, communication with the HS1CP is lost when this HS1CP shuts down.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the HS1CP is not shut down unless IGNORE_ERROR is specified.

If the IGNORE_ERRORS qualifier is specified, the HS1CP shuts down even if all customer data cannot be written to disk from the write-back cache.

CAUTION

Customer data may be lost or corrupted if the IGNORE_ERRORS qualifier is specified.

IMMEDIATE
NOIMMEDIATE (Default)
If IMMEDIATE is specified, the HS1CP shuts down immediately without checking for online devices or flushing user data from write-back cache to disk.

CAUTION

Customer data may be lost or corrupted if the IMMEDIATE qualifier is specified.

OVERRIDE_ONLINE
NOOVERRIDE_ONLINE (Default)
If any units are online to the HS1CP, the HS1CP is not shutdown unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the HS1CP shuts down after all customer data is written to disk.

CAUTION

Customer data may be lost or corrupted if the OVERRIDE_ONLINE qualifier is specified.

Examples

1. `HS1CP1> SHUTDOWN THIS_CONTROLLER`

   Shuts down this HS1CP as long as this HS1CP does not have any units online.

2. `HS1CP1> SHUTDOWN THIS_CONTROLLER OVERRIDE_ONLINE`

   Shuts down this HS1CP even if there are units online to this HS1CP.

B.3 CLI Messages

The following sections describe messages you can encounter during interactive use of the CLI.

B.3.1 Error Conventions

An Error *nnnn:* message means that the command did not complete. Except for a few of the failover messages (6000 series), no part of the command was executed. When encountering an error entering or exiting dual-redundant mode, some synchronization problems are unavoidable; the error message in such a case tells you what to do to get things back in synchronization.

Multiple error messages may result from one command.

Items in angle brackets (<>) are replaced at run time with names, numbers, and so on.

B.3.2 CLI Error Messages

Error 1000: Unit number must be from 0 to 4094

Explanation: This error results from an ADD UNIT command when the *n* in the D*n* or T*n* specified is out of range. The MSCP or TMSCP unit number after the "D" or "T" must be in the range of 0 to 4094.

Retry the ADD UNIT command with a correct number.

Error 1010: Maximum cached transfer size must be 1 through 1024 blocks

Explanation: This error results from a SET <unit number> or an ADD UNIT command when MAXIMUM_CACHED_TRANSFER_SIZE was specified. MAXIMUM_CACHED_TRANSFER_SIZE must be in the range 1 through 1024. Retry the SET or ADD command with a correct number.

Error 1020: CHUNKSIZE must be from <minimum> to <maximum>

Explanation: This error results from an INITIALIZE *storageset-container-name* command when CHUNKSIZE was specified. The chunk size must be DEFAULT or greater than 15. Retry the INITIALIZE command with DEFAULT or a correct number.

Error 1090: Tape unit numbers must start with the letter 'T'

Explanation: All tape unit numbers are of the form "T*n*." This error is displayed if you add a tape unit that does not begin the unit number with the letter "T."

Retry the ADD command with the letter "T" at the start of the unit number.

Error 1100: Disk unit numbers must start with the letter 'D'

Explanation: All disk unit numbers are of the form "D*n*." This error is displayed if you add a disk unit that does not begin the unit number with the letter "D."

Retry the ADD command with the letter "D" at the beginning of the unit number.

Error 1110: Unit numbers may not have leading zeros

Explanation: Tape and disk unit numbers may not be of the form "D03," for example, "D3" should be specified.

Retry the ADD command without any leading zeros.

Error 1120: LUN <lun> is already used

Explanation: Lun number <lun> has already been used by a disk or tape.

Retry the ADD command specifying a different LUN.

Error 1130: The unit number cannot exceed <max unit>

Explanation: You specified a unit number that was out-of-bounds.

Try to add the unit again using a unit number that is less than or equal to <max unit>.

Error 1140: Invalid unit number. Valid unit number range(s) are: <start> to <end>

Explanation: You attempted to create a unit out of the valid unit ranges. The valid unit ranges are given by the <start> and <end> values.

Retry the ADD command specifying a unit number in the correct range.

Error 1150: A restart of THIS_CONTROLLER is required before units may be added

Explanation: You changed the target IDs that the HS1CP supports without restarting the HS1CP, then tried to add a unit that is supported by the new target IDs. Before the new target ids may be used, a restart is required.

Restart the HS1CP.

Error 2000: Port must be 1 - <maximum port number>

Explanation: When adding a device, you specified a port less than 1 or greater than <maximum port number>.

Retry the command specifying a port within the range given.

Error 2010: Target must be 0 - <maximum target number>

Explanation: When adding a device, you specified a target greater than <maximum target number>.

In single HS1CP configurations, <maximum target number> is 6. In dual-redundant configurations, <maximum target number> is 5.

Error 2020: LUN must be 0 - 7

Explanation: When adding a device, you specified a LUN greater than 7.

Error 2030: This port, target LUN combination already in use by another device

Explanation: When adding a device, you specified PTL that is already specified by another device.

Error 2040: Cannot set TRANSPORTABLE when device in use by an upper layer

Explanation: A disk cannot be set to TRANSPORTABLE when it is being used by an upper level (unit or storageset).

Error 2050: Cannot set NOTRANSPORTABLE when device in use by an upper layer

Explanation: A disk cannot be set to NOTRANSPORTABLE when it is being used by an upper level (unit or storageset).

Error 2060: Can only clear UNKNOWN errors on a device

Explanation: You attempted to clear UNKNOWN on a storageset or a unit.

Check the name of the device and reissue the command.

Error 3050: <disk name> could not be initialized as a spare disk

Explanation: When adding spare disks to the spareset, they are initialized with special spare disk metadata. If the metadata cannot be written, error 3050 results.

Error 3060: <disk name> is not a member of the spareset

Explanation: You attempted to delete a disk drive from the spareset that was not a member of the spareset.

Error 3070: <disk name> is not a member of the failedset

Explanation: You attempted to delete a disk drive from the failedset that was not a member of the failedset.

Error 3080: <setname> can't be deleted

Explanation: You attempted to delete the spareset or the failedset. These containers cannot be deleted.

Error 3090: <licensable feature> support is not enabled on this HS1CP

Explanation: You attempted to use a feature that requires a license, and the license was not enabled on this HS1CP.

Error 3100: <licensable feature> support is not enabled on other HS1CP

Explanation: You attempted to use a feature that requires a license, and the license was not enabled on the other HS1CP.

Error 3110: <disk name> is not a member of <container name>, cannot remove it

Explanation: When issuing a SET <container name> REMOVE=<disk name>, the disk specified was not part of the container.

