StorageWorks™ Solutions

StorageWorks FDDI Server Service Manual

Order Number: EK–FSERV–SM. A01

This manual describes the procedures and tests used to troubleshoot and service the HS1*xx* and HS2*xx* StorageWorks FDDI Servers.

PHOTOGRAPHS DO NOT REPRODUCE
WELL ON MICROFICHE, THEREFORE
THEY MAY NOT BE READABLE

Digital Equipment Corporation
Maynard, Massachusetts

October 1995

While Digital believes the information included in this publication is correct as of the date of publication, it is subject to change without notice.

Digital Equipment Corporation makes no representations that the interconnection of its products in the manner described in this document will not infringe existing or future patent rights, nor do the descriptions contained in this document imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

© Digital Equipment Corporation, 1995.
All Rights Reserved.
Printed in the United States of America.

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.

Alpha, DECchip, DECevent, DECserver, Digital, HSC, HS1CP, OpenVMS, StorageWorks, VAXsimPLUS, VMS, VMScluster, and the DIGITAL logo are trademarks of Digital Equipment Corporation.

i960 and Intel are registered trademarks of Intel Corporation.

IBM is a registered trademark of International Business Machines Corporation.

NCR is a registered trademark of National Computer Systems, Inc.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Ltd.

PostScript is a registered trademark of Adobe Systems, Inc.

This document was prepared using VAX DOCUMENT Version 2.1.

Contents

Preface xxi

Manufacturer's Declarations xxv

1 Troubleshooting Strategy

1.1 Troubleshooting the System 1–1
1.2 Maintenance Features 1–4
1.3 Common Problems 1–5
1.4 Troubleshooting Basics 1–6
1.5 Service Tools and Utilities 1–6
1.5.1 Error Handling/Logging 1–6
1.5.2 ROM-Based Diagnostics 1–6
1.5.3 Loopback Tests 1–7
1.5.4 Firmware Console Commands 1–7
1.5.5 Operating System Exercisers (DEC VET) 1–7
1.5.6 Crash Dumps 1–7
1.6 Information Services 1–7
1.6.1 Training 1–8
1.6.2 Digital Assisted Services 1–8

2 Power-Up Diagnostics for the Server Processor

2.1 Interpreting Error Beep Codes 2–1
2.2 Power-Up Screen 2–5
2.3 Storage Device LEDs 2–6
2.4 EISA Bus Problems Indicated at Power-Up 2–7
2.4.1 Additional EISA Troubleshooting Tips 2–8
2.4.2 The Fail-Safe Loader 2–8
2.5 The Power-Up Sequence 2–9
2.5.1 AC Power-Up Sequence 2–10
2.5.2 DC Power-Up Sequence 2–10
2.6 Firmware Power-Up Diagnostics 2–10
2.6.1 Serial ROM Diagnostics 2–10
2.6.2 Console Firmware-Based Diagnostics 2–11
2.6.3 Operator Control Panel Indicators 2–11
2.6.4 Cycling Power with the OCP 2–13
2.6.5 Server Fan Fail Indicator 2–13

3 Power-Up Diagnostics for HS1CP

3.1 The Power-Up Sequence . . . 3–2
3.2 Firmware Power-Up Diagnostics . . . 3–2
3.2.1 HS1CP Front Panel Indicators . . . 3–2
3.2.2 HS1CP Normal Operation . . . 3–2
3.2.3 Fault Notification . . . 3–2
3.2.4 Device LEDs . . . 3–3
3.2.4.1 Storage SBB Status . . . 3–3
3.2.4.2 Device Shelf Status and Power Supply Status . . . 3–5

4 Running System Diagnostics

4.1 Diagnostic Procedures for the Server Processor . . . 4–1
4.1.1 Running ROM-Based Diagnostics . . . 4–1
4.1.2 Command Summary . . . 4–1
4.1.3 Command Reference . . . 4–2
4.1.3.1 The TEST Command . . . 4–3
4.1.3.2 The CAT EL and MORE EL Commands . . . 4–5
4.1.3.3 The MEMORY Command . . . 4–6
4.1.3.4 The NETEW Command . . . 4–7
4.1.3.5 The NETWORK Command . . . 4–8
4.1.3.6 The NET -S Command . . . 4–9
4.1.3.7 The NET -IC Command . . . 4–10
4.1.3.8 The KILL and KILL_DIAGS Commands . . . 4–11
4.1.3.9 The SHOW_STATUS Command . . . 4–12
4.1.4 Acceptance Testing and Initialization . . . 4–13
4.1.5 DEC VET . . . 4–13
4.2 Diagnostic Procedures for the HS1CP . . . 4–14

5 Normal Operation

5.1 HS241 Description . . . 5–2
5.2 HS280 Description . . . 5–3
5.3 HS221/121 Description . . . 5–4
5.4 HS211/111 Description . . . 5–5
5.5 HS210/110 Description . . . 5–6
5.6 Initialization . . . 5–6
5.6.1 HS1CP Initialization . . . 5–6
5.6.2 Dual-Redundant Configuration Initialization . . . 5–7
5.6.3 Subsystem Initialization . . . 5–7
5.7 HS1CP Front Panel . . . 5–7
5.8 Connecting a Maintenance Terminal . . . 5–7
5.9 Command Line Interpreter . . . 5–8
5.9.1 Accessing the CLI . . . 5–8
5.9.2 Exiting the CLI . . . 5–8
5.9.3 Command Sets . . . 5–9
5.9.4 Initial Configuration (Nonredundant HS1CP) . . . 5–9
5.9.5 Initial Configuration (Dual-Redundant HS1CPs) . . . 5–11
5.9.6 Configuring Storage Devices . . . 5–13
5.10 Acceptance Test . . . 5–14
5.11 Virtual Terminal . . . 5–14
5.12 VAXcluster Console System . . . 5–14
5.13 Operating Systems . . . 5–15
5.14 Failover . . . 5–16

5.14.1 Setting Failover 5–16
5.14.2 Failing Over 5–17
5.14.3 Exiting Failover 5–17
5.14.3.1 Before Failover Occurs 5–17
5.14.3.2 After Failover Occurs 5–18
5.14.4 Failover Setup Mismatch 5–18
5.14.5 Preferred Path 5–18
5.15 Moving Devices Between HS1CPs 5–19
5.16 Moving Devices Under the Same HS1CP 5–20
5.17 Command Disks 5–21
5.17.1 Uses for Command Disks 5–21
5.17.2 Creating a Command Disk 5–21
5.17.2.1 Controller Setup 5–22
5.17.2.2 Host Setup 5–22
5.17.3 Communicating with a Command Disk 5–23
5.17.4 Performance 5–24
5.17.5 Maintenance 5–24
5.18 Server Processor Configuration and Setup 5–24
5.18.1 System Firmware 5–24
5.18.2 Switching Between Interfaces 5–25
5.18.3 Verifying Configuration: SRM Console Commands for OpenVMS 5–25
5.18.3.1 The SHOW CONFIG Command 5–25
5.18.3.2 The SHOW DEVICE Command 5–26
5.18.3.3 The SHOW MEMORY Command 5–27
5.18.3.4 Setting and Showing Environment Variables 5–28
5.19 System Bus Options 5–31
5.19.1 CPU Daughterboard 5–33
5.19.2 Memory Modules 5–33
5.20 Motherboard 5–34
5.21 PCI Bus Options 5–34
5.22 EISA Bus Options 5–35
5.23 EISA Configuration Utility 5–35
5.23.1 Before You Run the ECU 5–35
5.23.2 How to Start the ECU 5–36
5.23.3 Configuring EISA Options 5–37
5.23.4 SCSI Buses 5–37
5.24 Console Port Configurations 5–37

6 Error Analysis and Fault Isolation

6.1 Special Considerations for the HS1CP 6–1
6.2 HS1CP Troubleshooting Basics 6–2
6.3 Types of HS1CP Error Reporting 6–3
6.3.1 HS1CP Front Panel Indicators 6–3
6.3.1.1 HS1CP Normal Operation 6–4
6.3.1.2 Fault Notification 6–5
6.3.2 Device LEDs 6–10
6.3.2.1 Storage SBB Status 6–10
6.3.2.2 Device Shelf Status and Power Supply Status 6–12
6.3.3 Error Messages 6–15
6.3.3.1 Conditions under Which CER Messages Appear 6–15
6.3.3.2 Clearing the Terminal of Error Messages 6–15
6.3.3.3 Types of Error Messages 6–16
6.3.4 Interactive CLI Messages 6–24

6.3.5 Last Failure Logging Messages ... 6–24
6.3.6 Event Logging Messages ... 6–25
6.4 Host Error Logs ... 6–26
6.4.1 Translation Utilities ... 6–26
6.4.2 Fault Management Utility ... 6–26
6.4.3 Host Error Log Translation ... 6–27

7 Diagnostics, Exercisers, and Utilities

7.1 Initialization and Self-Test Routines ... 7–1
7.1.1 Built-In Self-Test ... 7–2
7.1.2 Core Module Integrity Self-Test ... 7–2
7.1.3 Module Integrity Self-Test DAEMON ... 7–4
7.2 Disk Inline Exerciser ... 7–7
7.2.1 Invoking DILX ... 7–7
7.2.2 Interrupting DILX Execution ... 7–8
7.2.3 DILX Tests ... 7–8
7.2.3.1 Basic Function Test—DILX ... 7–8
7.2.3.2 User-Defined Test—DILX ... 7–9
7.2.4 DILX Test Definition Questions ... 7–10
7.2.5 DILX Output Messages ... 7–15
7.2.6 DILX End Message Display ... 7–19
7.2.7 DILX Error Information Packet Displays ... 7–20
7.2.8 DILX Data Patterns ... 7–22
7.2.9 DILX Examples ... 7–23
7.2.9.1 DILX Example—Using All Defaults ... 7–23
7.2.9.2 DILX Example—Using All Functions ... 7–24
7.2.9.3 DILX Examples—Auto-Configure with All Units ... 7–27
7.2.10 Interpreting the DILX Performance Summaries ... 7–29
7.2.11 DILX Abort Codes ... 7–31
7.2.12 DILX Error Codes ... 7–31
7.3 Tape Inline Exerciser ... 7–33
7.3.1 Invoking TILX ... 7–33
7.3.2 Interrupting TILX Execution ... 7–34
7.3.3 TILX Tests ... 7–34
7.3.3.1 Basic Function Test—TILX ... 7–34
7.3.3.2 User-Defined Test—TILX ... 7–35
7.3.3.3 Read Only Test—TILX ... 7–35
7.3.4 TILX Test Definition Questions ... 7–36
7.3.5 TILX Output Messages ... 7–40
7.3.6 TILX End Message Display ... 7–44
7.3.7 TILX Error Information Packet Displays ... 7–44
7.3.8 TILX Data Patterns ... 7–46
7.3.9 TILX Examples ... 7–47
7.3.9.1 TILX Example—Using All Defaults ... 7–47
7.3.9.2 TILX Example—Using All Functions ... 7–48
7.3.10 Interpreting the TILX Performance Summaries ... 7–50
7.3.11 TILX Abort Codes ... 7–51
7.3.12 TILX Error Codes ... 7–52
7.4 VTDPY Utility ... 7–53
7.4.1 How to Run VTDPY ... 7–53
7.4.2 Using the VTDPY Control Keys ... 7–53
7.4.3 Using the VTDPY Command Line ... 7–54
7.4.4 How to Interpret the VTDPY Display Fields ... 7–55

7.5 CONFIG Utility 7–71
7.5.1 CONFIG Conventions 7–71
7.5.2 Running the CONFIG Utility 7–71
7.6 Configuration Menu 7–73
7.6.1 Restrictions 7–73
7.6.2 Main Menu 7–73
7.6.3 Adding Devices 7–74
7.6.4 Adding Stripesets 7–76
7.6.5 Adding RAIDsets 7–76
7.6.6 Adding to Sparesets 7–78
7.6.7 Adding Passthrough Containers 7–79
7.6.8 Initializing Containers 7–80
7.6.9 Adding Units 7–81
7.6.10 Terminal Setup 7–82
7.6.11 Messages 7–82
7.6.12 Exiting CFMENU 7–84
7.7 Code Patch Utility 7–85
7.7.1 Invoking the Code Patch Utility 7–85
7.7.2 Code Patching 7–85
7.7.2.1 Code Patch Considerations 7–85
7.7.2.2 Entering a Patch 7–86
7.7.2.3 Listing Patches 7–87
7.7.2.4 Deleting a Patch 7–87
7.7.2.5 Messages 7–88
7.7.2.6 Exiting Code Patch 7–89
7.8 Firmware Licensing System 7–90
7.8.1 Enabling Options 7–90
7.8.2 Disabling Options 7–90
7.8.3 License Key 7–90
7.8.4 Using the Menu 7–90
7.8.5 Example 7–91
7.8.6 Messages 7–92
7.9 Fault Management Utility 7–94
7.9.1 The SET Command 7–94
7.9.2 The SHOW Command 7–97
7.9.3 The DESCRIBE Command 7–98
7.9.4 The EXIT Command 7–99
7.9.5 Examples 7–99
7.10 The CLONE Utility 7–104
7.10.1 Cloning a Single Device Unit 7–106
7.10.2 Cloning a Stripeset 7–108
7.10.3 Cloning a Mirrorset 7–111
7.10.4 Cloning a Striped Mirrorset 7–114

8 Removing and Replacing Field Replaceable Units

8.1 Electrostatic Discharge Protection 8–1
8.2 Accessing FDDI Server Components 8–2
8.3 Connecting a Terminal to the Server Processor 8–4
8.4 Replacing a Server Processor in a Dual Server Processor FDDI Server . . . 8–5
8.4.1 Removing the Internal Bus Cables 8–6
8.4.2 Exchanging the Server Processor 8–8
8.4.3 Reinstalling the Internal Bus Cables 8–9
8.5 Preparing to Replace Server Processor Components 8–11

8.5.1 Required Equipment 8–11
8.5.2 Shutting Down the System 8–11
8.5.3 Accessing the Server Processor 8–12
8.5.4 Server Processor Cover Panels 8–13
8.6 Replacing Internal Server Processor Components Accessible from the Rear 8–14
8.6.1 Removing and Installing SIMMs 8–16
8.6.1.1 Removing SIMMs 8–16
8.6.1.2 Preparing to Install SIMMs 8–18
8.6.1.3 Replacing SIMMs 8–18
8.6.1.4 Verifying Memory Configuration 8–19
8.6.1.5 Booting Your System 8–19
8.6.2 Removing and Replacing the EISA or PCI Option Cards 8–20
8.6.2.1 Preparing to Remove an EISA Device Bus Adapter Card 8–20
8.6.2.2 Removing an EISA or PCI Option Card 8–20
8.6.2.3 Replacing an EISA or PCI Option Card 8–21
8.6.2.4 Testing an EISA or PCI Option Card Installation 8–21
8.6.2.5 Configuring an EISA or PCI Option Card 8–22
8.6.3 Removing and Replacing the SCSI Hard Drive 8–22
8.6.4 Removing and Replacing the Speaker 8–24
8.6.5 Removing and Replacing the CPU Daughterboard 8–25
8.6.6 Removing and Replacing the FDDI Adapter 8–25
8.6.7 Removing and Replacing the Motherboard 8–25
8.7 Replacing Internal Server Processor Components Accessible from the Front 8–28
8.7.1 Removing and Replacing the Fan 8–30
8.7.2 Accessing the Media Tray 8–30
8.7.2.1 Replacing a CD–ROM Drive 8–30
8.7.2.2 Removing and Replacing a Floppy Drive 8–31
8.7.3 Removing and Replacing the OCP Module 8–32
8.7.4 Removing and Replacing the Power Supply 8–32
8.8 Reassembling the Server Processor 8–34
8.9 Reinstalling the Server Processor 8–35
8.10 Replacing the Cable Distribution Unit 8–36
8.11 Replacing FDDI Server Shelf Components 8–38
8.11.1 Warm Swap Procedure 8–38
8.11.1.1 Disk Drive SBB Warm Swap Removal 8–38
8.11.1.2 Disk Drive SBB Warm Swap Replacement 8–40
8.11.1.3 Tape Drive SBB Warm Swap Removal and Replacement 8–41
8.11.1.4 When the Quiesce Bus State Is Not Displayed 8–41
8.11.2 Removing and Replacing the BA350–M Shelf 8–42
8.11.3 Removing and Replacing a BA350–S Shelf 8–43
8.11.4 Replacing a Shelf Blower 8–44
8.11.5 Replacing a Shelf Power Supply 8–45
8.12 Removing and Replacing the HS1CP and HS1CP Components 8–46
8.12.1 Reconciling Firmware in Different Version HS1CPs 8–48
8.12.2 Using the Warm Swap Utility to Replace an HS1CP and HS1CP Components 8–48
8.12.2.1 When to Use C_SWAP 8–49
8.12.2.2 Functions of C_SWAP 8–49
8.12.3 Removing an HS1CP or Cache Module During a C_SWAP Operation 8–49
8.12.4 Replacing an HS1CP and/or Cache Module During a C_SWAP Operation 8–51

8.12.5 Removing an HS1CP 8–53
8.12.6 Removing a Cache Memory Module 8–54
8.12.7 Replacing an HS1CP 8–55
8.12.8 Replacing a Cache Module 8–56
8.12.9 Removing, Disposing of, and Replacing a Write-Back Cache Battery 8–57
8.12.9.1 Removing a Write-Back Cache Battery 8–57
8.12.9.2 Disposing of a Write-Back Cache Battery 8–57
8.12.9.3 Replacing a Write-Back Cache Battery 8–58
8.12.10 Verifying Firmware Versions and Write-Back Cache Size 8–58
8.12.11 Saving the Storage Configuration 8–60
8.12.12 Verifying Installation Procedures: Creating Redundant Data Paths 8–61
8.13 Removing and Replacing Fans 8–63
8.13.1 Removing the Cabinet Cover Fan Assembly 8–64
8.13.2 Installing/Replacing the Cover Fan 8–67
8.14 Replacing Cabinet Cables and Cords 8–67
8.14.1 Replacing a SCSI Cable 8–67
8.14.2 Replacing a DSSI Cable 8–69
8.14.2.1 Hot Swap Procedure for a Dual Server Processor FDDI Server 8–69
8.14.2.2 Standard Procedure 8–69
8.14.3 Replacing the Power Cord 8–71

A Functional Description

A.1 Cable Distribution Unit (CDU) A–3
A.2 Storage Devices A–3
A.3 Server Processor Components Accessible from the Front A–3
A.3.1 Operator Control Panel Controls and Indicators A–4
A.3.2 Media Tray A–5
A.3.3 Power Supplies A–6
A.3.4 Fan Assembly A–6
A.4 Server Processor Components Accessible from the Rear A–6
A.4.1 Rear Panel Ports and Slots A–8
A.4.2 Motherboard A–8
A.4.3 Daughterboard A–9
A.4.4 Memory Modules A–9
A.4.5 System Hard Disk Drive A–10
A.4.6 HS1AD Device Bus Adapter A–10
A.4.7 FDDI Adapter (DEFEA or DEFPA Module) A–11
A.5 The HS1CP A–11
A.5.1 Maintenance Terminal Port A–12
A.5.2 Program Card A–12
A.5.3 Policy Processor A–12
A.5.4 Diagnostic Registers A–12
A.5.5 Front Panel A–12
A.5.6 Dual Controller Port A–14
A.5.7 Nonvolatile Memory (NVMEM) A–14
A.5.8 Bus Exchangers A–15
A.5.9 Shared Memory A–15
A.5.10 Value-Added Functions A–15
A.5.11 Device Ports A–15
A.5.12 Write-Back Cache Module A–15
A.5.13 Host Interface A–17

A.5.14 Firmware ... A–19
A.5.14.1 Core Functions ... A–19
A.5.14.2 Host Interconnect Functions ... A–20
A.5.14.3 Operator Interface and Subsystem Management Functions ... A–20
A.5.14.4 Device Services Functions ... A–21
A.5.14.5 Value-Added Functions ... A–22
A.5.14.6 Addressing Storage Within the Subsystem ... A–22

B Field Replaceable Unit Part Numbers

B.1 Required Tools ... B–1
B.2 Field Replaceable Units ... B–2
B.3 Server Processor Cabling Interconnection ... B–4

C HS1CP Error Logging

C.1 Reading an HS1CP Error Log (ERF) ... C–1
C.2 Reading an HS1CP Error Log (DECevent) ... C–6
C.3 Event Log Formats ... C–11
C.3.1 Common Event Log Fields ... C–14
C.3.1.1 CI Host Interconnect Services Common Event Log Fields ... C–14
C.3.1.2 Host/Server Connection Common Fields ... C–15
C.3.1.3 Byte Count/Logical Block Number Common Fields ... C–16
C.3.1.4 Device Location/Identification Common Fields ... C–17
C.3.1.5 SCSI Device Sense Data Common Fields ... C–20
C.3.2 Specific Event Log Formats ... C–28
C.3.2.1 Last Failure Event Log (Template 01) ... C–28
C.3.2.2 Failover Event Log (Template 05) ... C–32
C.3.2.3 Firmware License Violation Event Log (Template 06) ... C–34
C.3.2.4 Nonvolatile Parameter Memory Component Event Log (Template 11) ... C–36
C.3.2.5 Backup Battery Failure Event Log (Template 12) ... C–38
C.3.2.6 Subsystem Built-In Self Test Failure Event Log (Template 13) ... C–40
C.3.2.7 Memory System Failure Event Log (Template 14) ... C–42
C.3.2.8 Internal Bus Port/Port Driver Event Log (Template 32) ... C–45
C.3.2.9 Internal Bus System Communication Services Event Log (Template 33) ... C–47
C.3.2.10 Device Services Nontransfer Error Event Log (Template 41) ... C–50
C.3.2.11 Disk Transfer Error Event Log (Template 51) ... C–53
C.3.2.12 Disk Bad Block Replacement Attempt Event Log (Template 57) ... C–55
C.3.2.13 Tape Transfer Error Event Log (Template 61) ... C–58
C.3.2.14 Media Loader Error Event Log (Template 71) ... C–60
C.3.2.15 Disk Copy Data Correlation Event Log ... C–63
C.4 Event Log Codes ... C–63
C.5 Event Notification/Recovery Threshold ... C–139
C.6 Recommended Repair Action ... C–140
C.7 Deskew Command Procedure ... C–146

D Configuration Rules and Restrictions

D.1 Standard FDDI Server Configurations . . . D–1
D.2 Electrostatic Discharge Protection . . . D–2
D.3 Connecting a Maintenance Terminal . . . D–3
D.4 Configuring the SW800 Cabinet . . . D–4
D.4.1 General Configuration Rules for FDDI Server Cabinets . . . D–4
D.4.2 Device Placement . . . D–9
D.4.3 Shelf Arrangement . . . D–10
D.4.4 Typical and Recommended 3.5-Inch and 5.25-Inch SBB Configurations . . . D–10
D.4.4.1 3.5-Inch SBB Restrictions . . . D–11
D.4.4.2 5.25-Inch SBB Restrictions . . . D–12
D.4.4.3 5.25-Inch SBB Recommended Configurations . . . D–13
D.4.4.4 Intermixing 5.25-Inch and 3.5-Inch SBBs . . . D–13
D.4.4.5 Atypical Configurations . . . D–14
D.5 Configuring Cabinet Power . . . D–15
D.5.1 Configuring for No Power Redundancy . . . D–15
D.5.2 Partial Power Redundancy Configuration . . . D–16
D.5.3 Full Power Redundancy Configuration . . . D–18
D.6 Configuring Internal Bus Cables for the HS1CP and Server Processor . . . D–20
D.6.1 Configuring HS111/211 Internal Bus Cable and Terminator . . . D–21
D.6.2 Configuring HS121/221 Internal Bus Cables . . . D–21
D.6.3 Configuring HS241 Internal Bus Cables . . . D–21
D.6.4 Configuring HS241/280 Internal Bus Cables . . . D–22
D.7 Configuring the HS1CP . . . D–24
D.7.1 HS1CP Configuration Guidelines and Considerations . . . D–27
D.7.2 Preconfigured and CTO Subsystems . . . D–30
D.7.3 Preset Configuration Parameters . . . D–31
D.8 Configuring the Server Processor . . . D–32
D.8.1 The HS1AD Adapter . . . D–35
D.8.2 System Disk SCSI Identifier . . . D–36
D.8.3 CD–ROM Drive SCSI Identifiers . . . D–39
D.9 Running the EISA Configuration Utility (ECU) . . . D–41
D.9.1 How to Run the ECU . . . D–41
D.9.2 Verifying Your Installation . . . D–43
D.10 Upgrading Existing FDDI Servers . . . D–44

E Command Line Interpreter

E.1 CLI Commands . . . E–1
ADD CDROM . . . E–2
ADD DISK . . . E–3
ADD LOADER . . . E–5
ADD MIRRORSET . . . E–6
ADD OPTICAL . . . E–8
ADD PASSTHROUGH . . . E–10
ADD RAIDSET . . . E–11
ADD SPARESET . . . E–13
ADD STRIPESET . . . E–14
ADD TAPE . . . E–15
ADD UNIT . . . E–16
CLEAR_ERRORS CLI . . . E–25

CLEAR_ERRORS INVALID_CACHE E–26
CLEAR_ERRORS LOST_DATA E–27
CLEAR_ERRORS UNKNOWN E–28
CLEAR_ERRORS UNWRITEABLE_DATA E–29
DELETE *container-name* E–30
DELETE FAILEDSET E–31
DELETE SPARESET E–32
DELETE *unit-number* E–33
DIRECTORY E–34
EXIT E–35
HELP E–36
INITIALIZE E–37
LOCATE E–39
MIRROR *disk-device-name1 container-name* E–41
REDUCE *disk-device-name1 [disk-device-nameN]* E–43
RENAME E–45
RESTART OTHER_CONTROLLER E–46
RESTART THIS_CONTROLLER E–48
RETRY_ERRORS UNWRITEABLE_DATA E–50
RUN E–51
SELFTEST OTHER_CONTROLLER E–52
SELFTEST THIS_CONTROLLER E–54
SET *disk-container-name* E–56
SET FAILOVER E–57
SET *mirrorset-container-name* E–59
SET NOFAILOVER E–62
SET OTHER_CONTROLLER E–63
SET *RAIDset-container-name* E–66
SET THIS_CONTROLLER E–68
SET *unit-number* E–70
SHOW CDROMS E–79
SHOW *cdrom-container-name* E–80
SHOW DEVICES E–81
SHOW DISKS E–82
SHOW *disk-container-name* E–83
SHOW FAILEDSET E–84
SHOW LOADERS E–85
SHOW *loader-container-name* E–86
SHOW MIRRORSETS E–87
SHOW *mirrorset-container-name* E–89
SHOW OPTICALS E–90
SHOW *optical-container-name* E–91
SHOW OTHER_CONTROLLER E–92
SHOW PASSTHROUGH E–93
SHOW *passthrough-container-name* E–94
SHOW RAIDSETS E–95
SHOW *raidset-container-name* E–97

SHOW SPARESET E–98
SHOW STORAGESETS E–99
SHOW STRIPESETS E–101
SHOW *stripeset-container-name* E–103
SHOW TAPES E–104
SHOW *tape-container-name* E–105
SHOW THIS_CONTROLLER E–106
SHOW UNITS E–107
SHOW *unit-number* E–108
SHUTDOWN OTHER_CONTROLLER E–109
SHUTDOWN THIS_CONTROLLER E–111
UNMIRROR *disk-device-name* E–113
E.2 CLI Messages E–114
E.2.1 Error Conventions E–114
E.2.2 CLI Error Messages E–114
E.2.3 Warning Conventions E–134
E.2.4 CLI Warning Messages E–134
E.3 Device Configuration Examples E–137

F SWEAT

F.1 Overview F–1
F.2 Features F–1
F.3 Installing SWEAT F–1
F.4 Commands and Syntax F–3
SWEAT F–4
F.5 Error Log Collation Techniques—Overview F–10
F.5.1 Error Log Collation F–10
F.5.2 Determining the Time Span of the Symptoms F–11
F.5.3 Collating Error Logs to One File F–12
F.5.4 Filtering the Error Log F–13
F.5.5 Using SWEAT to Break Out the Error Log F–13

Glossary

Index

Examples

7–1 DILX End Message Display 7–19
7–2 Controller Error 7–20
7–3 Memory Error 7–21
7–4 Disk Transfer Error 7–21
7–5 Bad Block Replacement Attempt Error 7–22
7–6 Using All Defaults—DILX 7–24
7–7 All Functions—DILX 7–25
7–8 Auto-Configuration with All Units 7–27
7–9 Auto-Configuration with Half of All Units 7–28
7–10 TILX End Message Display 7–44

7–11 Controller Error 7–45
7–12 Memory Error 7–45
7–13 Tape Error 7–45
7–14 Using All Defaults—TILX 7–48
7–15 Using All Functions—TILX 7–48
7–16 DEFAULT VTDPY DISPLAY 7–56
7–17 VTDPY DISPLAY STATUS 7–57
7–18 VTDPY DISPLAY DEVICE 7–58
8–1 SHOW MEMORY Command 8–19
8–2 The SHOW CONFIG Command 8–22
C–1 Disk Transfer Error Event Log C–2
C–2 ERF-Generated Error Event Log C–7
C–3 DECevent-Generated Error Event Log C–9
C–4 Deskew Command Procedure Example C–146
C–5 ERF Error Log Before Command Procedure C–148
C–6 ERF Error Log After Command Procedure C–149
E–1 Initial Single Controller Configuration of an HS1CP Controller E–137
E–2 Setting the Terminal Speed and Parity E–137
E–3 Creating a Unit from a Disk Device E–138
E–4 Creating a Unit from a Tape Device E–138
E–5 Creating a Unit from a Four-Member Stripeset E–138
E–6 Creating a Unit from a Five-Member RAIDset E–138
E–7 Creating a Unit From a Disk Device and Setting the Write Protection E–139
E–8 Setting the Write Protection for an Existing Unit E–139
E–9 Renumbering Disk Unit 0 to Disk Unit 100 . . . . . E–139
E–10 Creating a Transportable Unit from a Disk Device E–139
E–11 Changing the Replacement Policy of a RAIDset E–139
E–12 Deleting the Unit, Stripeset, and All Disks Associated with a Stripeset E–140

Figures

2–1 Jumper J1 on the CPU Daughterboard 2–4
2–2 Server Processor Motherboard Layout 2–5
2–3 Floppy Drive Activity LED 2–6
2–4 CD–ROM Drive Activity LED 2–7
2–5 Operator Control Panel 2–12
3–1 HS1CP Layout 3–1
3–2 3.5-Inch Storage SBB LEDs 3–4
3–3 5.25-Inch Storage SBB LEDs 3–4
3–4 Power Supply LEDs 3–6
5–1 HS241 Diagram 5–3
5–2 HS241 with HS280 Diagram 5–4
5–3 HS221/121 Diagram 5–5
5–4 HS211/111 Diagram 5–6
5–5 Device Name Convention 5–27

5–6	Card Slots and Bus Locations on the Motherboard	5–32
6–1	HS1CP Layout	6–4
6–2	Solid Front Panel Codes	6–6
6–3	Flashing Front Panel Codes	6–7
6–4	3.5-Inch Storage SBB LEDs	6–11
6–5	5.25-Inch Storage SBB LEDs	6–11
6–6	Power Supply LEDs	6–13
7–1	Controller Initialization	7–2
7–2	CFMENU Main Menu	7–74
7–3	CFMENU Device Menu (Before Adding Devices)	7–74
7–4	CFMENU Device Menu	7–75
7–5	CFMENU Stripeset Menu	7–76
7–6	CFMENU RAIDset Menu	7–77
7–7	CFMENU Spareset/Failedset Menu	7–78
7–8	CFMENU Passthrough Menu	7–79
7–9	CFMENU Initialization Menu	7–80
7–10	CFMENU Unit Menu	7–81
8–1	Exterior Cabinet Panel Removal and Installation	8–4
8–2	Back View of the Server Processor	8–5
8–3	Removing the Cover Panels	8–14
8–4	Server Processor Components Accessible from the Rear	8–15
8–5	Server Processor Motherboard (for HS1*xx*)	8–17
8–6	Removing SIMMs from the Motherboard	8–18
8–7	Installing a SIMM	8–19
8–8	System Disk Location	8–23
8–9	System Disk Cabling	8–23
8–10	System Disk Option Connector and SCSI ID 0	8–24
8–11	Motherboard Jumper Locations	8–27
8–12	Server Processor Components Accessible from the Front	8–29
8–13	Replacing a CD–ROM Drive	8–31
8–14	Power Supply Locations	8–32
8–15	Power Supply Cable Connections	8–33
8–16	Slide Insert Removal/Installation	8–34
8–17	Server Processor Installation	8–35
8–18	Shipping Brackets	8–36
8–19	SBB Warm Swap	8–39
8–20	FDDI Server Component Locations	8–42
8–21	Replacing Shelf Blowers	8–45
8–22	Removing a Power Supply	8–46
8–23	HS1CP and Cache Memory Boards	8–47
8–24	HS1CP Installation	8–54
8–25	HS1CP Shelf Rails	8–56
8–26	Write-Back Cache Batteries	8–58
8–27	FDDI Server Cabinet Component Locations	8–65
8–28	CDU Safety Bracket	8–66
8–29	BA350–M Shelf Layout (Front View)	8–68
8–30	Internal Bus Cable Connectors	8–69

8–31	Trilink Connector	8–70
8–32	Trilink Connector with Cable and Terminator	8–71
A–1	StorageWorks FDDI Server–Major Components	A–2
A–2	CDU Power Cord Connectors	A–3
A–3	Server Processor Front View—Exploded	A–4
A–4	Operator Control Panel	A–5
A–5	Server Processor Rear View—Exploded	A–7
A–6	Motherboard Layout	A–9
A–7	SCSI System Disk and Jumper Locations	A–10
A–8	HS1CP Front Panel	A–13
A–9	HS1CP Device Channel Processor Reset Button and LEDs	A–14
A–10	Common Hardware Block Diagram	A–17
A–11	DSSI Host Interface Hardware Block Diagram	A–18
A–12	SCSI–2 Host Interface Hardware Block Diagram	A–19
A–13	Controller Storage Addressing	A–23
B–1	Cable Connection Diagram	B–5
B–2	Cable Routing Diagram	B–6
C–1	Implementation Dependent Information Format	C–12
C–2	Instance Code Format	C–12
C–3	CI Host Interconnect Services Common Event Log Fields	C–14
C–4	Host/Server Connection Common Fields	C–16
C–5	Byte Count/Logical Block Number Common Fields	C–17
C–6	Device Location/Identification Common Fields	C–18
C–7	Device Locator Field Format	C–18
C–8	SCSI Device Sense Data Common Fields	C–20
C–9	Sense Data Qualifier Field Format	C–21
C–10	SCSI Sense Data Byte Zero ("ercdval") Field Format	C–22
C–11	SCSI Sense Data Byte Two ("snsflgs") Field Format	C–23
C–12	SCSI Sense Data Byte 0F through 11 ("keyspec") Field—Field Pointer Bytes Format	C–26
C–13	SCSI Sense Data Byte 0F through 11 ("keyspec") Field—Actual Retry Count Bytes Format	C–27
C–14	SCSI Sense Data Byte 0F through 11 ("keyspec") Field—Progress Indication Bytes Format	C–28
C–15	Last Failure Event Log (Template 01) Format	C–29
C–16	Last Failure Code Format	C–31
C–17	Failover Event Log (Template 05) Format	C–33
C–18	Firmware License Violation Event Log (Template 06) Format	C–35
C–19	Nonvolatile Parameter Memory Component Event Log (Template 11) Format	C–37
C–20	Backup Battery Failure Event Log (Template 12) Format	C–39
C–21	Subsystem Built-In Self Test Failure Event Log (Template 13) Format	C–41
C–22	Memory System Failure Event Log (Template 14) Format	C–43
C–23	Internal Bus Port/Port Driver Event Log (Template 32) Format	C–46
C–24	Internal Bus System Communication Services Event Log (Template 33) Format	C–48

C–25 Device Services Nontransfer Error Event Log (Template 41) Format C–51
C–26 Disk Transfer Error Event Log (Template 51) Format C–54
C–27 Disk Bad Block Replacement Attempt Event Log (Template 57) Format C–56
C–28 Tape Transfer Error Event Log (Template 61) Format C–59
C–29 Media Loader Error Event Log (Template 71) Format C–61
D–1 Connecting a Maintenance Terminal to the Server Processor D–3
D–2 HS121 Storage Configuration with Optimal Disk Storage and No TZ8*xx*–Series Tape Drives Installed D–5
D–3 Single Extension from Device Shelf to Device Shelf D–6
D–4 HS121 Storage Configuration with Four TZ8*xx*–Series Tape Drives Installed D–7
D–5 HS121 Storage Configuration with Two TZ8*xx*–Series Tape Drives Installed D–8
D–6 Adjacent Devices on a Single Port D–9
D–7 Single-Cabinet Power Configuration D–16
D–8 Dual Shelf AC Power Supply Configuration D–17
D–9 Safety Bracket Rotation D–18
D–10 Dual Cabinet Power Configuration D–19
D–11 Trilink Connector D–20
D–12 Cabling Diagram for HS121/221 ..... D–21
D–13 Cabling Diagram for HS241 ..... D–22
D–14 Cabling Diagram for HS241/280 ..... D–23
D–15 BA350–M Shelf Layout (Front View) D–26
D–16 HS1CP Installation D–28
D–17 Optimal Availability Configuration Example D–30
D–18 Layout of the Front Section of the Server Processor D–33
D–19 Layout of the Back Section of the Server Processor D–34
D–20 Server Processor–Rear View D–35
D–21 Device Bus Adapter Location D–36
D–22 System Disk Cabling D–37
D–23 System Disk Option Connector and SCSI ID 1 ..... D–38
D–24 System Disk Option Connector and SCSI ID 0 ..... D–38
D–25 CD–ROM Drive Option Connector and SCSI ID 4 ..... D–40

Tables

1 StorageWorks Related Documentation xxii
2 Acoustics—Preliminary Declared Values per ISO 9296 and ISO 7779 xxvi
3 Schallemissionswerte—Vorläufige Werteangaben nach ISO 9296 und ISO 7779/DIN EN27779 xxvi
1–1 Diagnostic Flow for Power Problems 1–2
1–2 Diagnostic Flow for Problems Getting to Console Mode 1–2
1–3 Diagnostic Flow for Problems Reported by the Console Program 1–3
1–4 Diagnostic Flow for Boot Problems 1–4
1–5 Diagnostic Flow for Errors Reported by the Operating System 1–4