Check the device and container names and reissue the command.

Error 3120: <container name> is already reduced. Another member cannot be removed

Explanation: When issuing a SET <container name> REMOVE=<disk name>, the container was already in a reduced state. Add another disk before removing another member.

Error 3130: Unable to remove <disk name> from <container name>

Explanation: When issuing a SET <container name> REMOVE=<disk name>, the HS1CP was unable to remove the device from the RAIDset.

Check for error conditions, and if none exist, contact Digital Multivendor Customer Services.

Error 3140: <disk name> is in a spareset. Remove it from the spareset first.

Explanation: When issuing a SET <container name> REPLACE=<disk name>, the disk specified was part of the spareset. A disk to be used as a replacement must not be part of any configuration.

Error 3150: <disk name> is still part of a configuration. Delete upper configuration first.

Explanation: When issuing a SET <container name> REPLACE=<disk name>, the disk specified was part of an existing configuration. A disk to be used as a replacement must not be part of any configuration.

Error 3160: <disk name> is not a disk. Can only use disks for replacement in a raidset.

Explanation: When issuing a SET <container-name> REPLACE=<disk-name>, the device identified by <disk name> was not a disk.

Error 3170: <container name> is not reduced. Cannot replace a member.

Explanation: When issuing a SET <container-name> REPLACE=<disk-name>, the container specified was not reduced.

Remove a member before replacing it.

Error 3180: <container name> has a replacement policy specified. Cannot manually replace a member.

Explanation: When issuing a SET <container-name> REPLACE=<disk-name>, it was discovered that the container specified already had a replacement policy specified. A manual replacement cannot be done on a container with an automatic replacement policy.

Set the replacement policy for the container to NOPOLICY and try the replacement again.

Error 3190: Unable to replace <disk name> in <container name>

Explanation: When issuing a SET <container name> REPLACE=<disk name>, the HS1CP was unable to replace the device into the RAIDset.

Check for error conditions, and if none exist, contact Digital Multivendor Customer Services.

Error 3200: No other switches may be specified on a REMOVE operation.

Explanation: When issuing a SET <container name> REMOVE=<disk name>, no other switches (such as POLICY) may be specified.

Error 3210: No other switches may be specified on a REPLACE operation.

Explanation: When issuing a SET <container name> REPLACE=<disk name>, no other switches (such as POLICY) may be specified.

Error 3220: A REPLACE may not be done on a raidset that is not configured as a unit

Explanation: A REPLACE operation may not be done on a RAIDset that has not been configured as a unit.

Error 3230: <container name> is reconstructing <disk name>. Only <disk name> may be removed

Explanation: When issuing a SET <container name> REMOVE=<disk name> on a RAIDset that is already reconstructing, only the disk drive that is being reconstructed may be removed.

Error 3240: <storageset type> can't be initialized

Explanation: Sparesets and failedsets cannot be initialized.

Check the name of the container that you wish to initialize and try again.

Error 3250: A REMOVE may not be done on a raidset that is not configured as a unit

Explanation: A RAIDset must be configured as a unit before a disk can be removed to reduce the RAIDset.

Create a unit from the RAIDset and then remove the member.

Error 3260: <disk name> is a TRANSPORTABLE disk. TRANSPORTABLE disks cannot be used by storagesets. Do a SET <disk name> NOTRANSPORTABLE before using this disk in a storageset

Explanation: You cannot place a TRANSPORTABLE disk into a reduced RAIDset.

Set the disk NOTRANSPORTABLE and retry the command.

Error 4000: The CLI prompt must have 1 to 16 characters.

Explanation: This error results from a SET THIS_CONTROLLER or SET OTHER_CONTROLLER command with the qualifier PROMPT=. The length of the CLI prompt must be at least one character and may not exceed 16 characters.

Retry the command with the correct number of characters.

Error 4010: Illegal character in CLI prompt.

Explanation: A nonprintable character was specified. Only ASCII characters space " " through tilde "~" may be specified (hex 20–7E).

Error 4020: Terminal speed must be 300, 1200, 2400, 4800, 9600 or 19200

Explanation: This error results from a SET THIS_CONTROLLER or SET OTHER_CONTROLLER command with the argument TERMINAL_SPEED=. The only valid baud rates that may be specified are 300, 1200, 2400, 4800, 9600 or 19200 baud.

Retry the command with a correct terminal speed.

Error 4030: HS1CP ID must be in the range 0 to <max nodes minus 1>.

Explanation: The ID= was specified with a number greater than <max nodes minus 1>.

If increasing the HS1CP's ID, set MAX_NODES first, then the HS1CP's ID.

Error 4040: SCS nodename length must be from 1 to 6 characters.

Explanation: This error results from a SET THIS_CONTROLLER or SET OTHER_CONTROLLER command with the argument SCS_NODENAME=. The SCS node name must consist of one to six alphanumeric characters enclosed in quotes with an alphabetic character first.

Retry the command with a correct SCS node name length.

Error 4050: SCS nodename must start with an alpha character and contain only A-Z and 0-9

Explanation: This error results from a SET THIS_CONTROLLER or SET OTHER_CONTROLLER command with the argument SCS_NODENAME=. The SCS node name must consist of alphanumeric characters enclosed in quotes with an alphabetic character first.

Retry the command with a correct SCS node name.

Error 4060: Allocation class must be from <minimum> to 255

Explanation: An illegal MSCP or TMSCP allocation class was specified. The <minimum> is 0 for a single HS1CP configuration, or 1 for a dual-redundant configuration.

Error 4070: Max nodes must be 2, 8, 16 or 32

Explanation: This error results from a SET THIS_CONTROLLER or SET OTHER_CONTROLLER command with the argument MAX_NODES=. Max nodes must be 2, 8, 16 or 32 nodes.

Retry the command with a correct max node number.

Error 4080: Current node ID too large for requested max nodes setting.

Explanation: This error results from a SET THIS_CONTROLLER or SET OTHER_CONTROLLER command with the arguments MAX_NODES= or ID=. MAX_NODES= was specified with a number less than the HS1CP's ID or the HS1CP's ID was specified with a number greater than MAX_NODES—1.

If decreasing MAX_NODES, set the HS1CP's ID first, then MAX_NODES.

Error 4090: Module has invalid serial number. This HS1CP cannot be used Call Digital Services.

Explanation: This error is typically the result of faulty Non-Volatile memory. This error cannot be fixed in the field.

A replacement HS1CP must be orderd. Contact Digital Multivendor Customer Services.

Error 4100: Unable to RESTART other HS1CP.

Explanation: A communication error occurred when trying to restart the other HS1CP.

Retry the RESTART command.