2–1	Interpreting Error Beep Codes	2–1
2–2	SROM Memory Tests, CPU Jumper J1	2–3
2–3	EISA Troubleshooting	2–7
2–4	OCP Controls and Indicators Description	2–12
3–1	Storage SBB Status LEDs	3–5
3–2	Shelf and Single Power Supply Status LEDs	3–6
3–3	Shelf and Dual Power Supply Status LEDs	3–7
4–1	Summary of Diagnostic and Related Commands	4–2
5–1	Standard HS241/HS280 HS1CP Device Channel Processor Identification	5–11
5–2	Transportable and Nontransportable Devices	5–20
5–3	Environment Variables Set During System Configuration	5–29
5–4	Summary of Procedure for Configuring an EISA Bus Option	5–37
6–1	Storage SBB Status LEDs	6–12
6–2	Shelf and Single Power Supply Status LEDs	6–13
6–3	Shelf and Dual Power Supply Status LEDs	6–14
7–1	Cache Module Testing	7–5
7–2	DILX Data Patterns	7–23
7–3	DILX Abort Codes and Definitions	7–31
7–4	DILX Error Codes and Definitions	7–32
7–5	TILX Data Pattern Definitions	7–47
7–6	TILX Abort Codes and Definitions	7–52
7–7	TILX Abort Codes and Definitions	7–52
7–8	VTDPY Control Keys	7–53
7–9	VTDPY Commands	7–54
7–10	Thread Description	7–63
8–1	Server Processor Components	8–13
8–2	Jumpers on the Motherboard	8–28
8–3	HS1CP Removal	8–53
8–4	Cache Module Removal	8–55
8–5	HS1CP Replacement	8–55
8–6	Cache Module Replacement	8–57
A–1	OCP Controls and Indicators Description	A–5
A–2	Power Supply Ratings	A–6
A–3	Rear Panel Ports and Slots	A–8
A–4	FDDI Adapter LED States	A–11
B–1	FDDI Server Replaceable Parts	B–2
B–2	CDU Parts List	B–2
B–3	BA350–S Shelf Assembly	B–3
B–4	BA350–M Shelf Assembly	B–3
B–5	Server Processor Replaceable Parts	B–4
C–1	Template Types	C–5
C–2	Firmware Component Identifier Codes	C–63
C–3	Host Interconnect Services Status Codes	C–64
C–4	CI Message Operation Codes	C–66
C–5	CI Virtual Circuit State Codes	C–66
C–6	Port/Port Driver Message Operation Codes	C–66

C–7	System Communication Services Message Operation Codes	C–67
C–8	CI Connection State Codes	C–67
C–9	Supported SCSI Device Type Codes	C–67
C–10	SCSI Command Operation Codes	C–68
C–11	SCSI Buffered Modes Codes	C–70
C–12	SCSI Sense Key Codes	C–70
C–13	SCSI ASC/ASCQ Codes for Direct-Access Devices (Such as Magnetic Disk)	C–71
C–14	SCSI ASC/ASCQ Codes for Sequential-Access Devices (Such as Magnetic Tape)	C–75
C–15	SCSI ASC/ASCQ Codes for CD-ROM Devices	C–78
C–16	SCSI ASC/ASCQ Codes for Medium Changer Devices (Such as Jukeboxes)	C–81
C–17	HS1CP Controller Vendor Specific SCSI ASC/ASCQ Codes	C–83
C–18	Last Failure Event Log (Template 01) Instance/MSCP Event Codes	C–85
C–19	Failover Event Log (Template 05) Instance/MSCP Event Codes	C–85
C–20	Firmware License Violation Event Log (Template 06) Instance/MSCP Event Codes	C–86
C–21	Nonvolatile Parameter Memory Component Event Log (Template 11) Instance/MSCP Event Codes	C–86
C–22	Backup Battery Failure Event Log (Template 12) Instance/MSCP Event Codes	C–87
C–23	Subsystem Built-In Self Test Failure Event Log (Template 13) Instance/MSCP Event Codes	C–87
C–24	Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes	C–88
C–25	Internal Bus Port/Port Driver Event Log (Template 32) Instance/MSCP Event Codes	C–99
C–26	Internal Bus System Communication Services Event Log (Template 33) Instance/MSCP Event Codes	C–100
C–27	Device Services Nontransfer Error Event Log (Template 41) Instance/MSCP Event Codes	C–102
C–28	Disk Transfer Error Event Log (Template 51) Instance/MSCP Event Codes	C–107
C–29	Disk Bad Block Replacement Attempt Event Log (Template 57) Instance/MSCP Event Codes	C–111
C–30	Tape Transfer Error Event Log (Template 61) Instance/MSCP Event Codes	C–111
C–31	Media Loader Error Event Log (Template 71) Instance/MSCP Event Codes	C–112
C–32	Disk Copy Data Correlation Event Log "Event Dependent Information" Values	C–113
C–33	Executive Services Last Failure Codes	C–113
C–34	Value-Added Services Last Failure Codes	C–115
C–35	Device Services Last Failure Codes	C–120
C–36	Fault Manager Last Failure Codes	C–126
C–37	Common Library Routines Last Failure Codes	C–128

C–38	Dual Universal Asynchronous Receiver/Transmitter Services Last Failure Codes	C–128
C–39	Failover Control Last Failure Codes	C–128
C–40	Nonvolatile Parameter Memory Failover Control Last Failure Codes	C–129
C–41	Command Line Interpreter Last Failure Codes	C–130
C–42	Host Interconnect Services Last Failure Codes	C–131
C–43	Host Interconnect Port Services Last Failure Codes	C–132
C–44	Disk and Tape MSCP Server Last Failure Codes	C–133
C–45	Diagnostics and Utilities Protocol Server Last Failure Codes	C–136
C–46	System Communication Services Directory Service Last Failure Codes	C–136
C–47	Disk Inline Exerciser (DILX) Last Failure Codes	C–136
C–48	Tape Inline Exerciser (TILX) Last Failure Codes	C–137
C–49	Device Configuration Utilities (CONFIG/CFMENU) Last Failure Codes	C–138
C–50	Clone Unit Utility (CLONE) Last Failure Codes	C–138
C–51	Controller Restart Codes	C–138
C–52	Event Notification/Recovery Threshold Classifications	C–139
C–53	Recommended Repair Action Codes	C–140
D–1	3.5-Inch SBB Configurations, 6-Port Controllers	D–11
D–2	5.25-Inch SBB Configurations for the HS1CP	D–13
D–3	Small Shelf Count Configurations, 6-Port Controllers	D–14
D–4	HS Operating Firmware Supported Devices	D–25
D–5	High-Performance Devices per Port	D–29
D–6	StorageWorks FDDI Server Migration Path Upgrade Matrix	D–44
F–1	ELE_TYPE Values (in Decimal)	F–13

Preface

Intended Audience

This manual is intended for use by Digital™ Multivendor Customer Services personnel or other qualified technicians who service the StorageWorks™ FDDI Servers.

This audience must be knowledgable in the installation and configuration of StorageWorks shelves, devices, and related components.

Structure

This guide is organized as follows:

Chapter 1	Provides an overview of troubleshooting strategies.
Chapter 2	Provides information about power-up diagnostics for the server processor.
Chapter 3	Provides information about power-up diagnostics for the device channel processor.
Chapter 4	Provides information about running system diagnostics.
Chapter 5	Provides information about the normal operation of the StorageWorks FDDI Server.
Chapter 6	Provides information about error analysis and fault isolation.
Chapter 7	Provides information about diagnostics, exercisers, and utilities.
Chapter 8	Provides information about removing and replacing field replacable units.
Appendix A	Provides a functional description of the FDDI Server.
Appendix B	Provides a list of field replaceable unit part numbers.
Appendix C	Provides examples of error logs.
Appendix D	Provides configuration rules and restrictions.
Appendix E	Provides details for CLI commands and their usage.

Related Documents

Table 1 lists the StorageWorks-related user documents organized by use, system, or product.

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 Array Controllers HS Family of Array Controllers User's Guide	EK–HSFAM–UG
StorageWorks™ Solutions Array Controllers HS Family of Array Controllers Service Manual	EK–HSFAM–SM
Storageworks FDDI Server Publications	
StorageWorks™ Solutions StorageWorks FDDI Server Installation Guide	EK–FSERV–IG
StorageWorks™ Solutions StorageWorks FDDI Server HS110 Upgrade Installation Guide	EK–HS110–IG
StorageWorks™ Solutions StorageWorks FDDI Server Upgrade Installation Guide	EK–FSUPG–IG
StorageWorks™ Solutions StorageWorks FDDI Server HS1AD Device Bus Adapter Installation Guide	EK–HS1AD–IG
StorageWorks™ Solutions StorageWorks FDDI Server HS1MM/HS2MM Server Processor Memory Installation Guide	EK–HS1MM–IG
StorageWorks™ Solutions StorageWorks FDDI Server HS1CP ™ Device Channel Processor Installation Guide	EK–HS1CP–IG
StorageWorks™ Solutions StorageWorks FDDI Server HS1PW Server Processor Power Supply Installation Guide	EK–HS1PW–IG
StorageWorks™ Solutions StorageWorks FDDI Server, Server Processor Component Manual	EK–HSASP–CM
StorageWorks Enclosures	
StorageWorks™ Solutions SW500 and SW800 Cabinet Metric Shelf Bracket Kit Installation Guide	EK–35XRD–IG
StorageWorks™ Solutions RETMA Shelf Rail Kit Installation Guide	EK–35XRB–IG
StorageWorks™ Solutions SW800-Series Data CDU Installation Guide	EK–SW8XP–BA
StorageWorks™ Solutions SW800-Series Data Center Cabinet Installation and User's Guide	EK–SW800–IG
Storage Devices	
StorageWorks™ Solutions 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

(continued on next page)

Table 1 (Cont.) StorageWorks Related Documentation

General Reference Publications	
Digital Systems and Options Catalog	†
Small Computer System Interface, An Overview	EK–SCSIS–OV
VAXcluster Console System User's Guide	†
VCS Software Manual	†
Guide to VMS Performance Management	†
Introduction to Performance Management	†
Introduction to VMS System Management	†
Guide to Setting Up a VMS System	†
Guide to VMS File Applications	†
VMS DCL Dictionary	†
DEC PS Guide	†
Error Reference & Utility Publications	
VMS Error Log Utility Reference Manual	†
DECevent Translation and Reporting Utility for OpenVMS User and Reference Guide	†
VMS System Generation Utility Manual	†
VMS Authorize Utility Manual	†
VMS Monitor Utility Manual	†

† Available from your Digital account representative.

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.
Ctrl/*x*	Ctrl/*x* indicates that you hold down the Ctrl key while you press another key, indicated by *x*. For DILX or TILX, the caret symbol (^) is equivalent to the Ctrl key and these same instructions apply.
italic type	Italic type indicates emphasis and complete manual titles. In the glossary, italic type also is used to indicate cross-references.
UPPERCASE	Words in uppercase text indicate a command, the name of a file, or an abbreviation for a system privilege.

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 corrective 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 guide 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 2 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**
HS2xx with only cabinet fans operating	7.6	7.6	59	59
HS2xx with 2 BA350–M shelves and 12 BA350–S shelves, each containing 6 RZ26–VA disk drives	7.6	7.6	59	59
Per device when installed in an HS2xx				
BA350–S shelf containing 6 RZ26–VA disk drives	5.7	5.7	39	39
BA350–M shelf	5.6	5.6	39	39

† Current values for specific configurations are available from Digital representatives.

‡ 1 B = 10 dBA.

Note

Table 3 contains the specifications in Table 2 translated into the German language.

Table 3 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**
HS2xx nur mit kabinett Lüftern in Betrieb	7,6	7,6	59	59
HS2xx mit 2 BA350–M shelves und 12 BA350–S shelves, jedes bestückt mit 6 RZ26–VA disk drives	7,6	7,6	59	59
Pro Gerät installiert im HS2xx				
BA350–S shelf mit 6 RZ26–VA disk drives	5,7	5,7	39	39
BA350–M shelf	5,6	5,6	39	39

† Aktuelle Werte für spezielle Ausrüstungsstufen sind über die Digital Equipment Vertretungen erhältlich.

‡ 1 B = 10 dBA.

1

Troubleshooting Strategy

This chapter describes the general troubleshooting processes and strategies for locating problems on the FDDI Server system. The following topics are discussed:

- Troubleshooting the system
- Maintenance features
- Common problems
- Troubleshooting basics
- Service tools and utilities
- Information services

1.1 Troubleshooting the System

Before troubleshooting any system problem, check the site maintenance log for the system's service history. Be sure to ask the system manager the following questions:

- Has the system been used before, and did it work correctly?
- Have other symptoms or minor problems been observed or reported?
- Have changes to hardware or updates to firmware or software been made to the system recently?
- What is the state of the system—is the operating system running?

  If the operating system is down and you are not able to reboot it, use the console environment diagnostic tools, such as the power-up display and ROM-based diagnostics (RBDs).

  If the operating system is running, use the operating system environment diagnostic tools, such as error logs, crash dumps, and exercisers (DEC VET).

Problem Categories

System problems can be classified into the following five categories. Use these categories to determine a starting point for diagnosis and eliminate the unlikely sources of the problem:

- Power problems (Table 1–1)
- No access to console mode (Table 1–2)
- Console-reported failures (Table 1–3)

- Boot failures (Table 1–4)
- Operating system-reported failures (Table 1–5)

Table 1–1 Diagnostic Flow for Power Problems

If...	Then...
The system does not power on	Do the following: • Check the power source circuit breakers and power cord connections. • Check the condition of all power cords. • If there are two power supplies, make sure both power supplies are plugged in. • Check the On/Off switch setting on the operator control panel. • Check that the ambient room temperature is within environmental specifications (10–40°C, 50–104°F). • Check that internal power supply cables are plugged in at both the power supply and system component.

Table 1–2 Diagnostic Flow for Problems Getting to Console Mode

If...	Then...
The power-up screen is not displayed	Do the following: • Interpret the error beep codes at power-up for a failure detected during self-tests. • Check that the keyboard and monitor are connected properly and turned on.

(continued on next page)

Table 1–2 (Cont.) Diagnostic Flow for Problems Getting to Console Mode

If...	Then...
If the power-up screen is not displayed, yet the system enters console mode when you press the Return key	Check that the console environment variable is set correctly, as follows: • If you are using a VGA console terminal, the console variable should be set to "graphics." • If you are using a serial console terminal, the console variable should be set to "serial." If the console is set to serial, the power-up screen is routed to the COM1 serial communication port and cannot be viewed from the VGA monitor. • Try connecting a console terminal to the COM1 serial communication port. If necessary, use an MMJ-to-9-pin adapter (H8571-J). Check the baud rate setting for the console terminal and the system. The system baud rate setting is 9600. When using the COM1 port, you must set the console environment variable to "serial." • For certain situations, power up using the fail-safe loader to load new console firmware from a diskette.

Table 1–3 Diagnostic Flow for Problems Reported by the Console Program

If...	Then...
The power-up tests do not complete	Interpret the error beep codes at power-up and check the power-up screen for a failure detected during self-tests.
The power-up display stops on e6	The problem is usually associated with an EISA board, for example, there could be an IRQ mismatch.
The console program reports error beep codes at power-up	Use the error beep codes and/or console terminal to determine the error.
The power-up screen includes error messages	Examine the console event log (enter the cat el command) or the power-up screen to check for embedded error messages recorded during power-up.
The power-up screen or console event log indicates problems with mass storage devices, or if storage devices are missing from the SHOW CONFIG display	Use the troubleshooting tables to determine the problem.
The power-up screen or console event log indicates problems with EISA devices, or if EISA devices are missing from the SHOW CONFIG display	Do the following: • Use the troubleshooting table to determine the problem. • Run the ROM-based diagnostic (RBD) tests to verify the problem.

See Chapter 6 for a discussion of error codes and messages.

See Appendix E for information on the SHOW CONFIG command and display.

Table 1–4 Diagnostic Flow for Boot Problems

If...	Then...
The system cannot find the boot device	Do the following: • Check the system configuration for the correct device parameters (node ID, device name, and so on) by using the SHOW CONFIG and SHOW DEVICE commands. • Check the system configuration for the correct environment variable settings by examining the auto_action, bootdef_dev, and os_type environment variables. • Run the device tests to check that the boot device is operating.
Device does not boot	Check the ew*0_protocols or er*0_protocols environment variable settings. For problems booting over a network, check the ew*0_protocols or er*0_protocols environment variable settings.

Table 1–5 Diagnostic Flow for Errors Reported by the Operating System

If...	Then...
The system is hung or has crashed	Examine the crash dump file. Refer to the *OpenVMS AXP Alpha System Dump Analyzer Utility Manual* for information on how to interpret OpenVMS crash dump files.
Errors have been logged and the operating system is up	Examine the operating system error log files to isolate the problem. If the problem occurs intermittently, run an operating system exerciser, such as DEC VET, to stress the system. Refer to the *DEC Verifier and Exerciser Tool User's Guide* for instructions on running DEC VET.

1.2 Maintenance Features

StorageWorks FDDI Servers have the following features to aid in troubleshooting and maintenance:

- **Initialization diagnostics**

  Various levels of initialization diagnostics execute on the controller. These tests ensure that the subsystem is ready to come online after it has been reset, powered on, and so forth. You can elect to rerun many of the diagnostics after initialization completes in order to test controller operation.

- **Terminal access**

  You can use a **virtual terminal** (host terminal) or a **maintenance terminal** to check status and set operating parameters. The terminal connection provides access to the following:

  - Command Line Interpreter (Appendix E)
  - Error messages (Chapter 6)
  - Error logs (Chapter 6 and Appendix C)

- **Utilities**

  The VTDPY utility displays current controller state and performance data, including processor utilization, host port activity and status, device state, **logical unit** state, and cache and I/O performance.

  The configuration utility (CONFIG) checks the **SCSI device** ports for any device not previously added. This utility adds and names these devices.

  The Firmware Licensing System (FLS) enables and disables licensed controller features, including **RAID** (Redundant Array of Independent Disks), **mirroring** and write-back caching.

  The Fault Management Utility (FMU) controls some spontaneous error displays and displays controller last failure and memory system failure information.

  See Chapter 7 for detailed information on each utility.

- **Exercisers**

  The controller can run both a disk inline exerciser (**DILX**) and a tape inline exerciser (**TILX**). These exercisers simulate high levels of user activity. Running them provides performance information you can use to determine the health of the controller and the devices attached to it. See Chapter 7 for more information about the exercisers.

- **Controller warm swap**

  You can safely remove and replace, or **warm swap**, one controller in a dual-redundant configuration. When you warm swap a controller, you are replacing a controller in the most transparent method available to the HS controller subsystem. Warm swapping a controller has minimal system and device impact. For more information on warm swapping, see Chapter 8.

- **HS1CP front panel**

  The seven buttons and LEDs on the front panel of the HS1CP control the SCSI ports and report fault and normal conditions. See Chapter 3 for a complete description of the HS1CP front panel.

1.3 Common Problems

If hardware failure occurs, check for the following problems:

- Loose or missing terminators
- Loose or damaged cables or connectors

1.4 Troubleshooting Basics

When you are presented with a component failure, do the following:

1. Make a note of all visual indicators (OCP, device LEDs, and/or error messages) available to you.
2. Extract and read host error logs.
3. For surviving components in dual-redundant pairs, try entering a RESTART command. The surviving component could be keeping its companion from operating.
4. Errors can be intermittent. Reseat a connection or cable to see if the error clears.
5. See if the error indication changes after reseating a component. If the error remains the same, look up information for that error. If the indication changes, look up information for the newer error.
6. Always consider reseating the controller and/or component when troubleshooting. Poor connections between a component and backplane can cause a variety of errors.

1.5 Service Tools and Utilities

This section lists the array of service tools and utilities available for acceptance testing, diagnosis, and serviceability and provides recommendations for their use.

1.5.1 Error Handling/Logging

OpenVMS operating systems provide recovery from errors, fault handling, and event logging. The DECevent™ Translation and Reporting Utility for OpenVMS™ and Error Report Formatter (ERF) provide bit-to-text translation of event logs for interpretation.

Recommended Use: Analysis of error logs is the primary method of diagnosis and fault isolation. If the system is up, or you are able to bring it up, look at this information first.

1.5.2 ROM-Based Diagnostics

Many ROM-based diagnostics and exercisers are embedded in server processor systems. ROM-based diagnostics execute automatically at power-up and can be invoked in console mode using console commands.

Recommended Use: ROM-based diagnostics are the primary means of testing the console environment and diagnosing the CPU, memory, Ethernet, I/O buses, and SCSI subsystems. Use ROM-based diagnostics in the acceptance test procedures when you install a system, add a memory module, or replace the following:

- CPU module
- Memory module
- Motherboard
- I/O bus device
- Storage device

1.5.3 Loopback Tests

Internal and external loopback tests are used to isolate a failure by testing segments of a particular control or data path. The loopback tests are a subset of the ROM-based diagnostics.

Recommended Use: Use loopback tests to isolate problems with the COM2 serial port, the parallel port, and Ethernet controllers.

1.5.4 Firmware Console Commands

Console commands are used to set and examine environment variables and device parameters, as well as to invoke ROM-based diagnostics and exercisers.

The following are examples of console commands:

- The SHOW MEMORY, SHOW CONFIG, and SHOW DEVICE commands are used to examine the configuration.
- The SET (bootdef_dev and auto_action) commands are used to set environment variables.
- The CDP command is used to configure DSSI parameters.

Recommended Use: Use console commands to set and examine environment variables and device parameters and to run RBDs.

1.5.5 Operating System Exercisers (DEC VET)

The Digital Verifier and Exerciser Tool (DEC VET) is supported by the OpenVMS operating system. DEC VET performs exerciser-oriented maintenance testing of both hardware and operating system.

Recommended Use: Use DEC VET as part of acceptance testing to ensure that the CPU, memory, disk, tape, file system, and network are interacting properly. Also use DEC VET to stress test the user's environment and configuration by simulating system operation under heavy loads to diagnose intermittent system failures.

1.5.6 Crash Dumps

For fatal errors, such as fatal bugchecks, the OpenVMS operating system saves the contents of memory to a crash dump file.

Recommended Use: Crash dump files can be used to determine why the system crashed. To save a crash dump file for analysis, you need to know the proper system settings. Refer to the *OpenVMS AXP Alpha System Dump Analyzer Utility Manual*.

1.6 Information Services

Several information services are available, including online information for servicers and customers, computer-based training, and maintenance documentation database services. The following sections contain a brief description of some of these services.

1.6.1 Training

The following computer-based training (CBT) and lecture lab courses are available from the Digital training center:

- Alpha Concepts
- ISA and EISA Bus Concepts: EY–I113E–P0
- RAID Concepts: EY–N935E
- SCSI Concepts and Troubleshooting: EY–P841E, EY–N838E

1.6.2 Digital Assisted Services

Digital Assisted Services (DAS) offers products, services, and programs to customers who participate in the maintenance of Digital computer equipment. Components of DAS include:

- Spare parts and kits
- Diagnostics and service information/documentation
- Tools and test equipment
- Parts repair services, including Field Change Orders (FCOs)

2

Power-Up Diagnostics for the Server Processor

This chapter provides information about power-up diagnostics for the server processor and discusses the following:

- Interpreting error beep codes
- Power-up screen
- Storage device LEDs
- EISA bus problems indicated at power-up
- The power-up sequence
- Firmware power-up diagnostics

2.1 Interpreting Error Beep Codes

If errors are detected at power-up, audible beep codes are emitted from the system.

The beep codes are the primary diagnostic tool for troubleshooting problems when console mode cannot be accessed.

Table 2–1 Interpreting Error Beep Codes

Beep Code	Problem	Corrective Action
1-1-4	The SROM code is unable to load the console code: Flash ROM header area or checksum error detected.	Do the following: 1. Use the Fail-Safe Loader to load new ARC/SRM console code. 2. If successfully loading new console firmware does not solve the problem, replace the motherboard.

(continued on next page)

Table 2–1 (Cont.) Interpreting Error Beep Codes

Beep Code	Problem	Corrective Action
3-3-1	Generic system failure. Possible problem sources include the following motherboard components: • Native SCSI controller (NCR 53C810) • Remote I/O chip (Intel 87312) • NVRAM chip (position E14)	Do the following: 1. Replace the NVRAM chip (E14) on the system motherboard (see Chapter 8). 2. If replacing the NVRAM chip did not solve the problem, replace the motherboard (see Chapter 8).
1-2-1	TOY NVRAM failure.	Replace the TOY NVRAM chip (E78) on the system motherboard (see Chapter 8).
1-3-3	No usable memory detected.	Do the following: 1. Verify that the memory modules are properly seated and try powering up again. 2. Swap bank 0 memory with known good memory and run SROM memory tests at power-up (see Section 2.1). 3. If populating bank 0 with known good memory does not solve the problem, replace the CPU daughterboard (see Chapter 8). 4. If replacing the CPU daughterboard does not solve the problem, replace the motherboard (see Chapter 8).

SROM Memory Power-Up Tests

To test SIMM memory and report the position of a failing SIMM, set SROM power-up tests by using jumper J1 on the CPU (see Figure 2–1). The progress and results of these tests are reported on the LCD display on the operator control panel (OCP).

To thoroughly test memory and data paths, complete the SROM tests in the order presented in Table 2–2. If a SIMM is reported as bad, replace the SIMM (see Chapter 8 and resume testing at bank 4 (Memory Test).

Table 2–2 SROM Memory Tests, CPU Jumper J1

Bank #	Test Description	Test Results
6	Backup Cache Tag Test	Test status displays on OCP: `1.2.3.done.` If the tests take longer than a few seconds between each number displayed in the test count, there is a problem with the cache—replace the CPU daughterboard (see Chapter 8).
2	Cache Test: Tests backup cache.	Test status displays on OCP: `....done.` If the test takes longer than a few seconds to complete, there is a problem with the backup cache—replace the CPU daughter board (see Chapter 8).

4	Memory Test: Tests memory with backup and data cache disabled.	Test status displays on OCP:
		`12345.done.`
		If an error is detected, the bank number and failing SIMM position are displayed. The following OCP message indicates a failing SIMM at bank 0, SIMM position 2:
		`FAIL B:0 S:2`
		Test duration: Approximately 10 seconds per 8 megabytes (MB) of memory.
		Figure 2–2 shows the memory bank and SIMM layout for the server processor. After determining the bad SIMM, refer to Chapter 8 for instructions on replacing FRUs.
		Note: The memory tests do not test the ECC SIMMs. If the operating system logs five or more single-bit correctable errors, swap the suspected ECC SIMMs with good SIMMs and repeat the memory test.

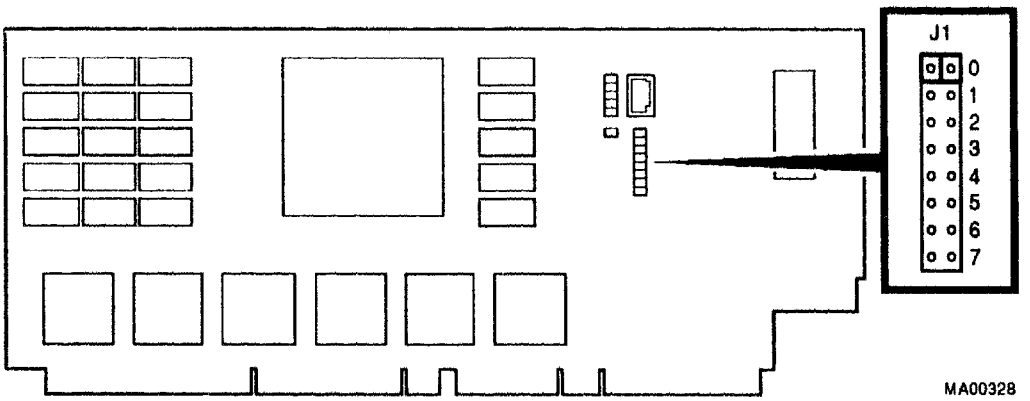
5	Memory Test, Cache Enabled: Tests memory with backup and data cache enabled.	Test status displays on OCP:
		`12345.done.`
		If an error is detected, the bank number and failing SIMM position are displayed. The following OCP message indicates a failing SIMM at bank 0, SIMM position 2:
		`FAIL B:0 S:2`
		Test duration: Approximately 2 seconds per 8 (MB) of memory.
		Figure 2–2 shows the bank and SIMM layout for the server processor. After determining the bad SIMM, refer to Chapter 8 for instructions on replacing FRUs.
		Note: The memory tests do not test the ECC SIMMs. If the operating system logs five or more single-bit correctable errors, swap the suspected ECC SIMMs with good SIMMs and repeat the memory test.

(continued on next page)

Table 2–2 (Cont.) SROM Memory Tests, CPU Jumper J1

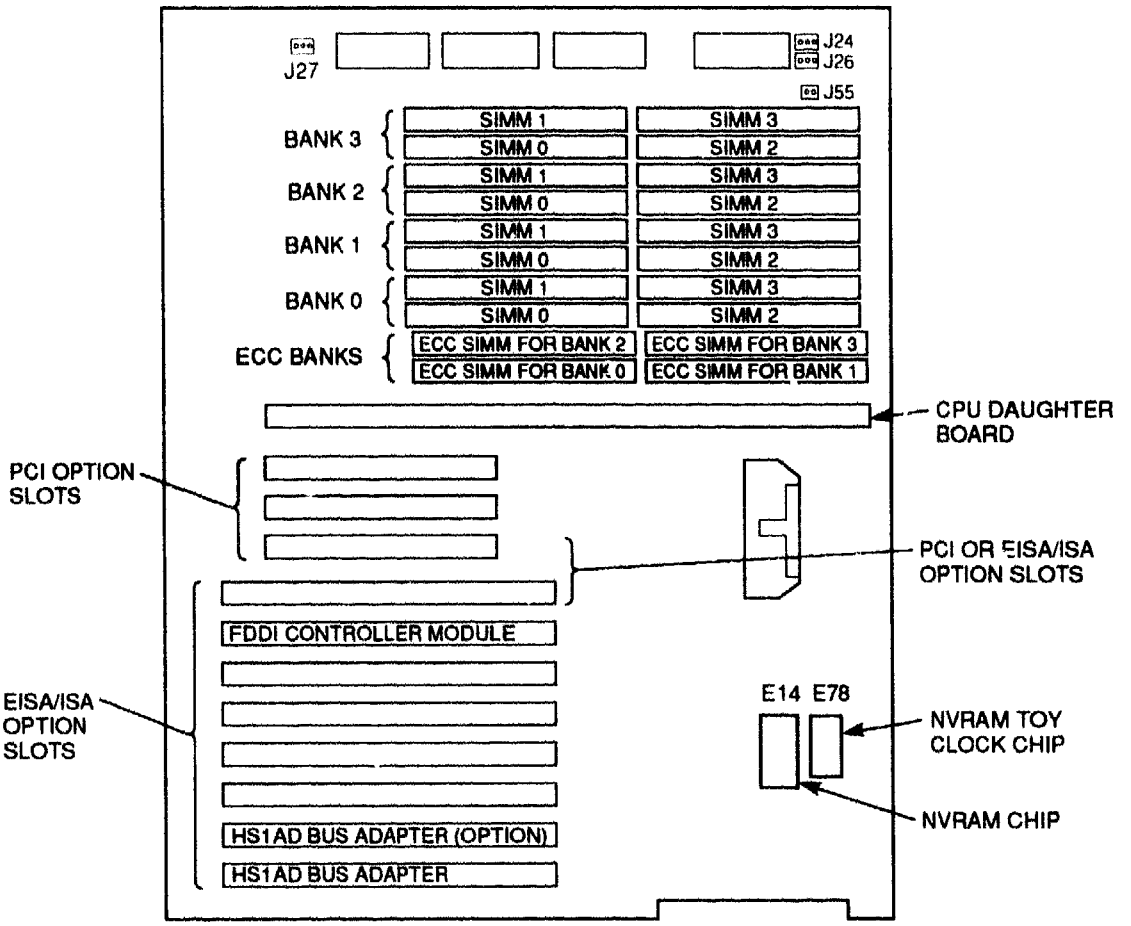
Bank #	Test Description	Test Results
3	Backup Cache Test: Tests backup cache alternatively with data cache enabled, then disabled.	Test status displays on OCP: d 12345.done. D 12345.done. D 12345.done. d 12345.done. If an error is detected, the bank number and failing SIMM position are displayed. The following OCP message indicates a failing SIMM at bank 0, SIMM position 2: FAIL B:0 S:2 Test duration: Approximately 2 seconds per 8 MB of memory. Figure 2–2 shows the bank and SIMM layout for the Server processor. After determining the bad SIMM, refer to Chapter 8 for instructions on replacing FRUs. Note: The memory tests do not test the ECC SIMMs. If the operating system logs five or more single-bit correctable errors, swap the suspected ECC SIMMs with good SIMMs and repeat the memory test.

Figure 2–1 Jumper J1 on the CPU Daughterboard



Bank	Jumper Setting
0	Standard boot setting (default)
1	Miniconsole setting: internal use only
2	SROM CacheTest: backup cache test
3	SROM BCacheTest: backup cache and memory test
4	SROM memTest: memory test with backup and data cache disabled
5	SROM memTestCacheOn: memory test with backup and data cache enabled
6	SROM BCache Tag Test: backup cache tag test
7	Fail-Safe Loader setting: selects fail-safe loader firmware

Figure 2–2 Server Processor Motherboard Layout



CXO-4626A-MC

2.2 Power-Up Screen

During power-up self-tests, the test status and result are displayed on the console terminal. Information similar to the following example should be displayed on the screen:

```
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
>>>
```

Note

If the power-up display stops on e6, an EISA or PCI board is causing the system to hang.

Console Event Log

The server processor maintains a console event log consisting of status messages received during power-on self-tests. If problems occur during power-up, standard error messages indicated by asterisks (***) may be embedded in the console event log. To display a console event log, use the cat el command.

Note

To stop the screen display from scrolling, press Ctrl/S. To resume scrolling, press Ctrl/Q.

You also can use the more el command to display the console event log one screen at a time.

The following example shows a console event log that contains a standard error message:

```
>>> cat el
*** keyboard not plugged in... ❶
ff.fe.fd.fc.fb.fa.f9.f8.f7.f6.f5.
ef.df.ee.f4.ed.ec.eb.ea.e9.e8.e7.e6.port pka0.7.0.6.0 initialized,
scripts are at 4f7faa0
resetting the SCSI bus on pka0.7.0.6.0
port pkb0.7.0.12.0 initialized, scripts are at 4f82be0
resetting the SCSI bus on pkb0.7.0.12.0
e5.e4.e3.e2.e1.e0.
V1.1-1, built on Nov  4 1994 at 16:44:07
device dka400.4.0.6.0 (RRD43) found on pka0.4.0.6.0
>>>
```

❶ The keyboard is not plugged in or is not working.

2.3 Storage Device LEDs

Storage device LEDs indicate the status of the device:

- Figure 2–3 shows the Activity LED for the floppy drive. This LED is on when the drive is in use.
- Figure 2–4 shows the Activity LED for the CD–ROM drive. This LED is on when the drive is in use.

For information on other storage devices, refer to the documentation provided by the manufacturer or vendor.

Figure 2–3 Floppy Drive Activity LED

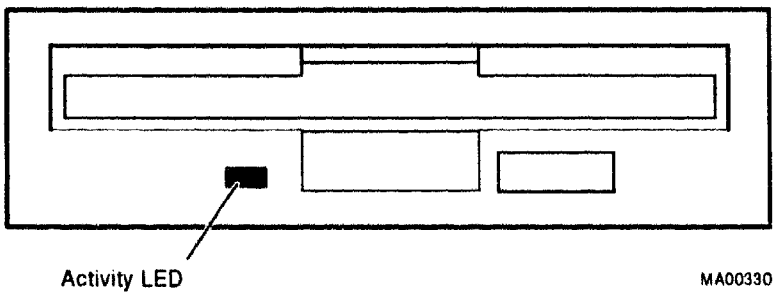
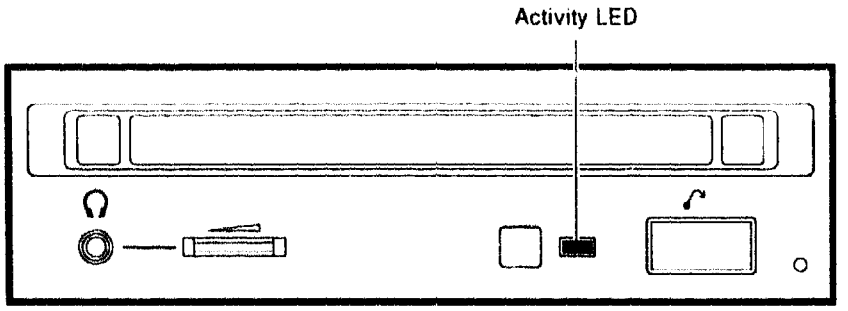


Figure 2–4 CD-ROM Drive Activity LED



2.4 EISA Bus Problems Indicated at Power-Up

EISA bus failures at power-up are usually indicated by the following messages displayed during power-up:

```
EISA Configuration Error. Run the EISA Configuration Utility.
```

Run the EISA Configuration Utility (ECU) (Section 5.22) when this message is displayed. Other problems are indicated by EISA devices missing from the SHOW CONFIG display.

Table 2–3 provides steps for troubleshooting EISA bus problems that persist after you run the ECU.

Table 2–3 EISA Troubleshooting

Step	Action
1	Confirm that the EISA module and any cabling are properly seated.
2	Run the ECU to: • Confirm that the system has been configured with the most recently installed controllers. • See what the hardware jumper and switch setting should be for each EISA controller. • See what the software setting should be for each EISA controller. • See if the ECU deactivated (<>) any controllers to prevent conflict. • See if any controllers are locked (!), which limits the ECU's ability to change resource assignments.
3	Confirm that the hardware jumpers and switches on EISA controllers reflect the settings indicated by the ECU. Start with the last EISA module tested.
4	Run ROM-based diagnostics for the type of option: • Storage adapter—Run `test` to exercise the storage devices off the EISA controller option (see Section 4.1.3.1). • Ethernet adapter—Run `netew` or `network` to exercise an Ethernet adapter (see Sections 4.1.3.4 and 4.1.3.5).

(continued on next page)

Table 2–3 (Cont.) EISA Troubleshooting

Step	Action
5	Check for a bad slot by moving the last installed controller to a different slot.
6	Call Digital Multivendor Customer Services for help.

2.4.1 Additional EISA Troubleshooting Tips

The following tips can aid in isolating EISA bus problems:

- Peripheral device controllers need to be seated (inserted) carefully, but firmly, into their slot to make all necessary contacts. Improper seating is a common source of problems for EISA modules.
- Be sure you run the correct version of ECU for the operating system.
- The CFG files supplied with the option you want to install may not work on StorageWorks FDDI Server systems. Some CFG files call overlay files that are not required on this system or may refer to inappropriate system resources (for example, BIOS addresses). Contact the option vendor to obtain the proper CFG file.
- Peripherals cannot share direct memory access (DMA) channels. Assignment of more than one peripheral to the same DMA channel can cause unpredictable results or even loss of function of the EISA module.
- Not all EISA products work together. EISA is an open standard, and not every EISA product or combination of products can be tested. Violations of specifications may matter in some configurations, but not in others.

  Manufacturers of EISA options often test the most common combinations and may have a list of EISA options that do not function in combination with particular systems. Be sure to check the documentation or contact the option vendor for the most up-to-date information.
- EISA systems do not function unless they first are configured using the ECU.
- The ECU does not notify you if the configuration program diskette is write-protected when it attempts to write the system configuration file (SYSTEM.SCI) to the diskette.