Error 4110: Unable to SHUTDOWN other HS1CP.

Explanation: A communication error occurred when trying to shutdown the other HS1CP.

Retry the SHUTDOWN command.

Error 4120: Unable to SELFTEST other HS1CP.

Explanation: A communication error occurred when trying to self-test the other HS1CP.

Retry the SELFTEST command.

Error 4130: Unable to setup HS1CP restart.

Explanation: A communication error occurred when trying to restart or self-test the other HS1CP.

Retry the RESTART or SELFTEST command.

Error 4140: Unable to lock the other HS1CP's NV memory

Explanation: Most configuration commands such as ADD, DELETE, and SET require both HS1CPs in a dual-redundant configuration to be running so configuration changes can be recorded in both HS1CPs. If one HS1CP is not running, the above message results when you attempt to change the configuration.

Restart the other HS1CP and try the command again, or SET NOFAILOVER on the remaining HS1CP.

Error 4150: Unable to rundown the following units on the other HS1CP: <list of problem units>

Explanation: When attempting to shut-down, restart or selftest the other HS1CP, some units could not be successfully run down. This can be caused either by online units or errors when trying to rundown the units. Either rectify the problems on the problem units or issue the SHUTDOWN, RESTART or SELFTEST command with the OVERRIDE_ONLINE or IGNORE_ERRORS qualifiers.

Error 4160: Unable to rundown the following units on this HS1CP: <list of problem units>

Explanation: When attempting to SHUTDOWN, RESTART or SELFTEST this HS1CP, some units could not be successfully run down. This can be caused either by online units or errors when trying to rundown the units.

Either rectify the problems on the problem units or issue the SHUTDOWN, RESTART or SELFTEST command with the qualifier OVERRIDE_ONLINE or IGNORE_ERRORS.

Error 4170: Only <max targets> targets may be specified

Explanation: When setting THIS_CONTROLLER or OTHER_ CONTROLLER ID=, you specified too many IDs; you can only specify up to <max targets> IDs.

Retry the SET THIS_CONTROLLER ID= command with no more than <max targets> IDs specified.

Error 4180: Invalid unit number(s) still present that must be deleted before the HS1CP ID may be changed. All unit numbers must be in the range(s): <start> to <end>

Explanation: You attempted to change the HS1CP IDs when there were still units using those IDs. The current valid unit ranges are given by the <start> and <end> values.

Either delete the units that use the ID that will no longer be specified, or Retry the SET THIS_CONTROLLER ID= specifying the ID being used by the existing units.

Error 4190: The time must be specified in the format dd-mmm-yyyy:hh:mm:ss

Explanation: The time must be specified as shown.

Retry the command using the correct time format.

Error 4200: CACHE_FLUSH_TIMER must be in the range 1 to 65535

Explanation: The value given for the CACHE_FLUSH_TIMER is out of range.

Reissue the command specifying a number in the range shown.

Error 4210: IDs specified as preferred must be a subset of the IDs specified by the ID= argument first

Explanation: The PREFERRED_IDs specified must be a subset of the IDs (targets) supported by the HS1CP. When changing either the supported targets or the preferred ids, it was found that the PREFERRED_IDs were not a subset of the IDs.

Reissue the command with valid PREFERRED_ID= arguments or change the IDs supported by the HS1CP.

Error 5000: A program name must be from 1 to 6 characters in length

Explanation: This error results from a "RUN <program name>."

Error 5010: The requested program is currently busy.

Explanation: This error results from a "RUN <program name>." The program requested is being run by someone else.

Error 5020: The requested program is unknown.

Explanation: This error results from a "RUN <program name>."

Enter "DIR" to get a list of available programs.

Error 5030: Insufficient memory for request.

Explanation: This error results from a "RUN <program name>" resource problem. Retry the command later.

Error 6000: Communication failure with the other HS1CP.

Explanation: There was a communication problem with the other HS1CP. This typically happens if the other HS1CP is shutting down. If these messages happen often when the other HS1CP is not shutting down, call Digital Multivendor Customer Services.

Error 6010: Other HS1CP not present

Explanation: When asked to communicate with another HS1CP (the result of any one of a number of commands), the other HS1CP was found not to be running.

If the other HS1CP is in the process of restarting, retry the command later. If the other HS1CP is shut down or turned off, start it. If the other HS1CP is no longer present, enter a SET NOFAILOVER command to take it out of dual-redundant mode.

Error 6020: Initial failover handshake not yet complete

Explanation: For a short period of time after start up, the two HS1CPs must communicate to set up a dual-redundant mode. This setup time is typically less than 1 minute. If commands that require HS1CP-to-HS1CP communication are entered during this setup time, error 6020 results.

Retry the command later.

Error 6030: Unable to communicate with the other HS1CP to setup FAILOVER

Explanation: Could not setup FAILOVER due to communication problems between the HS1CPs.

Retry the command later.

Error 6040: The write of the other HS1CP's configuration information did not succeed; information may be in an inconsistent state. Before further use both HS1CPs should be removed from dual-redundant mode (SET NOFAILOVER) and then placed back into dual-redundant mode (SET FAILOVER) to assure consistency

Explanation: Communication was lost in the middle of a SET FAILOVER command.

Follow the instructions included in the error message.

Error 6050: Communication failure with other HS1CP while putting HS1CPs into dual-redundant mode. Reissue SET FAILOVER command

Explanation: Communication was lost in the middle of a SET FAILOVER command.

Follow the instructions included in the error message.

Error 6070: Illegal command—this HS1CP not configured for dual-redundancy

Explanation: A command was entered to a single HS1CP configuration that requires two HS1CPs to be in dual-redundant mode.

If two HS1CPs are supposed to be in dual-redundant mode, enter a SET FAILOVER command. If not, do not enter the command that resulted in the error.

Error 6080: Illegal command—this HS1CP not currently in dual-redundant mode

Explanation: A command was entered to a dual-redundant-configured HS1CP, but the other HS1CP was not available for communication.

Restart the other HS1CP and wait until it is communicating with this HS1CP. If this HS1CP is no longer supposed to be in dual-redundant mode, enter a SET NOFAILOVER command.

Error 6090: In failover no device may be configured at target 6 <device type> <device name> is at PTL <port> <target> <lun>

Explanation: Target addresses 6 and 7 are used by the HS1CPs when in a dual-redundant configuration. When in a single HS1CP configuration, target 6 is available for use by devices. If devices are configured at target 6 and you attempt to install a dual-redundant configuration, this error is displayed for all devices that use target 6 and the HS1CPs will not be placed in a dual-redundant configuration.

Reconfigure the drives both logically and physically so that target 6 is not used.

Error 6100: Allocation classes cannot be zero for a dual-redundant configuration. Set MSCP and TMSCP allocation classes to non-zero.

Explanation: If in a dual-redundant configuration, the allocation class must not be set to zero.

Error 6110: HS1CPs already configured for failover

Explanation: A SET FAILOVER cannot be issued on a HS1CP already in failover.

Error 6130: RAID5 in use on this HS1CP but not enabled on the other HS1CP.

Explanation: When trying to SET FAILOVER, it was discovered that there were RAID5 configurations on this HS1CP but the other HS1CP did not have the RAID5 feature enabled.

If RAID5 is licensed on the other HS1CP, enable it. If it is not licensed, either contact Digital Multivendor Customer Services for licensing information, or do not use the two HS1CPs in dual-redundant mode, or do not use a RAID5 configuration.

Error 6140: Writeback cache in use on this HS1CP but not enabled on the other HS1CP.

Explanation: When trying to SET FAILOVER, it was discovered that there were write-back cache switches set on this HS1CP but the other HS1CP did not have the write-back cache feature enabled.

If write-back cache is licensed on the other HS1CP, enable it. If it is not licensed, either contact Digital Multivendor Customer Services for licensing information, or do not use the two HS1CPs in dual-redundant mode, or do not use the write-back cache switches.

Error 6150: RAID5 in use on other HS1CP but not enabled on this HS1CP

Explanation: When trying to SET FAILOVER, it was discovered that there were RAID5 configurations on the other HS1CP but this HS1CP did not have the RAID5 feature enabled.

If RAID5 is licensed on this HS1CP, enable it. If it is not licensed, contact Digital Multivendor Customer Services for licensing information or do not use the two HS1CPs in dual-redundant mode, or do not use a RAID5 configuration.

Error 6160: Writeback cache in use on other HS1CP but not enabled on this HS1CP.

Explanation: When trying to SET FAILOVER, it was discovered that there were write-back cache switches set on the other HS1CP but this HS1CP did not have the write-back cache feature enabled.

If write-back cache is licensed on this HS1CP, enable it. If it is not licensed, either contact Digital Multivendor Customer Services for licensing information or do not use the two HS1CPs in dual-redundant mode, or do not use the write-back cache switches.

Error 6170: An <<REFERENCE>(MIKASA_CONN_SHORT) type> and <<REFERENCE>(MIKASA_CONN_SHORT) type> cannot configured for failover

Explanation: Two different HS1CPs cannot be configured for failover.

Replace the other HS1CP with the same model as this one and reissue the command.

Error 7000: Can only clear LOST_DATA cache errors on a unit.

Explanation: you specified something other than a unit for clearing the LOST_DATA cache error.

Error 7010: Can only clear UNWRITEABLE_DATA cache errors on a unit.

Explanation: You specified something other than a unit for clearing the UNWRITEABLE_DATA cache error.

Error 7020: Can only retry UNWRITEABLE_DATA cache errors on a unit

Explanation: You specified something other than a unit for retrying a write on a UNWRITEABLE_DATA cache error.

Error 7030: Unable to force write of unwriteable data

Explanation: A RETRY UNWRITEABLE_DATA command could not write the UNWRITEABLE_DATA.

Error 7040: Unable to rundown unit before clearing error

Explanation: To clear UNWRITEABLE_DATA and LOST_DATA errors, the unit must be rundown before the error is cleared. If the unit could not be rundown, the above error results. If this error persists, call Digital Multivendor Customer Services.

Error 7050: Unable to runup unit after clearing error. This HS1CP must be restarted

Explanation: To clear UNWRITEABLE_DATA and LOST_DATA errors, the unit must be rundown before the error is cleared. If the unit was rundown and the error was cleared and then the unit was unable to be run back up, the unit will remain unavailable until the HS1CP is restarted.

Error 9000: Cannot rename a unit

Explanation: Only devices and storagesets may be renamed. If you attempt to rename a unit, the above message results.

Error 9010: <name> is an illegal name, it must be from 1 to 9 characters.

Explanation: This error results from an ADD command with an illegal name given.

Error 9020: <name> is an illegal name, it must start with A-Z

Explanation: This error results from an ADD command with an illegal name given.

Error 9030: <name> is an illegal name, characters may consist only of A-Z, 0-9, ., - or _

Explanation: This error results from an ADD command with an illegal name given.

Error 9040: <name> conflicts with keyword <keyword>

Explanation: The name given in an ADD command conflicts with a CLI keyword.

Specify another name.

Error 9050: Configuration area full

Explanation: The total number of units, devices, and storagesets that can be configured is 195 in any combination. This error results when you exceed that number of nodes.

Delete some units or devices in order to recover some configuration nodes.

Error 9060: <name> does not exist

Explanation: Some operation (SET, DELETE, INITIALIZE, and so forth) specified a name that does not exist.

Check the name and retry the command.

Error 9070: <name> is part of a configuration

Explanation: Devices may not be deleted if they are still in use by storagesets or units. Storagesets may not be deleted if they are still used by units.

Delete configurations from the top down; delete units, then stripesets, and RAIDsets (if any), and then finally devices.

Error 9080: <name> is already used

Explanation: An ADD command specified a name that is already in use.

Specify another name.

Error 9090: A <device type> cannot be used in a <storageset type>

Explanation: The device specified cannot be used in the storageset specified, for example, tapes cannot be bound into a stripeset.

Reexamine the configuration and correct the incompatibility.

Error 9100: A <storageset type> must have from <minimum> to <maximum> entities

Explanation: The wrong number of devices was specified for this storageset. Different storagesets require different numbers of devices.

Reexamine the configuration, then correct the number of devices.

Error 9130: Cannot delete ONLINE unit

Explanation: Unit specified in a DELETE command is online to a host.

Dismount the unit at the host then retry the command. Or add the OVERRIDE_ONLINE qualifier to the DELETE command.

Error 9140: Cannot delete exclusive access unit

Explanation: Unit specified in a DELETE command is set up for exclusive access.

Take the unit out of exclusive access mode and retry the command.

Error 9150: INITIALIZE is no longer supported at the unit level. You must INITIALIZE the container that makes up this unit

Explanation: You tried to initialize a unit. Units may no longer be initialized. The container that makes up the unit must be initialized before a unit is created out of the container.

Error 9160: Non-disk devices cannot be INITIALIZED

Explanation: Tapes and CDROMS may not be initialized.

Error 9170: <device type> <device name> at PTL <port> <target> <lun> No device installed

Explanation: When a unit is added or initialized, the configuration of the devices that makes up the unit is checked. If no device is found at the PTL specified, this error is displayed.