2.4.2 The Fail-Safe Loader

The fail-safe loader (FSL) allows you to attempt to recover when one of the following is the cause of a problem getting to the console program under normal power-up:

- A power failure or accidental power down during a firmware upgrade
- A checksum failure or flash ROM header error while the SROM code is trying to load the SRM/ARC console firmware

Note

The fail-safe loader should be used only when a failure at power-up prohibits you from getting to the console program. You cannot boot an operating system from the fail-safe loader.

If a checksum error is detected when the SRM/ARC console is loading at power-up (error beep code 1-1-4), you need to activate the fail-safe loader and reinstall the firmware.

From the FSL program, you can update or load new console firmware.

Activating the Fail-Safe Loader

To activate the FSL, do the following:

1. Install the jumper at bank 7 of the J1 jumper on the CPU daughterboard (Figure 2–1). The jumper is usually installed in the standard boot setting (bank 0).
2. Install the console firmware floppy diskette and turn on the system.
3. Reinstall the console firmware from a firmware diskette.
4. When you have finished, power down and return the J1 jumper to the standard boot setting (bank 0).

Bank	Jumper Setting
0	Standard boot setting (default)
1	Miniconsole setting: Internal use only
2	SROM CacheTest: Backup cache test
3	SROM BCacheTest: Backup cache and memory test
4	SROM memTest: Memory test with backup and data cache disabled
5	SROM memTestCacheOn: Memory test with backup and data cache enabled
6	SROM BCache Tag Test: Backup cache tag test
7	Fail-Safe Loader setting: Selects fail-safe loader firmware

2.5 The Power-Up Sequence

During the StorageWorks FDDI Server power-up sequence, the power supplies are stabilized and the system is initialized and tested through the firmware power-on self-tests.

The power-up sequence includes the following:

- Power supply power-up:
  - AC power-up
  - DC power-up
- Two sets of power-on diagnostics:
  - Serial ROM diagnostics
  - Console firmware-based diagnostics

2.5.1 AC Power-Up Sequence

The following power-up sequence occurs when AC power is applied to the system (system is plugged in) or when electricity is restored after a power outage:

1. The front end of the power supply begins operation and energizes.
2. The power supply then waits for the DC power to be enabled.

2.5.2 DC Power-Up Sequence

Apply dc power to the system with the DC On/Off button on the operator control panel.

A summary of the dc power-up sequence follows:

1. When the DC On/Off button is pressed, the power supply checks for a POK_H condition.
2. 12V, 5V, 3.3V, and -12V outputs are energized and stabilized. If the outputs do not come into regulation, the power-up sequence is aborted and the power supply enters the latching-shutdown mode.

2.6 Firmware Power-Up Diagnostics

After successful completion of ac and dc power-up sequences, the processor performs its power-up diagnostics. These tests verify system operation, load the system console, and test the core system (CPU, memory, and motherboard), including all boot path devices. These tests are performed as two distinct sets of diagnostics:

- Serial ROM diagnostics—These tests are loaded from the serial ROM located on the CPU daughterboard into the CPU's instruction cache (I-cache). The tests check the basic functionality of the system and load the console code from the FEPROM on the motherboard into system memory.

  Failures during these tests are indicated by audible error beep codes. Failures of customized SROM tests, set using the J1 jumper on the CPU daughterboard, are displayed on the operator control panel.

- Console firmware-based diagnostics—These tests are executed by the console code. They test the core system, including all boot path devices.

  Failures during these tests are reported to the console terminal through the power-up screen or console event log.

2.6.1 Serial ROM Diagnostics

The serial ROM diagnostics are loaded into the CPU's instruction cache from the serial ROM on the CPU daughterboard. The diagnostics test the system in the following order:

1. Test the CPU and backup cache located on the CPU daughterboard.
2. Test the CPU module's system bus interface.
3. Test the system bus to the EISA bus bridge. If the EISA bridge fails, an audible error beep code sounds. The power-up tests continue despite these errors.
4. Configure the memory in the system and test only the first 4 MB of memory. If there is more than one memory module of the same size, the lowest numbered memory module (one closest to the CPU) is tested first.

If the memory test fails, the failing bank is mapped out and memory is reconfigured and retested. Testing continues until good memory is found. If good memory is not found, an error beep code (1-3-3) is generated and the power-up tests are terminated.

5. Check the data path to the FEPROMs on the motherboard.
6. The console program is loaded into memory from the FEPROM on the motherboard. A checksum test is executed for the console image. If the checksum test fails, an error beep code (1-1-4) is generated and the power-up tests are terminated.

   If the checksum test passes, control is passed to the console code, and the console firmware-based diagnostics are run.

2.6.2 Console Firmware-Based Diagnostics

Console firmware-based tests are executed once control is passed to the console code in memory. They check the system in the following order:

1. Perform a complete check of system memory. (Steps 2–5 may be completed in parallel.
2. Start the I/O drivers for mass storage devices and tapes. At this time, a complete functional check of the machine is made. After the I/O drivers are started, the console program continuously polls the bus for devices (approximately every 20 or 30 seconds).
3. Check that EISA configuration information is present in NVRAM for each EISA module detected and that no information is present for modules that have been removed.
4. Run exercisers on the drives currently seen by the system.

   Note

   This step does not ensure that all disks in the system are tested or that any device drivers are completely tested. Spin-up time varies for different drives, so not all disks may be on line at this point in the power-up sequence. To ensure complete testing of disk devices, use the TEST command (Section 4.1.3.1).

5. Enter console mode or boot the operating system. This action is determined by the `auto_action` environment variable.

2.6.3 Operator Control Panel Indicators

The OCP consists of a 16-character liquid crystal display (LCD), DC power-enable switch, momentary Halt switch, momentary Reset switch, and two light-emitting diodes (LEDs), indicating DC POWER ON and SYSTEM HALT. The controls and indicators on the OCP are shown in Figure 2–5.

Figure 2–5 Operator Control Panel

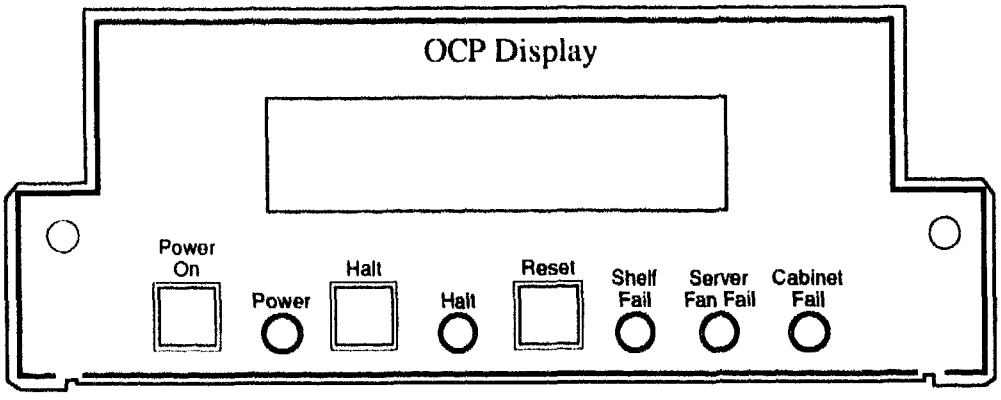


Table 2–4 describes the controls and indicators.

Table 2–4 OCP Controls and Indicators Description

Control or Indicator	Function
Power On	Switches the system unit on and off.
Power indicator	Lights when the system unit is turned on. In a redundant system, the light displays amber if one power supply is failed.
Halt switch	Halts an OpenVMS system, returning it to console mode control.
Halt indicator	Lights when the system is halted.
Reset switch	Reinitializes the system and performs startup tests.
Shelf fail indicator	Not used.
Server fan fail indicator	Lights when any system fan is not working.
Cabinet fail indicator	When a redundant power supply fails, this light, along with the power indicator, displays red for three to four seconds. It then goes off, and the power indicator displays amber.
Operator control panel display	Displays diagnostic and error messages.

2.6.4 Cycling Power with the OCP

During the installation procedure, you are directed to cycle (remove, then restore) power to the server processor. Use the Power On switch on the OCP to perform this step when directed to cycle the server processor.

2.6.5 Server Fan Fail Indicator

The Server Fan Fail indicator lights when either fan 1 or fan 2 fails. Periodically check to make sure the Server Fan Fail indicator is not lighted.

Failure of fan 3 (the fan cooling the microprocessor) causes the server processor to shut down. The indicator does not light in this instance. In the event of a fan 3 failure, the following message appears on the server processor terminal:

```
System Machine Check Through Vector 00000660
logout frame address 0x6048 code 0x223
Machine Check Code---> 0x223 Fan 1 Failure
```

Note

Fan 3 is the hardware designation for the fan that provides the primary cooling air to the microprocessor. This fan is designated as fan 1 to the server processor software. Because of this, the error message shown indicates fan 1 failure.

In the event of a fan failure, call your service representative.

3

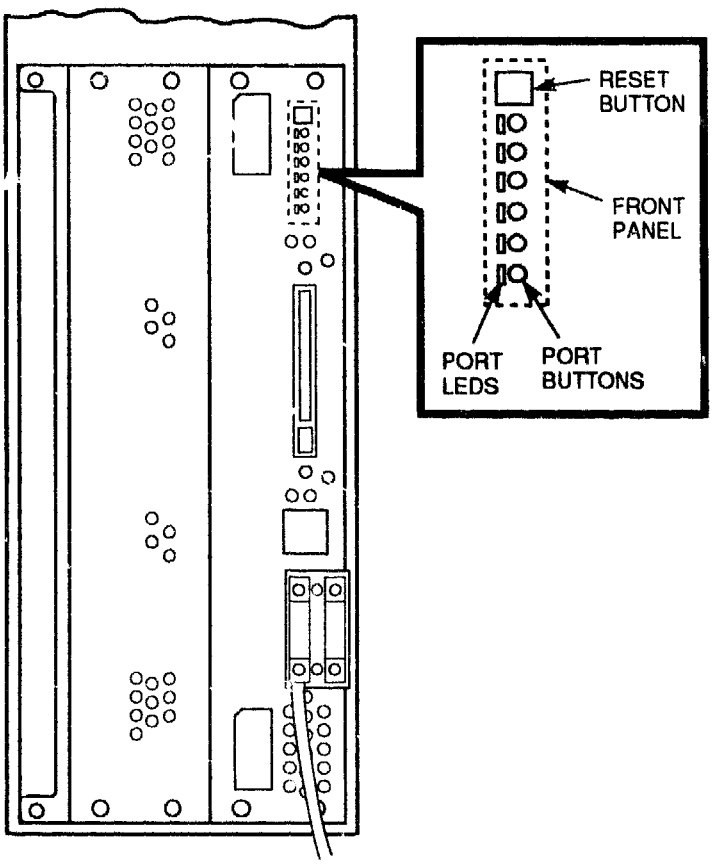
Power-Up Diagnostics for HS1CP

This chapter provides information on power-up diagnostics for the HS1CP and discusses the following:

- The power-up sequence
- Firmware power-up diagnostics

Figure 3–1 shows the layout of the HS1CP.

Figure 3–1 HS1CP Layout



CXO-4461B-MC

3.1 The Power-Up Sequence

The power-up sequence includes the following:

AC Power-Up Sequence

AC power is applied when the CDU breaker is in the On position.

DC Power-Up Sequence

DC power is applied to the system with the DC On/Off button on the OCP of the server processor.

3.2 Firmware Power-Up Diagnostics

After successful completion of ac and dc power-up sequences, the HS1CP performs its firmware diagnostics. The following sections describe the HS1CP front panel indicators, normal operation, and fault notification.

3.2.1 HS1CP Front Panel Indicators

The front panel includes the following indicators:

- One Reset button with an embedded green LED
- One button per SCSI port
- Six amber LEDs

Figure 3–1 shows an example of an HS1CP and the front panel indicators. The buttons and LEDs serve different functions with respect to controlling the SCSI ports and/or reporting fault and normal conditions.

3.2.2 HS1CP Normal Operation

The green LED button reflects the state of the HS1CP and the host interface. Once HS1CP initialization completes and its firmware is functioning, the green button flashes continuously at 1 Hz (once per second). Pressing the green button during normal operation resets the HS1CP.

Under normal operation, the amber LEDs indicate the state of the respective SCSI–2 device buses attached to the HS1CP. When the devices on the buses are functioning correctly, the amber LEDs do not light or flash.

Pressing one of the port buttons during normal operation lights its corresponding amber LED and momentarily quiesces its SCSI–2 port. You must quiesce a port to remove or warm swap a device on the SCSI–2 bus for that port. (Use a small pointed object to press the port button.)

See Chapter 8 for a detailed description of removing and replacing devices.

3.2.3 Fault Notification

The front panel LEDs display information when the HS1CP encounters a problem with a device configuration, a device, or the HS1CP itself.

Should a configuration mismatch or a device fault occur, the amber LED for the affected device's bus lights continuously.

For HS1CP problems, LED codes determined by internal diagnostics and operating firmware indicate either HS1CP faults or HS1CP operating firmware program card faults. In either case, the single green Reset LED lights continuously when an error is detected. The remaining amber LEDs display the error codes in two different ways:

- The error code lights continuously for faults detected by internal diagnostic and initialization routines. See Figure 6–2 to determine what these codes mean.
- The error code flashes at 3 Hz representing faults that occur during normal controller operation. See Figure 6–3 to determine what these codes mean.

Note

If the controller finds errors during initialization or self-test diagnostics, any flashing codes momentarily stop flashing while the error messages are posted to the console. Normal flashing resumes once the display completes. Do not attempt to look up the codes while this information is printing at the terminal.

3.2.4 Device LEDs

The storage devices (SBBs) and their power supplies have LEDs to indicate power and status. You can use these LEDs in conjunction with the front panel indicators to isolate certain faults, as discussed in the following sections.

3.2.4.1 Storage SBB Status

Device shelves monitor the status of the storage SBBs. When a fault occurs, the fault and the SBB device address (SCSI target ID) are reported to the HS1CP for processing. The SBB internal fault/identity bus controls the fault (lower) LED.

As shown in Figures 3–2 and 3–3, each storage SBB has two LED indicators that display the SBB's status. These LEDs have three states: on, off, and flashing.

- The upper LED (green) is the device activity LED and is on or flashing when the SBB is active.

CAUTION

Do not remove a storage SBB when the upper LED is on or flashing. This can cause the loss or corruption of data.

- The lower LED (amber) is the storage SBB fault LED and indicates an error condition when it is either on or flashing. When this LED indicates a fault, the HS1CP's amber LED for the device's port is lit continuously as well.

CAUTION

You should record which devices have lit/flashing fault LEDs before resetting the HS1CP; resetting it can clear this LED temporarily, even though the fault remains.

Figure 3–2 3.5-Inch Storage SBB LEDs

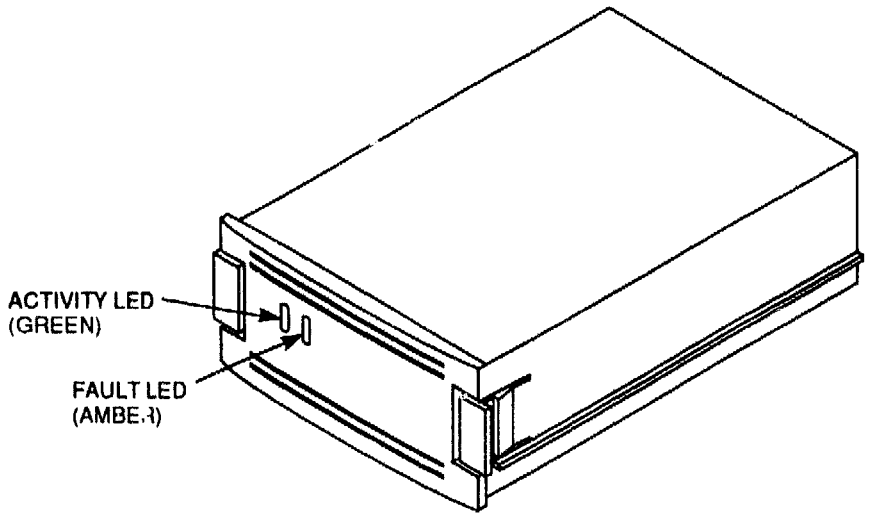


Figure 3–3 5.25-Inch Storage SBB LEDs

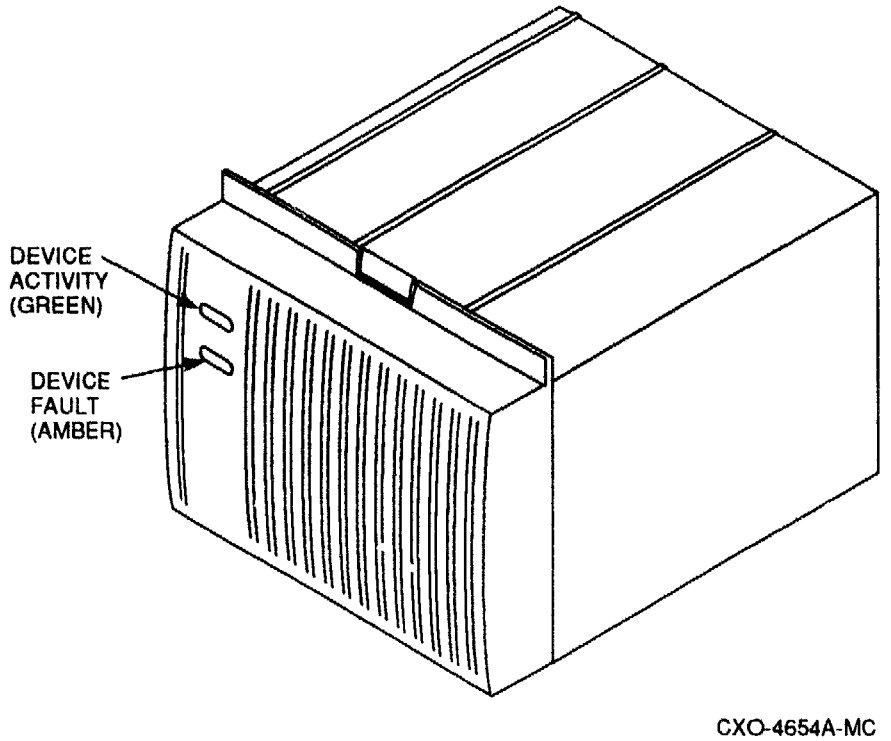


Table 3–1 explains the SBB LED codes.

Table 3–1 Storage SBB Status LEDs

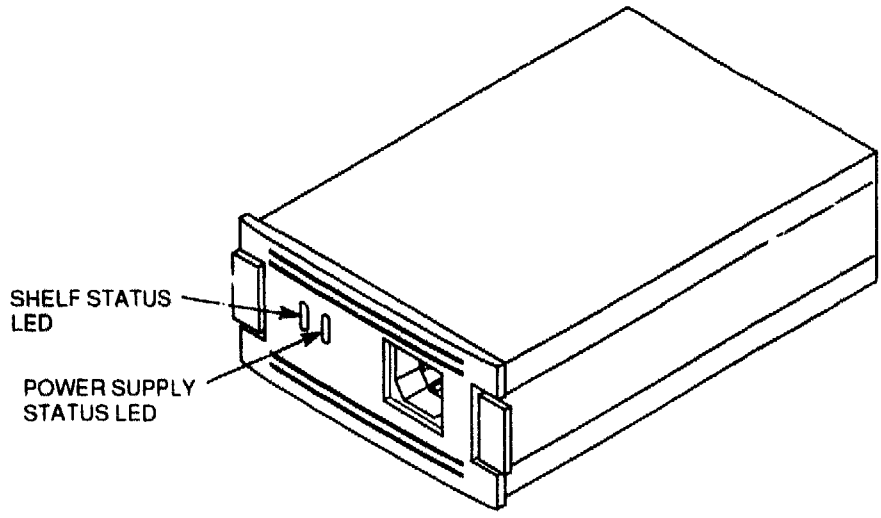
LED	Status	Indication
Device activity Device fault	On Off	SBB is operating normally.
Device activity Device fault	Flashing Off	SBB is operating normally.
Device activity Device fault	Off Off	SBB is operating normally. The SBB is inactive, and there is no fault.
Device activity Device fault	On On	Fault status. SBB is probably not responding to control signals. Digital recommends that you replace the SBB.
Device activity Device fault	Off On	Fault status. SBB is inactive and spun down. Digital recommends that you replace the SBB.
Device activity Device fault	On Flashing	Fault status. SBB is active and is spinning down because of the fault.

3.2.4.2 Device Shelf Status and Power Supply Status

The status of both the device shelf blowers and power supplies is displayed on the power supply LEDs, as shown in Figure 3–4. The upper LED displays the shelf status and the lower LED displays the power supply status.

- When the upper LED is on, both the shelf blowers and the power supplies are functioning properly.
- When the upper LED is off, either a shelf blower or a power supply is not functioning properly.
- When the lower LED is off, either there is an input power problem or the power supply is not functioning.

Figure 3–4 Power Supply LEDs



For a detailed explanation of the power supply LED codes, see Tables 3–2 and 3–3.

Table 3–2 Shelf and Single Power Supply Status LEDs

Status LED	State	Indication
Shelf (upper) PS (lower)	On On	System is operating normally.
Shelf (upper) PS (lower)	Off On	Fault status. There is a shelf fault; there is no power supply fault. Replace blower as described in Chapter 8.
Shelf (upper) PS (lower)	Off Off	Fault status. Shelf and power supply fault. Replace power supply as described in Chapter 8.

Note

When a shelf has two power supplies, you must observe the LEDs on *both* power supplies to determine the status (see Table 3–3).

Table 3–3 Shelf and Dual Power Supply Status LEDs

Status LED	PS1†	PS2‡	Indication
Shelf (upper) PS (lower)	On On	On On	Normal status. System is operating normally.
Shelf (upper) PS (lower)	Off On	Off On	Fault status. There is a shelf fault; there is no power supply fault. Replace blower as described in Chapter 8.
Shelf (upper) PS (lower)	Off On	Off Off	Fault status. PS1 is operational. Replace PS2 as described in Chapter 8.
Shelf (upper) PS (lower)	Off Off	Off On	Fault status. PS2 is operational. Replace PS1 as described in Chapter 8.
Shelf (upper) PS (lower)	Off Off	Off Off	Fault status. Possible PS1 and PS2 fault or input power problem.

† Shelf power supply installed in slot 7.

‡ Redundant power supply installed in slot 6.

4

Running System Diagnostics

This chapter provides information on how to run system diagnostics for the server processor and for the HS1CP. This chapter covers the following topics:

- Diagnostic procedures for the server processor
- Diagnostic procedures for the HS1CP

4.1 Diagnostic Procedures for the Server Processor

The following sections provide information on diagnostic procedures for the server processor.

4.1.1 Running ROM-Based Diagnostics

ROM-based diagnostics (RBDs), which are part of the console firmware that is loaded from the FEPROM on the system motherboard, offer many powerful diagnostic utilities, including the ability to examine error logs from the console environment and run system- or device-specific exercisers.

The FDDI Server RBDs rely on exerciser modules to isolate errors rather than on functional tests. The exercisers are designed to run concurrently, providing a maximum bus interaction between the console drivers and the target devices.

The multitasking ability of the console firmware allows you to run diagnostics in the background (using the background operator "&" at the end of the command). You run RBDs by using console commands.

Note

ROM-based diagnostics, including the TEST command, are run from the console. RBDs report errors to the console terminal and/or the console event log.

4.1.2 Command Summary

Table 4–1 provides a summary of the diagnostic and related commands.

Table 4–1 Summary of Diagnostic and Related Commands

Command	Function	Reference
Acceptance Testing		
TEST	Quickly tests the core system. The TEST command is the primary diagnostic for acceptance testing and console environment diagnosis.	Section 4.1.3.1
Error Reporting		
CAT EL	Displays the console event log.	Section 4.1.3.2
MORE EL	Displays the console event log one screen at a time.	Section 4.1.3.2
Extended Testing/Troubleshooting		
MEMORY	Runs memory exercises each time the command is entered. These exercises run concurrently in the background.	Section 4.1.3.3
NET -IC	Initializes the MOP counters for the specified Ethernet port.	Section 4.1.3.7
NET -S	Displays the MOP counters for the specified Ethernet port.	Section 4.1.3.6
NETEW	Runs external mop loopback tests for specified EISA-based ew* (DECchip 21040, TULIP) Ethernet ports.	Section 4.1.3.4
NETWORK	Runs external mop loopback tests for specified EISA-based er* (DEC 4220, LANCE) Ethernet ports.	Section 4.1.3.5
Loopback Testing		
TEST LB	Conducts loopback tests for COM2 and the parallel port in addition to quick core system tests.	Section 4.1.3.1
NETEW	Runs external mop loopback tests for specified EISA-based ew* (DECchip 21040, TULIP) Ethernet ports.	Section 4.1.3.4
NETWORK	Runs external mop loopback tests for specified EISA-based er* (DEC 4220, LANCE) Ethernet ports.	Section 4.1.3.5
Diagnostic-Related Commands		
KILL	Terminates a specified process.	Section 4.1.3.8
KILL_DIAGS	Terminates all currently executing diagnostics.	Section 4.1.3.8
SHOW_STATUS	Reports the status of currently executing test/exercisers.	Section 4.1.3.9

4.1.3 Command Reference

The following sections provides detailed information on the diagnostic commands and related commands.

4.1.3.1 The TEST Command

The TEST command runs firmware diagnostics for the entire system. The tests are run concurrently in the background. Fatal errors are reported to the console terminal.

Use the CAT EL command in conjunction with the TEST command to examine test/error information reported to the console event log.

Because the tests are run concurrently and indefinitely (until you stop them with the KILL_DIAGS command), they are useful in flushing out intermittent hardware problems.

Note

By default, no write tests are performed on disk and tape drives. Media must be installed to test the floppy drive and tape drives. A loopback connector is required for the COM2 (9-pin loopback connector, 12–27351–01) port.

Use the KILL command to terminate an individual diagnostic or the KILL_DIAGS command to terminate all diagnostics. Use the SHOW_STATUS display to determine the process ID when terminating an individual diagnostic test.

Note

A serial loopback connector (12–27351–01) must be installed on the COM2 serial port for the KILL_DIAGS command to successfully terminate system tests.

The TEST script tests devices in the following order:

1. **Console loopback tests if lb argument is specified: COM2 serial port and parallel port.**
2. **Network external loopback tests for E*A0. This test requires that the Ethernet port be terminated or connected to a live network; otherwise, the test fails.**
3. **Memory tests (one pass).**
4. **Read-only tests: DK* disks, DR* disks, DU* disks, MK* tapes, DV* floppy.**
5. **VGA console tests. These tests are run only if the console environment variable is set to "serial." The VGA console test displays rows of the letter "H."**

Synopsis:

TEST [lb]

Argument:

[lb]	The loopback option includes console loopback tests for the COM2 serial port and the parallel port during the test sequence.

Examples:

The system is tested and the tests complete successfully.

Note

Examine the console event log after running tests.

```
>>> TEST
Requires diskette and loopback connectors on COM2 and parallel port
type kill_diags to halt testing
type show_status to display testing progress
type cat el to redisplay recent errors
Testing COM2 port
Setting up network test, this will take about 20 seconds
Testing the network
     128 Meg of System Memory
     Bank 0 = 64 Mbytes(16 MB Per Simm) Starting at 0x00000000
     Bank 1 = 64 Mbytes(16 MB Per Simm) Starting at 0x01000000
     Bank 2 = No Memory Detected
     Bank 3 = No Memory Detected

Testing the memory
Testing parallel port
Testing the SCSI Disks
Non-destructive Test of the Floppy starteddka400.4.0.6.0 has no media
present or is disabled via the RUN/STOP switch
file open failed for dka400.4.0.6.0
Testing the VGA(Alphanumeric Mode only)
Printer offline
file open failed for para

>>> SHOW_STATUS

ID       Program     Device       Pass  Hard/Soft Bytes Written Bytes Read
-------- ---------- ------------ ------ --------- ------------- ------------
00000001       idle system           0    0    0             0            0
0000002d   exer_kid tta1             0    0    0             1            0
0000003d    nettest era0.0.0.2.1     43    0    0          1376         1376
00000045    memtest memory            7    0    0     424673280    424673280
00000052   exer_kid dka100.1.0.6     0    0    0             0      2688512
00000053   exer_kid dka200.2.0.6     0    0    0             0       922624
>>> KILL_DIAGS
```

The system is tested and the system reports a fatal error message. No network server responded to a loopback message. Ethernet connectivity on this system should be checked.

```
>>> TEST
Requires diskette and loopback connectors on COM2 and parallel port
type kill_diags to halt testing
type show_status to display testing progress
type cat el to redisplay recent errors
Testing COM2 port
Setting up network test, this will take about 20 seconds
Testing the network

*** Error (era0), Mop loop message timed out from:
08-00-2b-3b-42-fd

*** List index: 7 received count: 0 expected count 2
```

4.1.3.2 The CAT EL and MORE EL Commands

The CAT EL and MORE EL commands display the current contents of the console event log. Status and error messages (if problems occur) are logged to the console event log at power-up, during normal system operation, and while running system tests.

Standard error messages are indicated by asterisks (***).

When you invoke the CAT EL command, the contents of the console event log scroll by. Enter Ctrl/S to stop the screen from scrolling, and enter Ctrl/Q to resume scrolling.

The MORE EL command allows you to view the console event log one screen at a time.

Synopsis:

CAT EL

or

MORE EL

Examples:

The following examples show abbreviated console event logs that contain standard error messages:

❶ The first error message indicates the keyboard is not plugged in or is not working.

```
>>> CAT EL
*** keyboard not plugged in... ❶
ff.fe.fd.fc.fb.fa.f9.f8.f7.f6.f5.
ef.df.ee.f4.ed.ec.eb.ea.e9.e8.e7.e6.port pka0.7.0.6.0 initialized,
scripts are at 4f7faa0
resetting the SCSI bus on pka0.7.0.6.0
port pkb0.7.0.12.0 initialized, scripts are at 4f82be0
resetting the SCSI bus on pkb0.7.0.12.0
e5.e4.e3.e2.e1.e0.
V1.1-1, built on Nov  4 1994 at 16:44:07
device dka400.4.0.6.0 (RRD43) found on pka0.4.0.6.0
>>>
```

4.1.3.3 The MEMORY Command

The MEMORY command tests memory by running a memory exerciser each time the command is entered. The exercisers are run in the background and nothing is displayed unless an error occurs.

The number of exercisers, as well as the length of time for testing, depends on the context of the testing. Generally, running three to five exercisers for 15 minutes to 1 hour is sufficient for troubleshooting most memory problems.

Use the KILL command to terminate an individual diagnostic or the KILL_DIAGS command to terminate all diagnostics. Use the SHOW_STATUS display to determine the process ID when terminating an individual diagnostic test.

Synopsis:

MEMORY

Examples:

Example with no errors.

```
>>> MEMORY
>>> SHOW_STATUS

 ID       Program      Device       Pass  Hard/Soft Bytes Written  Bytes Read
-------- ------------ ------------ ------ --------- ------------- -------------
00000001         idle system           0    0    0             0             0
0000006b      memtest memory           1    0    0      53477376      53477376
00000071      memtest memory           1    0    0      31457280      31457280
00000077      memtest memory           1    0    0      24117248      24117248
>>> KILL_DIAGS
>>>
```

Example with a memory compare error indicating bad SIMMs.

```
>>> MEMORY

*** Hard Error - Error #44 - Memory compare error

Diagnostic Name        ID             Device  Pass  Test  Hard/Soft   1-JAN-2066
memtest          000000c8               brd0     1     1     1    0     12:00:01
Expected value:  00000004
Received value:  80000001
Failing addr:    800001c

*** End of Error ***

>>> KILL_DIAGS
>>>
```

4.1.3.4 The NETEW Command

The NETEW command is used to run MOP loopback tests for any EISA-based ew* (DECchip 21040, TULIP) Ethernet ports. The command also can be used to test a port on a "live" network.

The loopback tests are set to run continuously (-p pass_count set to 0). Use the KILL command or Ctrl/C to terminate an individual diagnostic or the KILL_DIAGS command to terminate all diagnostics. Use the SHOW_STATUS display to determine the process ID when terminating an individual diagnostic test.

Note

While some results of network tests are reported directly to the console, you should examine the console event log using the CAT EL or MORE EL commands for complete test results.

Synopsis:

NETEW

When the NETEW command is entered, the following script is executed:

```
net -sa ew*0>ndbr/lp_nodes_ew*0
set ew*0_loop_count 2 2>nl
set ew*0_loop_inc 1 2>nl
set ew*0_loop_patt ffffffff 2>nl
set ew*0_loop_size 10 2>nl
set ew*0_lp_msg_node 1 2>nl
net -cm ex ew*0
echo "Testing the network"
nettest ew*0 -sv 3 -mode nc -p 0 -w 1 &
```

The script builds a list of nodes for which to send MOP loopback packets, sets certain test environment variables, and tests the Ethernet port by using the following variation of the nettest exerciser:

nettest ew*0 -sv 3 -mode nc -p 0 -w 1 &

Testing an Ethernet Port:

```
>>> NETEW
>>> SHOW_STATUS
 ID       Program      Device      Pass  Hard/Soft Bytes Written Bytes Read
-------- ---------- ------------ ------ --------- ------------- ------------
00000001        idle system          0     0    0             0            0
000000d5     nettest ewa0.0.0.0.0    13     0    0        308672       308672

>>> KILL_DIAGS
>>>
```

4.1.3.5 The NETWORK Command

The NETWORK command is used to run MOP loopback tests for any EISA-based er* (DEC 4220, LANCE) Ethernet ports. The command also can be used to test a port on a "live" network.

The loopback tests are set to run continuously (-p pass_count set to 0). Use the KILL command (or Ctrl/C) to terminate an individual diagnostic or the KILL_DIAGS command to terminate all diagnostics. Use the SHOW_STATUS display to determine the process ID when terminating an individual diagnostic test.

Note

While some results of network tests are reported directly to the console, you should examine the console event log (using the CAT EL or MORE EL commands) for complete test results.

Synopsis:

NETWORK

When the NETWORK command is entered, the following script is executed:

```
echo "setting up the network test, this will take about 20 seconds"
net -stop er*0
net -sa er*0>ndbr/lp_nodes_er*0
net ic er*0
set er*0_loop_count 2 2>nl
set er*0_loop_inc 1 2>nl
set er*0_loop_patt ffffffff 2>nl
set er*0_loop_size 10 2>nl
set er*0_lp_msg_node 1 2>nl
set er*0_mode 44 2>nl
net -start er*0
echo "Testing the network"
nettest er*0 -sv 3 -mode nc -p 0 -w 1 &
```

The script builds a list of nodes for which to send MOP loopback packets, sets certain test environment variables, and tests the Ethernet port by using the following variation of the nettest exerciser:

nettest er*0 -sv 3 -mode nc -p 0 -w 1 &

Testing an Ethernet Port:

```
>>> NETWORK
>>> SHOW_STATUS
 ID       Program      Device       Pass  Hard/Soft Bytes Written Bytes Read
-------- ----------- ----------- ------ --------- ------------- ------------
00000001         idle system           0    0    0             0            0
000000d5      nettest era0.0.0.0.0    13    0    0        308672       308672
>>> KILL_DIAGS
>>>
```

4.1.3.6 The NET -S Command

The NET -S command displays the MOP counters for the specified Ethernet port.

Synopsis:

NET -S EWA0

Example:

```
>>> NET -S EWA0

Status counts:
ti: 72 tps: 0 tu: 47 tjt: 0 unf: 0 ri: 70 ru: 0
rps: 0 rwt: 0 at: 0 fd: 0 lnf: 0 se: 0 tbf: 0
tto: 1 lkf: 1 ato: 1 nc: 71 oc: 0

MOP BLOCK:
 Network list size: 0

MOP COUNTERS:
Time since zeroed (Secs): 42

TX:
 Bytes: 0 Frames: 0
 Deferred: 1 One collision: 0 Multi collisions: 0
TX Failures:
 Excessive collisions: 0 Carrier check: 0 Short circuit: 71
 Open circuit: 0 Long frame: 0 Remote defer: 0
 Collision detect: 71

RX:
 Bytes: 49972 Frames: 70
 Multicast bytes: 0 Multicast frames: 0
RX Failures:
 Block check: 0 Framing error: 0 Long frame: 0
 Unknown destination: 0 Data overrun: 0 No system buffer: 0
 No user buffers: 0
>>>
```

4.1.3.7 The NET -IC Command

The NET -IC command initializes the MOP counters for the specified Ethernet port.

Synopsis:

NET -IC EWA0

Example:

```
>>> NET -IC EWA0
>>> NET -S EWA0

Status counts:
ti: 72 tps: 0 tu: 47 tjt: 0 unf: 0 ri: 70 ru: 0
rps: 0 rwt: 0 at: 0 fd: 0 lnf: 0 se: 0 tbf: 0
tto: 1 lkf: 1 ato: 1 nc: 71 oc: 0

MOP BLOCK:
 Network list size: 0

MOP COUNTERS:
Time since zeroed (Secs): 3

TX:
 Bytes: 0 Frames: 0
 Deferred: 0 One collision: 0 Multi collisions: 0
TX Failures:
 Excessive collisions: 0 Carrier check: 0 Short circuit: 0
 Open circuit: 0 Long frame: 0 Remote defer: 0
 Collision detect: 0

RX:
 Bytes: 0 Frames: 0
 Multicast bytes: 0 Multicast frames: 0
RX Failures:
 Block check: 0 Framing error: 0 Long frame: 0
 Unknown destination: 0 Data overrun: 0 No system buffer: 0
 No user buffers: 0
>>>
```

4.1.3.8 The KILL and KILL_DIAGS Commands

The KILL and KILL_DIAGS commands terminate diagnostics that are currently executing.

Note

A serial loopback connector (12–27351–01) must be installed on the COM2 serial port for the KILL_DIAGS command to successfully terminate system tests.

- The KILL command terminates a specified process.
- The KILL_DIAGS command terminates all diagnostics.

Synopsis:

KILL_DIAGS

kill [PID . . .]

Argument:

[PID . . .] The process ID of the diagnostic to terminate. Use the SHOW_STATUS command to determine the process ID.

4.1.3.9 The SHOW_STATUS Command

The SHOW_STATUS command reports one line of information per executing diagnostic. The information includes ID, diagnostic program, device under test, error counts, passes completed, bytes written, and bytes read.

Many of the diagnostics run in the background and provide information only if an error occurs. Use the SHOW_STATUS command to display the progress of diagnostics.

The following command string is useful for periodically displaying diagnostic status information for diagnostics running in the background:

```
>>> while true;show_status;sleep n;done
```

where *n* is the number of seconds between SHOW_STATUS displays.