Check both the logical and physical configuration of the unit and correct any mismatches.

Error 9180: <device type> <device name> at PTL <port> <target> <lun> Incorrect device type installed

Explanation: When a unit is added or initialized, the configuration of the devices that make up the unit is checked. If a non disk device is found at the PTL specified, this error is displayed.

Check both the logical and physical configuration of the unit and correct any mismatches.

Error 9190: Unit <unum> is currently online

Explanation: When a SHUTDOWN, RESTART, or SELFTEST command is entered without the OVERRIDE_ONLINE qualifier and online devices are found, the command is aborted and the units currently online are listed.

Either retry the command with OVERRIDE_ONLINE qualifier or dismount all devices from the hosts.

Error 9200: <name> conflicts with unit names

Explanation: This error results from an ADD command. Names in the format of D*n* and T*n*, when *n* is a number from 0 to 4094, are reserved for units. Rename the storageset or device that is being added so it does not conflict with the unit names and retry the command.

Error 9210: Cannot check if drives are online to the other HS1CP

Explanation: When trying to check for online drives on the other HS1CP, there was a communication failure.

Retry the command.

Error 9230: Unable to modify switches requested

Explanation: This error results from a SET command. The system is currently busy.

Retry the SET command later.

Error 9240: Cannot delete unit in maintenance mode

Explanation: When trying to delete a unit, the unit was found to be in maintenance mode. This is typically the result of trying to delete a unit that is in use by DILX or TILX.

Ensure that DILX and TILX is not being run against the unit that is to be deleted, and retry the command.

Error 9250: Initialize of disk failed

Explanation: Unable to write metadata on disk.

Make sure the disk is not broken.

Error 9260: Cannot INITIALIZE a container that is still part of a configuration. Delete upper configuration first

Explanation: A container cannot be initialized that is part of another configuration or is being used by a unit.

Delete the upper configuration and reissue the INITIALIZE command.

Error 9270: No metadata found on container, unit not created. An INITIALIZE <container name> must be issued before this container may be used

Explanation: You attempted to create a unit from a container that did not have valid metadata.

INITIALIZE the metadata on the container, then create a unit out if it.

Error 9290: Communication failure with other HS1CP, cannot check other HS1CP's licensing

Explanation: Unable to communicate with the other HS1CP to check licensing before creating a RAIDset or enabling write-back cache.

Check to make sure that both HS1CPs are running. If one is broken, take this HS1CP out of failover (SET NOFAILOVER) and reissue the command.

Error 9320: Bad write-back cache or battery on HS1CP

Explanation: If you attempt to set write-back cache on a device or create a RAIDset, and there is a bad cache or cache battery on a HS1CP, write-back is not set, or the RAIDset is not created and the above message is displayed.

Error 9330: NV memory write collision. Please try again

Explanation: Two people were trying to configure the CLI at the same time.

Check the configuration you were trying to modify to make sure it's unchanged and retry the command.

Error 9340: Reduced raidsets cannot be INITIALIZED

Explanation: You cannot INITIALIZE a RAIDset that is running in reduced state.

Replace a member and try again.

Error 9360: A tape is not installed at the PTL <port> <target> <lun>. Cannot set tape switches unless a tape is installed

Explanation: A SET or ADD command specified a tape format, but there was no tape installed at the tape's PTL.

Install a tape and retry the command.

Error 9370: A <tape name> is an unsupported device. Tape switches cannot be set on unsupported devices

Explanation: The tape installed is not currently supported by the HS1CP.

Replace the tape with a supported device and retry the command.

Error 9380: Unable to allocate unit for NORUN to RUN transition

Explanation: The unit could not be allocated so the HS1CP could do a RUN/NORUN transition.

Retry the command. If this error persists, call Digital Multivendor Customer Services.

Error 9390: Cannot change default tape format while tape drive online to host

Explanation: The default tape format cannot be changed when the tape drive is online to a host.

Dismount the tape drive from the host and retry the command.

Error 9400: Cannot rundown or allocate unit in order to delete it

Explanation: Retry the command. If this error persists, call Digital Multivendor Customer Services.

Error 9410: Cannot delete unit—<type> error exists on unit that must be cleared first. To clear error type: <clear error string>

Explanation: Units cannot be deleted if cache errors exist. Any cache errors must be cleared before a unit can be deleted.

Issue the <clear error string> command and then delete the unit.

Error 9420: Unit <unit number> has unflushed data or a cache error and must be deleted on this HS1CP

Explanation: When trying to set failover a unit with unflushed data or a cache error was detected on this HS1CP.

Delete the unit as requested and then retry the SET FAILOVER command.

Error 9430: Cannot check if drives have unflushed data or cache errors on the other HS1CP

Explanation: Communication error when trying to SET FAILOVER.

Retry the command. If this error persists, call Digital Multivendor Customer Services.

Error 9440: Unit <unit number> has unflushed data or a cache error and must be deleted on the other HS1CP

Explanation: When trying to set failover a unit with unflushed data or a cache error was detected on the other HS1CP.

Delete the unit as requested and then retry the SET FAILOVER command.

B.3.3 Warning Conventions

A Warning *nnnn:* message means that the command completed, but there is a situation that you should be aware of. Typically, but not always, a warning will result in an unusable configuration; you will have to either logically reconfigure the cabinet using the CLI or physically reconfigure the cabinet by moving the disks around.

Multiple warning messages may result from one command.

Items in angle brackets (<>) are replaced at run time with names, numbers, and so on.

B.3.4 CLI Warning Messages

Warning 1000: You should have read the user guide and fully understand the implications of setting WRITEBACK_CACHE

Explanation: Using write-back cache introduces behaviors that you should completely understand before using. See the full documentation on write-back cache in the user guide.

Warning 3000: This storageset is configured with more than one disk per port. This causes a degradation in performance

Explanation: This warning results from an ADD *storageset-type* command. The storageset specified has more than one member per port. One method of increasing the HS1CP's performance is through parallel transfers to members of a storageset. If multiple members of a storageset are on one port, transfers must be done in serial to those members.

Though multiple storageset members on one port will work, it is strongly recommended that the storageset be deleted and reconfigured with one member per port.

Warning 3010: Unable to check all device types that make up this storageset. If the storageset is made up of different device types, it may result in a storageset of reduced size

Explanation: This warning results from an ADD *storageset-type* command. Device types being added to a storageset are checked to assure that they are the correct device types. If one or more devices could not be checked, the above warning is displayed.

You should check all the devices to assure that they are correctly installed and configured.

Warning 3020: This storageset is configured with different device types. This may result in a storageset of reduced size

Explanation: This warning results from an ADD *storageset-type* command. Device types being added to a storageset are checked to assure that they are the same types. If all devices are not the same, the above warning is reported. Storageset size is determined by the size of the smallest device, so the storageset configured is of reduced size.

If a reduced size storageset is acceptable, nothing needs to be done in response to the above warning. To realize the maximum storageset size, all devices that make up the storageset should be identical.

Warning 4000: A restart of this HS1CP is required before all the parameters modified will take effect

Explanation: This warning results from a SET THIS_CONTROLLER command. Some HS1CP parameters require a restart before they can take effect. If any of those parameters are changed, this warning is displayed.

It is recommended that a restart via the "RESTART THIS_CONTROLLER" command be done as soon as possible.

Warning 4010: A restart of the other HS1CP is required before all the parameters modified will take effect

Explanation: This warning results from a SET OTHER_CONTROLLER command. Some HS1CP parameters require a restart before they can take effect. If any of those parameters are changed, this warning is displayed.

Restart the HS1CP and retry the command.

Warning 4020: A restart of both this and the other HS1CP is required before all the parameters modified will take effect

Explanation: This warning results from a SET THIS_CONTROLLER or a SET OTHER_CONTROLLER command. Some HS1CP parameters require a restart of both HS1CPs before they can take effect. If any of those parameters are changed, this warning is displayed. Restart both HS1CPs and retry the command.

Warning 6000: Communication failure with the other HS1CP while taking HS1CPs out of dual-redundant mode. Issue a SET NOFAILOVER command on the other HS1CP

Explanation: This warning results from a SET NOFAILOVER command. This HS1CP was unable to communicate with the other HS1CP to notify it that it is no longer in dual-redundant mode. Typically, this occurs when the other HS1CP has already been removed prior to the SET NOFAILOVER command.

Enter a SET NOFAILOVER command on the other HS1CP as soon as possible.

Warning 6010: Licensing different between the two HS1CPs

Explanation: If the licensing is not identical on both HS1CPs in a dual-redundant configuration, the above warning is displayed.

You should check the licensing on both HS1CPs and make sure they are identical.

Warning 7000: Data written successfully before clearing unwriteable data error

Explanation: As a result of a CLEAR UNWRITEABLE_DATA, if the last-ditch attempt to write data before clearing the error was successful, the above warning is displayed.

Note

This means that no customer data was lost, so this warning is actually good.

Warning 7010: Unable to clear LOST_DATA on other HS1CP

Explanation: When trying to clear LOST_DATA on the other HS1CP, a communication error occurred.

Retry the command. If the failure persists, contact Digital Multivendor Customer Services.

Warning 7020: Unable to clear UNWRITEABLE_DATA on other HS1CP

Explanation: When trying to clear UNWRITEABLE_DATA on the other HS1CP, a communication error occurred.

Retry the command. If the failure persists, contact Digital Multivendor Customer Services.

Warning 9030: Cannot determine if the correct device type is at the PTL specified

Explanation: When a device is added, the location specified is checked to see if the correct device type is present. This warning results when no device responds from the location specified.

Check the physical configuration and the PTL that was specified.

Warning 9040: There is currently a <device type> at the PTL specified

Explanation: When a device is added, the location specified is checked to see if the correct device type is present. This warning results when a device different from the one specified is found at the location specified (for example, a tape is found where a disk was added).

Check the physical configuration and the PTL that was specified.

Warning 9050: <device type> <device name> at PTL <port> <target> <lun> No device installed

Explanation: When a unit is added, the configuration of the disks that make up the unit is checked. If no device is found at the PTL specified, this warning is displayed.

Check both the logical and physical configuration of the devices that make up the unit and correct any mismatches.

Warning 9060: <device type> <device name> at PTL <port> <target> <lun> Incorrect device type installed

Explanation: When a unit is added, the configuration of the disks that make up the unit is checked. If a non disk device is found at the PTL specified, this warning is displayed.

Check both the logical and physical configuration of the devices that make up the unit and correct any mismatches.

Warning 9080: <license> support is not licensed on HS1CP. Any use of this feature requires licensing. Continued use does not comply with the terms and conditions of licensing for this product.

Explanation: You have a licensed feature enabled on this HS1CP but it is not licensed. This is against the contractual agreement between Digital and your company. Please disable the licensed feature and contact Digital Multivendor Customer Services if you wish to purchase it.

Warning 9090: Metadata found on container. Are you sure this is a TRANSPORTABLE container?

Explanation: When a transportable disk was initialized, metadata was found.

Verify that this disk in fact should be marked transportable. No action is required to correct this warning.

Warning 9100: Bad or low battery or bad write cache on HS1CP writeback cache will not be used

Explanation: The battery is low or bad on the specified HS1CP. The unit specified will not use write-back cache until the battery is charged or repaired.

Warning 9110: Bad or low battery or bad write cache on HS1CP this unit cannot be used by HS1CP

Explanation: The battery is low or bad on the specified HS1CP. The unit specified requires the use of write-back cache, so its use has been disabled until the battery is charged or repaired.

B.4 Device Configuration Examples

The following examples cover the majority of configurations and the method of defining those configurations.

Example B–2 Initial Single HS1CP Configuration

```
HS1CP1> SET THIS_CONTROLLER ID=5 SCS_NODENAME="HS1CP1"
HS1CP1> SET THIS_CONTROLLER MSCP_ALLOCATION_CLASS=4 TMSCP_ALLOCATION_CLASS=4
HS1CP1> RESTART THIS_CONTROLLER
[this HS1CP restarts at this point]
HS1CP1> SET THIS_CONTROLLER PATH
```

Example B–3 Setting the Terminal Speed and Parity

```
HS1CP1> SET THIS_CONTROLLER TERMINAL_SPEED=19200 NOTERMINAL_PARITY
```

Note

Garbage will appear on the terminal after setting the HS1CP's terminal speed until you set the terminal's speed to match the HS1CP's new terminal speed.