Synopsis:

SHOW_STATUS

Example:

```
>>> SHOW_STATUS
```

```
    ❶        ❷          ❸         ❹    ❺        ❻             ❼
>>> SHOW_STATUS
 ID       Program      Device       Pass  Hard/Soft Bytes Written  Bytes Read
-------- ------------ ------------ ------ --------- ------------- -------------
00000001         idle system            0    0    0             0             0
0000002d     exer_kid ttal              0    0    0             1             0
0000003d      nettest era0.0.0.2.1     43    0    0          1376          1376
00000045      memtest memory            7    0    0     424673280     424673280
00000052     exer_kid dka100.1.0.6      0    0    0             0       2688512
>>>
```

❶ Process ID

❷ Program module name

❸ Device under test

❹ Diagnostic pass count

❺ Error count (hard and soft): Soft errors are not usually fatal; hard errors halt the system or prevent completion of the diagnostics.

❻ Bytes successfully written by diagnostic

❼ Bytes successfully read by diagnostic

4.1.4 Acceptance Testing and Initialization

You must perform acceptance tests after installing a system or when adding or replacing the following:

- CPU modules
- Memory modules
- Motherboard
- CPU daughterboard
- Storage devices
- EISA options

Perform the following acceptance tests on the previous modules after installing, adding, or replacing them:

1. Run the RBD acceptance tests using the TEST command.
2. If you have added, moved, or removed an EISA option, run the EISA Configuration Utility (ECU).
3. Bring up the operating system.
4. Run DEC VET to test that the operating system is correctly installed. Refer to Section 4.1.5 for information on DEC VET.

4.1.5 DEC VET

Digital's DEC Verifier and Exerciser Tool (DEC VET) software is a multipurpose system maintenance tool that performs exerciser-oriented maintenance testing. DEC VET runs on OpenVMS operating systems, and consists of a manager and exercisers. The DEC VET manager controls the exercisers. The exercisers test system hardware and the operating system.

DEC VET supports various exerciser configurations, ranging from a single device exerciser to full system loading (that is, simultaneous exercising of multiple devices).

Refer to the *DEC Verifier and Exerciser Tool User's Guide* for instructions on running DEC VET.

4.2 Diagnostic Procedures for the HS1CP

Diagnostic procedures for the HS1CP include the following:

- The disk inline exerciser (DILX)
- The tape inline exerciser (TILX)
- The VTDPY utility
- The CONFIG utility
- The configuration menu (CFMENU)
- The Code Patch utility
- The Firmware Licensing System (FLS)
- The Fault Management Utility (FMU)

- The CLONE utility

For more information about these diagnostic procedures, see Chapter 7.

5

Normal Operation

This chapter describes operating conditions and testing procedures for the FDDI Server. Information is included about the HS241, HS280, HS221/121, and HS211/111 systems.

The configurations discussed in this chapter are those set by the system administrator employing user interfaces such as the Command Line Interface (CLI). See Appendix D for configuration rules and restrictions of the subsystem hardware. See Appendix E for information on the CLI.

This chapter provides information on the following topics:

- HS241 description
- HS280 description
- HS221/121 description
- HS211/111 description
- Initialization
- HS1CP front panel
- Connecting a maintenance terminal
- Command Line Interpreter
- Acceptance test
- Virtual terminal
- VAXcluster console system
- Operating systems
- Failover
- Moving devices between HS1CPs
- Moving devices under the same HS1CP
- Command disks
- Server processor configuration and setup
- System bus options
- Motherboard
- PCI bus options
- EISA bus options

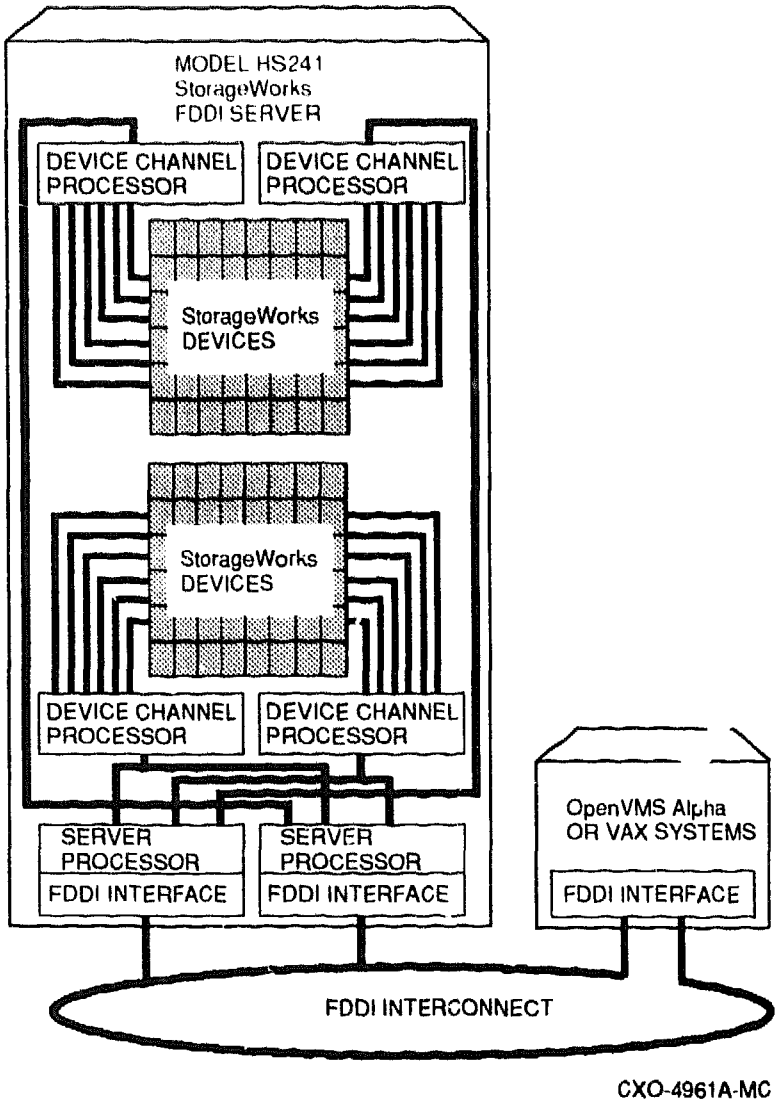
- EISA Configuration utility
- Power supply configurations
- Console port configurations

5.1 HS241 Description

Figure 5–1 shows a conceptual model of the HS241 StorageWorks FDDI Server. Ideal for large networks running "bet-your-business" applications, the HS241 server runs at 4300 I/Os per second and provides high server availability via dual-redundant server processors and four dual-redundant device channel processors. Housed in a single StorageWorks SW800 cabinet, this server can connect up to 72 redundant devices. Maximum HS241 online disk capacity is 309 GB using RZ29 disk drives. Consult the Software Product Description, *HS1CP Device Channel Processor Operating Firmware Version 2.5 for the StorageWorks FDDI Server*, SPD 64.19.01, for a complete list of supported StorageWorks devices. Nearline storage of 10-plus terabytes is possible by adding up to two StorageWorks tape libraries for backups and archiving.

During normal operation, I/O performance is enhanced by balancing the I/O load across the redundant system components in a customer-definable manner. If a component failure occurs, the redundant partner takes over, providing continued service until the failed component can be repaired or replaced. All HS2*xx* redundant StorageWorks FDDI Server components can be hot swapped.

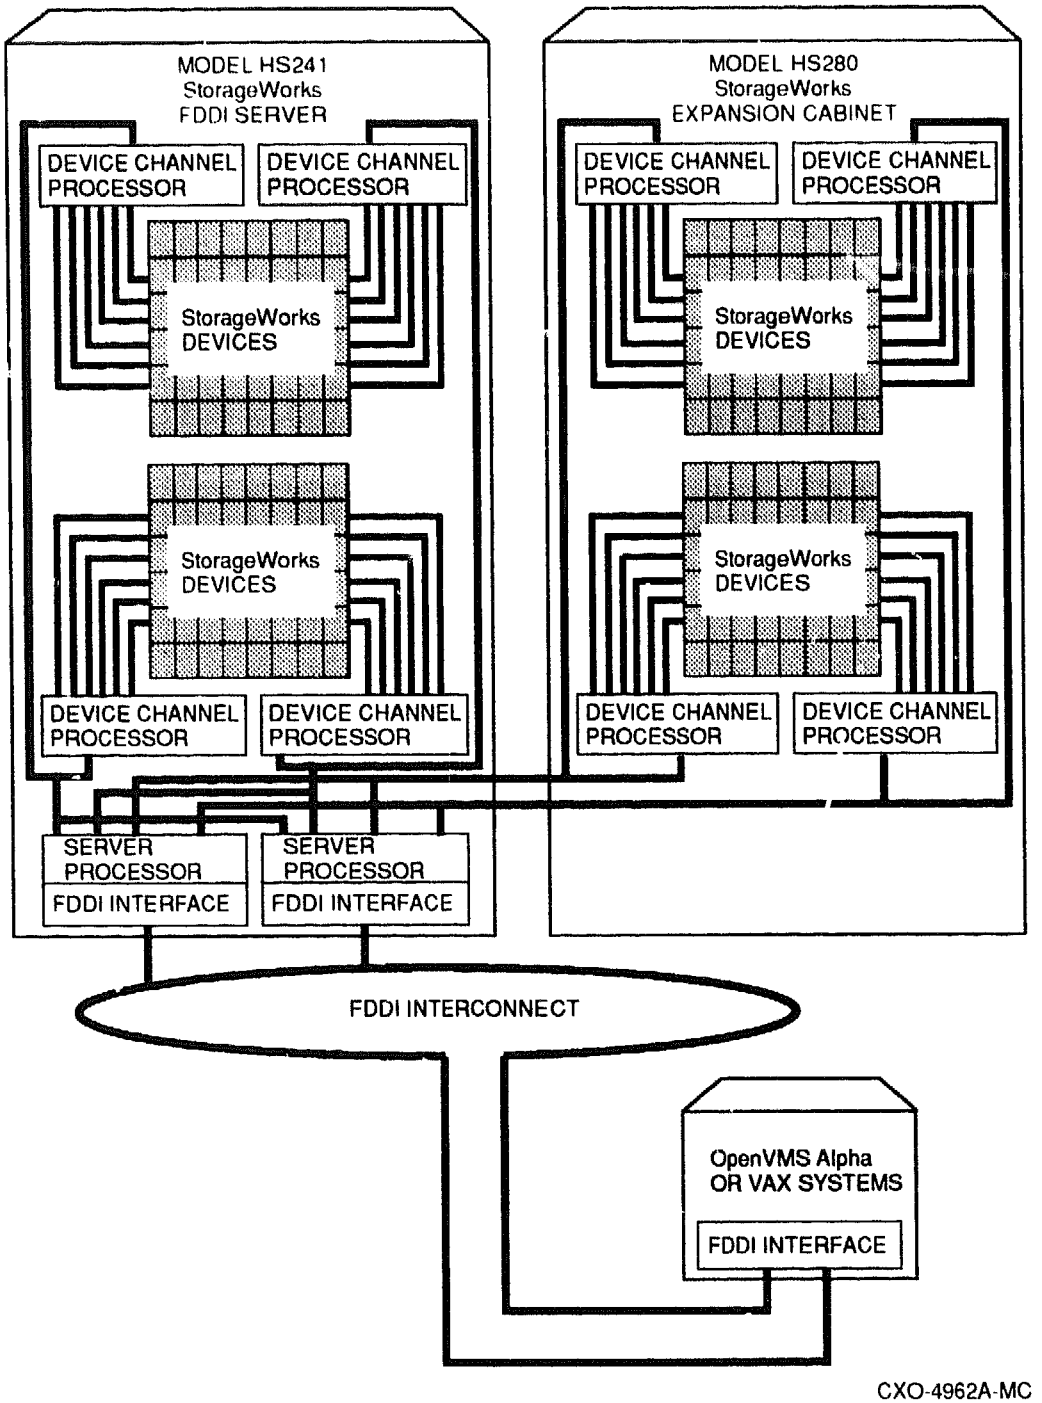
Figure 5–1 HS241 Diagram



5.2 HS280 Description

Double the device support (to 144 devices) and online disk capacity to 619 GB by connecting the HS241 to an optional HS280 expansion cabinet. The HS280 includes a second StorageWorks SW800 cabinet and four additional redundant device channel processors. Redundant power and cooling also are standard at the SW800 cabinet level. Figure 5–2 shows a conceptual model of the HS241 StorageWorks FDDI Server with the HS280 expansion option installed.

Figure 5–2 HS241 with HS280 Diagram



5.3 HS221/121 Description

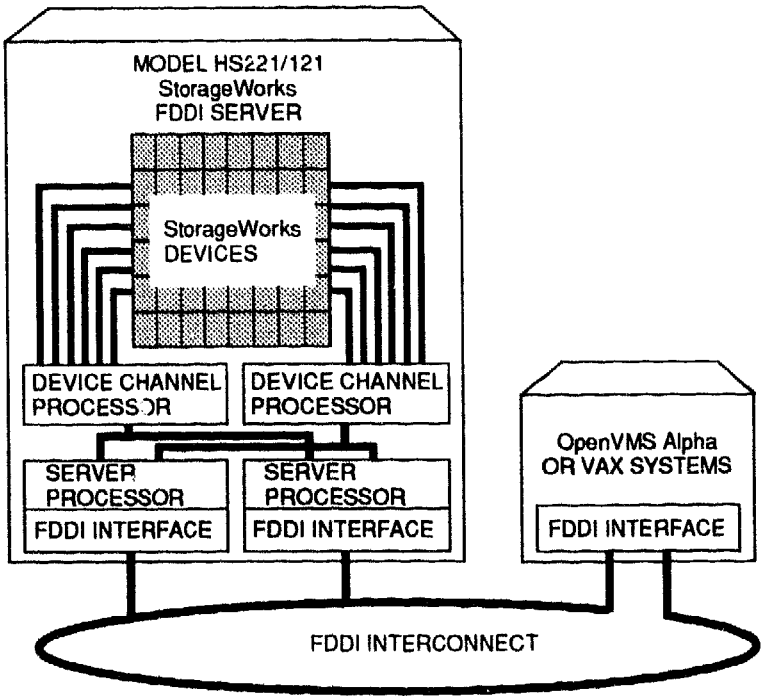
Figure 5–3 shows a conceptual model of the HS221/121 StorageWorks FDDI Server. It is a fully-redundant, high-availability FDDI storage solution offering full protection against any single component failure. These servers are suited to medium-sized operations running "mission-critical" applications. High server availability is assured with fully redundant server processors, two paths to the FDDI interface, two paths to every storage device and standard redundant

power and cooling. The full range of StorageWorks redundant power and cooling features for devices are also available.

The basic HS221/121 StorageWorks FDDI Server provides fully redundant access to up to 36 redundant StorageWorks device connections. Consult the Software Product Description, *HS1CP Device Channel Processor Operating Firmware Version 2.5 for the StorageWorks FDDI Server*, SPD 64.19.01, for a complete list of supported StorageWorks devices. Using the RZ29 disk drive, 154 GB of online, redundant storage per server is available. The model HS221/121 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-definable manner. If a component failure occurs, the redundant partner takes over, providing continued service until the failed component can be repaired or replaced. All HS2*xx* redundant StorageWorks FDDI Server components can be hot swapped.

Figure 5–3 HS221/121 Diagram

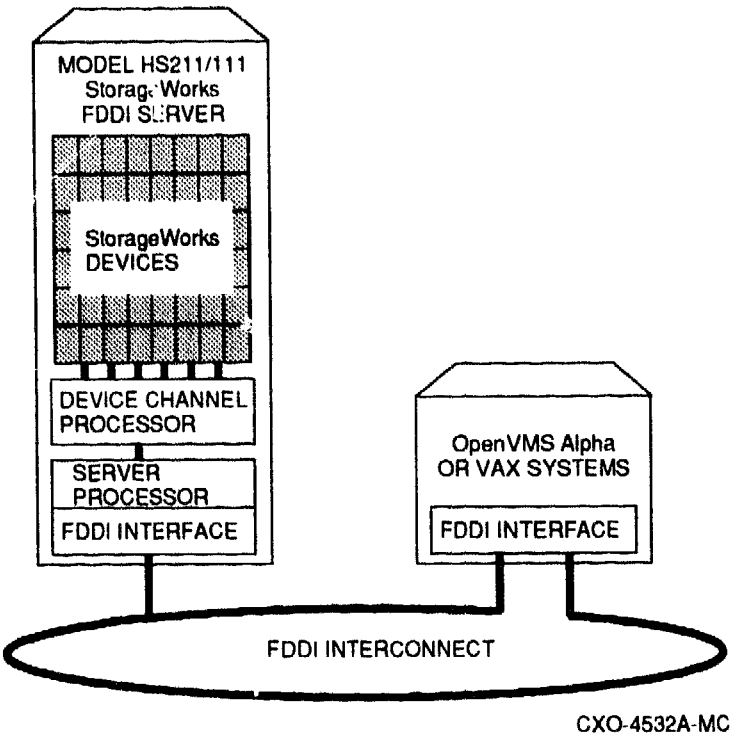


5.4 HS211/111 Description

Figure 5–4 shows a conceptual model of the HS211/111 StorageWorks FDDI Server. Suited to medium-sized operations where you need distributed access and storage management benefits, these servers consist of a single Alpha-based server processor with a FDDI DAS (HS211) or SAS (HS111) network interface and a device channel processor (HS1CP) capable of connecting the server processor to as many as 42 StorageWorks devices. Consult the Software Product Description, *HS1CP Device Channel Processor Operating Firmware Version 2.5 for the StorageWorks FDDI Server*, SPD 64.19.01, for a complete list of supported StorageWorks devices. Using the RZ29 disk drive, up to 180 GB of nonredundant

online storage per server is available. The model HS211/111 StorageWorks FDDI Server and its attached storage devices are housed in a single StorageWorks SW800 cabinet.

Figure 5–4 HS211/111 Diagram



5.5 HS210/110 Description

The HS210/110 contains all the major components of an HS211/111 except for the StorageWorks SW800 cabinet. These options provide for the conversion of an existing SW800 cabinet to a StorageWorks HS211/111 FDDI Server. The *StorageWorks FDDI Server Upgrade Installation Guide* describes the HS210 installation while the *StorageWorks FDDI Server HS110 Upgrade Installation Guide* describes the HS110 installation.

5.6 Initialization

The following sections discuss the operating conditions surrounding initialization of the HS1CP.

5.6.1 HS1CP Initialization

The HS1CP controller initializes after any of the following conditions:

- Power is turned on.
- The firmware resets the HS1CP.
- The operator presses the green Reset button on the front panel of the HS1CP.
- The host clears the HS1CP.
- The program card is removed.

See Chapter 7 for a description of the initialization of both the controller and its cache module. (The process is described in Chapter 7, because some of the initialization diagnostics are available as a controller self-test function for the operator.)

5.6.2 Dual-Redundant Configuration Initialization

The HS1CPs in a dual-redundant configuration run the same initialization sequence that is described in Chapter 7, except that they exchange signals during their individual initialization sequences. The first signal occurs after one HS1CP starts initializing. The signal informs the other HS1CP that an initialization has occurred. This way, the other HS1CP does not assume that the initializing HS1CP has malfunctioned and does not attempt to disable it.

5.6.3 Subsystem Initialization

Full StorageWorks subsystem initialization takes place when the subsystem is switched on for the first time. If a reset occurs due to one of the following conditions, a subset of the initialization sequence is run:

- A partial or complete power failure
- Equipment failure
- An error condition

A complete StorageWorks subsystem initialization includes the following:

1. When the subsystem is turned on, all shelves in the subsystem are reset. When reset, the entities in the shelves (including storage devices, HS1CPs, and cache modules) run their individual initialization and self-test sequences.
2. During initialization, the HS1CP interrogates the entities with which it has connections, including other HS1CPs in the subsystem.
3. When the initialization sequence in each of the entities is completed, the HS1CP begins data transfer and other operations with the host.