Example B–4 Creating a Unit from a Disk Device

```
HS1CP1> ADD DISK DISK0 2 0 0
HS1CP1> INITIALIZE DISK0
HS1CP1> ADD UNIT D0 DISK0
```

Example B–5 Creating a Unit from a Tape Device

```
HS1CP1> ADD TAPE TAPE0 3 0 0
HS1CP1> ADD UNIT T0 TAPE0
```

Example B–6 Creating a Unit from a Four-Member Stripeset

```
HS1CP1> ADD DISK DISK0 1 0 0
HS1CP1> ADD DISK DISK1 2 0 0
HS1CP1> ADD DISK DISK2 3 0 0
HS1CP1> ADD DISK DISK3 1 1 0
HS1CP1> ADD STRIPESET STRIPE0 DISK0 DISK1 DISK2 DISK3
Warning 3000: This storageset is configured with more than one disk per port.
              This causes a degradation in performance
HS1CP1> INITIALIZE STRIPE0
HS1CP1> ADD UNIT D0 STRIPE0
```

Example B–7 Creating a Unit from a Five-Member RAIDset

```
HS1CP1> ADD DISK DISK0 1 0 0
HS1CP1> ADD DISK DISK1 2 0 0
HS1CP1> ADD DISK DISK2 3 0 0
HS1CP1> ADD DISK DISK3 1 1 0
HS1CP1> ADD DISK DISK4 2 1 0
HS1CP1> ADD RAIDSET RAID9 DISK0 DISK1 DISK2 DISK3 DISK4
Warning 3000: This storageset is configured with more than one disk per port.
              This causes a degradation in performance
HS1CP1> INITIALIZE RAID9
HS1CP1> ADD UNIT D0 RAID9
```

Example B–8 Creating a Unit from a Disk Device and Setting the Write Protection

```
HS1CP1> ADD DISK DISK0 2 0 0
HS1CP1> INITIALIZE DISK0
HS1CP1> ADD UNIT D0 DISK0 WRITE_PROTECT
```

Example B–9 Setting the Write Protection for an Existing Unit

```
HS1CP1> ADD DISK DISK0 2 0 0
HS1CP1> INITIALIZE DISK0
HS1CP1> ADD UNIT D0 DISK0
HS1CP1> SET D0 WRITE_PROTECT
```

Example B–10 Renumbering Disk Unit 0 to Disk Unit 100

```
HS1CP1> ADD DISK DISK0 2 0 0
HS1CP1> INITIALIZE DISK0
HS1CP1> ADD UNIT D0 DISK0
HS1CP1> DELETE D0
HS1CP1> ADD UNIT D100 DISK0
```

Note

No INITIALIZE is required because DISK0 has already been initialized.

Example B–11 Creating a Transportable Unit from a Disk Device

```
HS1CP1> ADD DISK DISK0 2 0 0 TRANSPORTABLE
HS1CP1> INITIALIZE DISK0
HS1CP1> ADD UNIT D0 DISK0
```

[or]

```
HS1CP1> ADD DISK DISK0 2 0 0
HS1CP1> SET DISK0 TRANSPORTABLE
HS1CP1> INITIALIZE DISK0
HS1CP1> ADD UNIT D0 DISK0
```

Example B–12 Changing the Replacement Policy of a RAIDset

```
HS1CP1> ADD DISK DISK0 1 0 0
HS1CP1> ADD DISK DISK1 2 0 0
HS1CP1> ADD DISK DISK2 3 0 0
HS1CP1> ADD DISK DISK3 4 0 0
HS1CP1> ADD DISK DISK4 5 0 0
HS1CP1> ADD RAIDSET RAID9 DISK0 DISK1 DISK2 DISK3 DISK4
HS1CP1> INITIALIZE RAID9
HS1CP1> ADD UNIT D0 RAID9
HS1CP1> SET RAID9 POLICY=BEST_FIT
```

Note

The replacement policy can be changed at any time.

Example B–13 Deleting the Unit, Stripeset, and All Disks Associated with a Stripeset

```
HS1CP1> DELETE D0
HS1CP1> DELETE STRIPE0
HS1CP1> DELETE DISK0
HS1CP1> DELETE DISK1
HS1CP1> DELETE DISK2
HS1CP1> DELETE DISK3
```

StorageWorks Glossary

ac distribution

The method of distributing ac power in a cabinet.

ac power supply

A power supply designed to produce dc power from an ac input.

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. *See* **signal converter**.

American National Standards Institute

See **ANSI**.

ANSI

American National Standards Institute. An organization that develops and publishes electronic and mechanical standards.

array controller

A hardware/software device that facilitates communications between a host and one or more devices organized in an array. HS-family controllers are examples of array controllers.

BA350–M*x* controller shelf

The StorageWorks controller shelf used for HS-family controller modules, cache modules, and shelf power units.

BA350–S*x* SBB shelf

A StorageWorks shelf used for only power units and SBBs.

bad block

A disk drive data block containing a physical defect.

bad block replacement

See **BBR**.

battery backup unit

See **BBU**.

BBR

Bad block replacement. The substitution of defect-free disk blocks for those found to have defects.

BBU

Battery backup unit. Extends power availability after the loss of primary ac power or a power supply to protect against the corruption or loss of data.

block

A stream of data transferred as a unit. (Used interchangeably with the term **sector** for disk drives.)

blower

An airflow device mounted in a StorageWorks shelf.

cable distribution unit

See **CDU**.

carrier

A standard, StorageWorks shelf-compatible, plastic shell into which a device can be installed. Sometimes called SBB carrier.

CDU

Cable distribution unit. The AC power entry device for StorageWorks cabinets. The unit provides the connections necessary to distribute ac power to cabinet shelves and fans.

CLI

Command line interpreter. Operator command line interface for the HS1CP firmware.

cold swap

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

command line interpreter

See **CLI**.

controller

A hardware/software device that facilitates communications between a host and one or more devices. A controller translates bus protocols and hardware interfaces and adds functionality to the host/device communications.

controller shelf

A StorageWorks shelf designed to contain controller and cache memory modules.

DAT

Digital Audio Tape. A format for recording digital data on a cartridge tape.

data center cabinet

A generic reference to the large cabinets, such as the SW800 series, in which StorageWorks components can be mounted.

deskside expansion unit

A pedestal enclosure designed to house one or two StorageWorks shelves in a vertical position.

Diagnostics and Utilities Protocol

See **DUP**.

digital audio tape

See **DAT**.

DIGITAL Storage Architecture

See **DSA**.

DIGITAL Storage System Interconnect

See **DSSI**.

DILX

Disk inline exerciser. Diagnostic firmware used to test the data transfer capabilities of disk drives in a way that simulates a high level of user activity.

Disk Inline Exerciser

See **DILX**.

DSA

DIGITAL Storage Architecture. A set of specifications and interfaces describing standards for designing mass storage products. DSA defines the functions performed by host computers, controllers, and drives. It also specifies how they interact to accomplish mass storage management.

DSSI

Digital Storage System Interconnect. A Digital-specific data bus with an 8-bit data transfer rate of 4 to 5 MB/s.

dual cabinet power configuration

A cabinet ac power configuration in which two ac sources and two ac power supplies are used to provide redundant dc power to each of the cabinet's SBB shelves.

dual shelf power configuration

A cabinet ac power configuration in which one ac source and two ac power supplies are used to provide redundant dc power to each of the cabinet's SBB shelves.

dual universal asynchronous receiver transmitter

See **DUART**.

dual-redundant configuration

A controller 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.

DUART

Dual Universal Asynchronous Receiver Transmitter. An integrated circuit containing two serial, asynchronous transceiver circuits.

DUP

Diagnostic and Utility Protocol. Host application software that allows a host operator terminal to connect to the controller's command line interpreter. *See also* **virtual terminal**.

ECC

Error correction code. One or more cyclic redundancy check (CRC) words that allow detection of a mismatch between transmitted and received data in a communications system, or between stored and retrieved data in a storage system. The ECC allows for location and correction of an error in the received/retrieved data. All ECCs have limited correction power.

EDC

Error detection code. One or more checksum words that allow detection of a mismatch between transmitted and received data in a communications system, or between stored and retrieved data in a storage system. The EDC has no data correction capability.

electromagnetic interference

See **EMI**.

electrostatic discharge

See **ESD**.

EMI

Electromagnetic interference. The impairment of a signal by an electromagnetic disturbance.

error correction code

See **ECC**.

error detection code

See **EDC**.

ESD

Electrostatic discharge. The discharge of a potentially harmful static electric voltage as a result of improper grounding.

failover

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

fan

An airflow device mounted in a StorageWorks cabinet.

field replaceable unit

See **FRU**.

filler panel

A sheet metal or plastic panel used to cover unused mounting areas in StorageWorks cabinets and shelves.

FRU

field replaceable unit. A hardware component that can be replaced in the field.

full-height device

A single device that occupies an entire 5.25 inch SBB carrier. StorageWorks full-height devices have an order number suffix of "–VA."

half-height device

A device that occupies half of a 5.25 inch SBB carrier. Two half-height devices can be mounted in a 5.25 inch SBB carrier. The first half-height device is normally mounted in the lower part of the carrier. The second device is normally mounted in the upper part of the carrier.

Hierarchical Storage Operating Firmware

See **HSOF**.

HIS

Host Interconnect Services. The firmware that communicates with the host in HS-family controllers.

host

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

Host Interconnect Services

See **HIS**.

HSOF

Hierarchical Storage Operating Firmware. HS-family controller firmware is contained on a removable PCMCIA card.

KILL line

The controller-to-controller disable signal used in a dual-redundant configuration.

least recently used

See **LRU**.

logical units

A group of devices addressable as a virtual unit.

LRU

Least recently used. The block replacement algorithm for the read cache.

maintenance terminal

The operator terminal used to identify an HS-family controller, to enable its host paths, to define its subsystem configuration, and to check its status. The HS-family maintenance terminal interface is designed to accept any terminal conforming to EIA–423. The maintenance terminal is only required to configure a storage subsystem and is not required for normal operations.

nonredundant

A configuration in which there is no backup hardware in place for the hardware that is present.

nonvolatile

See **NV**.

NV

Nonvolatile. A term used to describe memory, the contents of which survive loss of power.

OCP

Operator control panel. The control/indicator panel associated with a device. The OCP is usually mounted on the device and is accessible to the operator.

operator control panel

See **OCP**.

port

The hardware and software used to connect a host controller to a communication bus, such as a CI, SCSI, or SDI bus.

qualified device

A device that has been fully tested in all appropriate StorageWorks hardware and software configurations, and is in complete compliance with Digital and country-specific standards,for example, FCC and TÜV.

quiesce

To make a bus inactive or dormant. The operator must quiesce SCSI bus operations, for example, during a device warm swap.

radio frequency interference

See **RFI**.

RAID

Redundant array of independent disks. A set of storage techniques devised to increase the performance and availability of a storage subsystem.

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. *See also* **write through cache**.

Redundant Array of Independent Disks

See **RAID**.

RFI

Radio frequency interference. The impairment of a signal by an unwanted radio signal or radio disturbance.

SBB

StorageWorks building block. A device housed in a standard StorageWorks SBB carrier. An SBB has a standard physical and electrical interface that is compatible with those of StorageWorks shelves and enclosures.

SBB shelf

StorageWorks building block shelf. A StorageWorks shelf, such as the BA350–S*x*, designed to house plug-in SBB modules.

SCS

System Communication Services. A delivery protocol for packets of information (commands or data) to or from the host.

SCSI

Small Computer System Interface. An ANSI interface defining the physical and electrical parameters of a parallel I/O bus used to connect hosts to a maximum of seven devices. The StorageWorks device interface is implemented according to the SCSI–2 standard, allowing the synchronous transfer of 8-bit data at rates of up to 10 MB/s.

shelf brackets

Sheet metal components designed to attach and position StorageWorks shelves in their associated enclosures.

signal converter

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

single cabinet power configuration

A cabinet ac power configuration in which only one ac source and one ac power supply is used to supply dc power to the cabinet's SBB shelves.

skirt

A trim panel designed to mount around the base of the cabinet.

Small Computer System Interface

See **SCSI**.

storage sets

A grouping of disk drives that make up a new distinct container.

StorageWorks

Digital's family of modular data storage products that allows customers to design and configure their own storage subsystems. Components include power, packaging, cabling, devices, controllers, and software. Customers can integrate devices and array controllers in StorageWorks enclosures to form storage subsystems.

StorageWorks building block

See **SBB**.

stripesets

In a RAID configuration, a virtual disk drive with its physical data spread across multiple physical disks. Stripeset configurations do not include a data recovery mechanism.

supported device

A device tested as functionally compatible with an approved StorageWorks hardware and software configuration.

System Communications Services

See **SCS**.

Tape Inline Exerciser

See **TILX**.

TILX

Tape Inline Exerciser. Diagnostic firmware used to test the data transfer capabilities of tape drives in a way that simulates a high level of user activity.

VAXcluster System Console

See **VCS**.

VCS

VAXcluster System Console. This terminal allows access to hosts (by networks). Another method of accessing the controller. *See also* **DUP**.

virtual terminal

A software path from an operator terminal on the host to the controller's CLI interface. The path can be established via the host port on the controller (using DUP) or via the maintenance port through an intermediary host (VCS). A virtual terminal is also sometimes called a host console.

warm swap

A controller function that allows devices to be added, removed, or replaced while the subsystem remains operational. All activity on the device's SCSI bus must normally be halted for the duration of the warm swap operation.

write through cache

A technique for handling host write requests in read caches. When the host requests a write operation, the cache 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. *See also* **read cache**.

write-back cache

A cache configuration that increases the performance of host write requests. When the host requests a write operation, the cache 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.

Index

A

B

C

D

E

F

G

H

I

L

M

N

P

Q

R

S