5.7 HS1CP Front Panel

You should use the front panel of the HS1CP to reset the controller, quiesce the SCSI–2 device buses attached to the HS1CP, and interpret error conditions that result in LED error codes.

5.8 Connecting a Maintenance Terminal

A maintenance terminal is an EIA–423 compatible terminal connected directly to the maintenance port on the front of the HS1CP. You do not need a maintenance terminal for normal HS1CP operations. However, you must connect a maintenance terminal for initial controller configuration. Thereafter, use either a maintenance terminal or a host (virtual) terminal to communicate with the HS1CP.

Follow this procedure to connect a maintenance terminal:

1. Make sure the power switch on the back of the terminal is in the Off (○) position.
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 maintenance port on the front of the HS1CP.

4. Set your terminal's communication setup to 9600 baud, 8 data bits, 1 stop bit, and no parity. Refer to your terminal documentation for terminal setup instructions.

5.9 Command Line Interpreter

The Command Line Interpreter (CLI) is the user interface to the controller. The CLI allows you to set storage and controller configurations through commands. The following sections explain how to use the CLI, and how it defines and modifies configurations. A detailed description of CLI commands is provided in Appendix E.

5.9.1 Accessing the CLI

You can access the CLI through a maintenance terminal (see Section 5.8) or through a virtual terminal (see Section 5.11).

To access the CLI through a maintenance terminal, connect the terminal and press the Return key. You must use a maintenance terminal to set the controller initial configuration. This is because a new controller arrives with an invalid ID, and host ports initially off.

Thereafter, you can use a virtual (host) terminal to modify the configuration. The method of establishing the virtual terminal connection varies depending on your operating system and interface. The following command connects a host terminal to the CLI:

Note

The following command requires the DIAGNOSE privilege. Also, the controller SCS node name *must* be specified.

```
$ SET HOST/DUP/SERVER=MSCP$DUP/TASK=CLI/
  LOG=CONFIGURATION.INFO SCS_nodename
```

where *SCS_nodename* is the node name of the controller. The qualifier /LOG=CONFIGURATION.INFO creates a log file of your CLI session. Reviews allow you to remember such information as how you configured your subsystem, results from diagnostics, and so on.

Note

Your CLI prompt may be factory set to reflect your controller model. This manual uses the designation CLI> to identify the CLI prompt.

Appendix E provides details on how to change the prompt.

5.9.2 Exiting the CLI

When exiting the CLI, keep the following guidelines in mind:

- If you are using a maintenance terminal, you cannot exit the CLI. Entering the EXIT command restarts the CLI and redisplays the controller type and any last fail error information,

- If you are using the DUP connection, enter the following command to exit the CLI and return to the host:

```
CLI> EXIT
```

5.9.3 Command Sets

The CLI consists of the following command sets:

- Failover commands

  Failover commands support dual-redundant HS1CP configurations.

- HS1CP commands
  - Set and show the basic HS1CP parameters.
  - Set the HS1CP ID (CI or DSSI node number or SCSI target ID).
  - Set the resident terminal characteristics.
  - Restart the HS1CP.
  - Run resident diagnostics and utilities (see Chapter 7).
- Device commands

  Device commands specify and show the location of *physical* SCSI–2 devices attached to the HS1CP. Devices locations are specified using their SCSI PTL designation.

- Storageset commands

  Storageset commands add, modify, rename, and show storagesets (stripesets, mirrorsets, and RAIDsets). These commands also apply (to some extent) to the **spareset** and **failedset**.

- Logical unit commands

  Logical unit commands add, modify, and show logical units built from devices and storagesets.

- Utility commands

  The RUN command invokes various exercisers and utilities, such as DILX, TILX, and CLONE.

5.9.4 Initial Configuration (Nonredundant HS1CP)

After installing a nonredundant HS1CP, use the CLI to define its parameters in the following order (from a maintenance terminal).

1. Enter the following command to set a valid HS1CP ID:

   ```
   CLI> SET THIS_CONTROLLER ID=n
   ```

   where *n* is the one-digit DSSI node number (0 through 7). Each HS1CP DSSI node number must be unique among all devices on its DSSI interconnect.

2. Enter the following command to set the SCS node:

   ```
   CLI> SET THIS_CONTROLLER SCS_NODENAME="xxxxxx"
   ```

   where *xxxxxx* is a one- to six-character alphanumeric name for this node. The node name must be enclosed in quotes with an alphabetic character first. Each SCS node name must be unique within its VMScluster.

3. Enter the following command to set the MSCP **allocation class**:

   ```
   CLI> SET THIS_CONTROLLER MSCP_ALLOCATION_CLASS=nn
   ```

   where *nn* is 0 through 255.

4. Enter the following command to set the TMSCP allocation class:

   ```
   CLI> SET THIS_CONTROLLER TMSCP_ALLOCATION_CLASS=nn
   ```

   where *nn* is 0 through 255.

Note

Always restart the HS1CP after setting the ID, SCS node name, or allocation classes.

5. Restart the HS1CP either by pressing the green Reset button on the front panel or by entering the following command:

   ```
   CLI> RESTART THIS_CONTROLLER
   ```

6. Enter the following command to verify the preceding parameters were set:

   ```
   CLI> SHOW THIS_CONTROLLER FULL
   ```

CAUTION

Unless a trilink connector is installed on the HS1CP, you must disconnect power from all devices on a DSSI bus before making/breaking a host connection on the front panel of the HS1CP.

Without the trilink connector installed, if you accidentally short pins during aligning and inserting/removing a connector, you risk blowing the fuses of *all* members on the DSSI bus. With the insulator installed on the trilink connector, you can make or break a connection with power applied on the host bus. However, Digital recommends using care *at all times* when inserting or removing connectors, even if the trilink connector has been installed.

Note

Be aware that the trilink connector is the interface between the HS1CP and host cable connections. In a power-on situation, you must work around cable or terminator connections to the SCSI cables without removing them.

Some cable and terminator connectors do not provide enough access to the trilink screws for you to install the trilink (without first removing the cable and/or terminator). In these cases, you must disconnect power from all bus members and remove cables and terminators before installing the trilink.

7. Connect the host port cable to the front of the HS1CP (see Chapter 8).

8. Enter the following command to enable the host port path:

   ```
   CLI> SET THIS_CONTROLLER PATH
   ```

5.9.5 Initial Configuration (Dual-Redundant HS1CPs)

In a dual-redundant configuration, you can use one terminal to set both HS1CP configurations. After installing both HS1CPs, use CLI commands to define configuration parameters in the following order (from a maintenance terminal connected to one HS1CP):

Note

Use Table 5–1 to identify the appropriate HS1CP device channel processor with the appropriate prompt and SCSI ID.

Table 5–1 Standard HS241/HS280 HS1CP Device Channel Processor Identification

BA350–M Shelf	Slot 7	SCSI ID	Slot 6	SCSI ID
HS241 Front	HS1CP1	1	HS1CP2	2
HS241 Rear	HS1CP3	3	HS1CP4	4
HS280 Front	HS1CP5	1	HS1CP6	2
HS280 Rear	HS1CP7	3	HS1CP8	4

1. Enter the following command to set a valid HS1CP ID:

   ```
   CLI> SET THIS_CONTROLLER ID=n
   ```

   where *n* is a one-digit DSSI node number (0 through 7). Each HS1CP DSSI node number must be unique among all devices on its DSSI interconnect.

2. Enter the following command to set the SCS node:

   ```
   CLI> SET THIS_CONTROLLER SCS_NODENAME="nn"
   ```

   where *nn* is a one- to six-character alphanumeric name for this node. The node name must be enclosed in quotes with an alphabetic character first. Each SCS node name must be unique within its VMScluster.

3. Enter the following command to set the MSCP allocation class:

   ```
   CLI> SET THIS_CONTROLLER MSCP_ALLOCATION_CLASS=nn
   ```

   where *nn* is 1 through 255.

   Digital recommends providing a unique allocation class value for every pair of dual-redundant HS1CPs in the same cluster.

4. Enter the following command to set the TMSCP allocation class:

   ```
   CLI> SET THIS_CONTROLLER TMSCP_ALLOCATION_CLASS=nn
   ```

   where *nn* is 1 through 255.

CAUTION

The SET FAILOVER command establishes controller-to-controller communication and copies configuration information. Enter this command on one controller only. COPY=*configuration-source* specifies where the *good* configuration data is located. *Never* blindly specify SET FAILOVER.

Know where your good configuration information resides before entering the command.

---

5. Enter the following command to copy parameters to the other controller (the one *not* connected to):

```
CLI> SET FAILOVER COPY=THIS_CONTROLLER
```

Note

Always restart the HS1CPs after setting the ID, SCS node name, or allocation classes.

---

6. Restart both HS1CPs either by pressing the green reset buttons on the front panel or by entering the following commands:

```
CLI> RESTART OTHER_CONTROLLER CLI> RESTART THIS_CONTROLLER
```

7. Enter the following commands to verify the preceding parameters were set:

```
CLI> SHOW THIS_CONTROLLER FULL
CLI> SHOW OTHER_CONTROLLER FULL
```

CAUTION

Unless a trilink connector is installed on the HS1CP, you must disconnect power from all devices on a DSSI bus before making or breaking a host connection on the front of the HS1CP.

Without the trilink connector installed, if you accidentally short pins during aligning and inserting or removing a connector, you risk blowing the fuses of *all* members on the DSSI bus. With the insulator installed on the trilink connector, you can make or break a connection with power applied on the host bus. However, Digital recommends using care when inserting or removing connectors *at all times*, even if the trilink connector is installed.

Note

Be aware that the trilink connector is the interface between the HS1CP and two DSSI cable connections. It is the *cable connector* that you can safely remove and replace when a trilink connector is installed. In a power-on situation, you must work around any DSSI cable or terminator connections to the trilink without removing them.

Some cable and terminator connectors do not provide enough access to the trilink screws for you to remove the connector (without first removing the cable and/or terminator). In these cases, you must disconnect power from all bus members and remove cables and terminators before removing the trilink.

---

8. Connect the host port cables to the front of the HS1CPs (see Chapter 8). You can connect two dual-redundant HS1CPs to different DSSI buses.

9. Enter the following commands to enable the host port path:

```
CLI> SET THIS_CONTROLLER PATH
CLI> SET OTHER_CONTROLLER PATH
```

5.9.6 Configuring Storage Devices

To automatically configure devices on the HS1CP, use either the CONFIG or CFMENU utility. These utilities are described in Chapter 7.

For manual configuration, the following steps add devices, storagesets, and logical units. Use the CLI to complete these steps so that the host does not recognize the storage device. (These steps can be run from a virtual terminal.)

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

   ```
   CLI> ADD device-type device-name SCSI-location
   ```

   For example:

   ```
   CLI> ADD DISK DISK100 1 0 0
   CLI> ADD TAPE TAPE510 5 1 0
   CLI> ADD CD--ROM CD--ROM0 6 0 0
   ```

   where:

   device-type is the type of device to be added. This can be DISK, TAPE, LOADER, OPTICAL, or CD–ROM.

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

   SCSI-location is the PTL for the device. When entering the PTL, separate the three numbers with spaces.

2. Add the storagesets for the devices. Storagesets include stripesets, mirrorsets, and RAIDsets.

3. Enter the following command to initialize the containers (devices, storagesets, or both) prior to adding logical units to the configuration.

   CAUTION

   The INITIALIZE command destroys all data on a **container**. See Appendix E for specific information on this command.

   ```
   CLI> INITIALIZE container-name
   ```

   where *container-name* is a device or storageset that becomes a part of a unit.

   When initializing a single-device container:

   - If NOTRANSPORTABLE (the default) is specified when the device was added, a small amount of disk space was made inaccessible to the host and used for **metadata**. The metadata is now initialized.
   - If TRANSPORTABLE was specified, any metadata on the device is destroyed. (See Appendix E for more information on metadata and when INITIALIZE is required.)

4. Add the units that use either the devices or the storagesets by entering the following command:

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

where:

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

container-name identifies the device or the storageset.

5.10 Acceptance Test

After you install and configure your HS1CP, follow the guidelines in this section to perform an acceptance test on your subsystem.

1. Turn your system on. This resets all shelves and starts the spin-up cycle on devices within the shelves. This includes the initialization (diagnostics) on the HS1CPs and device self-tests.
2. Run DILX using the default answers to the test questions (see Chapter 7). This utility tests all disk devices in your subsystem.
3. Run TILX using the default answers to the test questions (see Chapter 7). This utility tests all tape devices in your subsystem.

5.11 Virtual Terminal

After setting initial controller parameters through a maintenance terminal, you can access the controllers from a virtual host terminal through a DUP connection.

Establishing a virtual terminal session under the OpenVMS VAX and OpenVMS Alpha operating systems (SET HOST/DUP) requires the FYDRIVER. The following error indicates that the FYDRIVER is not loaded:

```
%HSCPAD-F-DRVNOTLOAD, FYDRIVER not loaded
-SYSTEM-W-NOSUCHDEV, no such device available
```

If you receive this message, load the FYDRIVER as follows:

- For OpenVMS VAX:

```
$ MCR SYSGEN
SYSGEN> LOAD SYS$LOADABLE_IMAGES:FYDRIVER
SYSGEN> CONNECT FYA0 /NOADAPTER
SYSGEN> EXIT
$
```

- For OpenVMS Alpha:

```
$ MCR SYSMAN
SYSMAN> IO CONNECT FYA0 /NOADAPTER/DRIVER=SYS$FYDRIVER
SYSMAN> EXIT
$
```

Once FYDRIVER is loaded, you should make the virtual terminal connection as follows:

```
$ SET HOST/DUP/SERVER=MSCP$DUP/TASK=CLI/LOG=CONFIGURATION.INFO SCS_nodename
```

5.12 VAXcluster Console System

You can run VAXcluster Console System (VCS) with any HS1CP. If you are unfamiliar with VCS, refer to the *VCS Software Manual* for instructions. You can attach a VCS to the maintenance terminal port via a DECserver™ or terminal.

5.13 Operating Systems

The following sections describe particulars associated with host operating systems, which can help in understanding and servicing the HS1CP controllers.

The primary operating system that supports the HS1CP controller is the OpenVMS operating system.

Refer to your firmware release notes for restrictions and updates regarding operating system support.

Although certain specifics regarding operating systems are covered here, see the *StorageWorks Array Controllers HS Family of Array Controllers User's Guide* for complete information on operating system support.

Other Conditions (OpenVMS)

The following conditions and recommendations also apply to controllers running under the OpenVMS operating system:

- Storageset size

  The OpenVMS VAX operating system maximum capacity restriction for file-structured volumes, 16,777,216 blocks (about 8.5 GB), remains in effect for operating system versions prior to Version 6.0.

- The CLUSTER_SIZE qualifier for large devices or storagesets

  Digital recommends that the formula displayed by the OpenVMS HELP DEVICE INIT/CLUSTER_SIZE command be used to determine the proper OpenVMS file system cluster size. Using too small a file system cluster size can prevent some of the device or storageset capacity from being accessed; too large a cluster size usually wastes storage capacity by allocating large blocks of storage for small files.

- Shadowset operation

  In OpenVMS VAX operating system versions prior to Version 6.0, timed-out I/O requests to shadowset members can lead to member disks attached to controllers being dropped from shadowsets. In some cases, this could lead to host crashes.

  To avoid this possibility, Digital recommends changing the value of the SYSGEN parameter SHADOW_MBR_TMO to at least 120 (seconds) for systems running operating system versions earlier than Version 6.0. (Be aware that your system may pause temporarily during the 120 second interval.) Version 6.0 of the OpenVMS VAX operating system avoids this problem by retrying timed-out operations to shadowset members several times.

- PAPOLLINTERVAL and PANUMPOLL parameters

  Digital recommends that the SYSGEN parameters PAPOLLINTERVAL and PANUMPOLL be set such that all nodes in the cluster are polled within 30 seconds or less. Failure to set this value can result in MSCP command timeouts. The default values are set to poll 16-node clusters every 5 seconds and 32-node clusters every 10 seconds.

5.14 Failover

Failover occurs when one HS1CP fails in a dual-redundant configuration. To support failover, the following information is shared between the two HS1CPs:

- Physical device PTL configurations
- Storageset names
- Logical unit definitions

Prior to failover, all resources are considered unbound to a particular HS1CP, until a logical unit is brought online by the host through (one of) the HS1CPs. At this point, all containers used by the logical unit become accessible solely through the one HS1CP.

In a failover configuration, all commands are shared between the two HS1CPs except the following:

SET THIS_CONTROLLER
SET OTHER_CONTROLLER
SHOW THIS_CONTROLLER
SHOW OTHER_CONTROLLER
RESTART THIS_CONTROLLER
RESTART OTHER_CONTROLLER
SHUTDOWN THIS_CONTROLLER
SHUTDOWN OTHER_CONTROLLER

In these cases, the command is directed to the correct HS1CP:

- THIS_CONTROLLER refers to the HS1CP to which the terminal is connected.
- OTHER_CONTROLLER refers to the other HS1CP in the dual-redundant pair.

5.14.1 Setting Failover

To place two HS1CPs into failover configuration, enter the following command:

```
CLI> SET FAILOVER COPY=configuration-source
```

where *configuration-source* is either THIS_CONTROLLER or OTHER_CONTROLLER, depending on where the "good" copy of device configuration information resides.

CAUTION

Digital recommends that the HS1CPs be set for failover before any device configuration commands are entered. Then, as devices, storagesets, and units are added to one HS1CP's configuration, they are added automatically to the other HS1CP's configuration.

Given two HS1CPs, it is possible to fully configure one HS1CP, and then enter the SET FAILOVER command, but if the wrong *configuration-source* has been specified, all device configuration information is lost (overwritten). *Never* blindly specify SET FAILOVER. Know where your good configuration information resides before entering the command. (A considerable amount of work and effort could be lost easily by overwriting good information.)

Note

Due to the amount of information that must be passed between the two HS1CPs, the SET FAILOVER command could take up to 1 minute to complete.

When setting dual-redundant HS1CPs for failover, make sure the target HS1CP (the controller you are copying configuration to) has no cache errors or unwritten cache data. Delete any remaining, configured units on the target HS1CP to verify there are no cache errors before entering the SET FAILOVER command. Deleting units from the target HS1CP does not create problems because, after setting failover, you do not need to access those units anyway. (The target HS1CP accesses only the copied units.)

5.14.2 Failing Over

A failed or unresponsive HS1CP in a dual-redundant configuration is disabled by its companion HS1CP. The functioning HS1CP sends a signal to the other HS1CP to induce failover. The functioning HS1CP assumes control of the storage devices that were online to the disabled HS1CP. Maintenance now can take place on the failed HS1CP.

Failover normally completes in 30 seconds or less (15 seconds or less for three-port HS1CPs). If there is no outstanding drive I/O activity at the time of HS1CP failure, failover should require substantially less than 30 seconds. If drive I/O is in progress at the time of failure, the surviving HS1CP must reset any SCSI buses with outstanding I/O. These bus resets can require up to 5 seconds per port to complete.

Whenever you need to revive an HS1CP that is disabled, you must enter the following command from a terminal connected to the functioning HS1CP:

```
CLI> RESTART OTHER_CONTROLLER
```

Then initialize the controller by pressing the Reset button on the controller to be revived.

You can test failover by removing the program card from one of the HS1CPs. The other HS1CP assumes service to the dormant HS1CP's devices until you reinsert the program card and reinitialize/restart the HS1CP.

5.14.3 Exiting Failover

You rarely should force a dual-redundant HS1CP pair out of the failover configuration. However, the circumstances when you can do this safely are described below.

5.14.3.1 Before Failover Occurs

If you have two normally operating dual-redundant HS1CPs set for failover, and you want to remove one HS1CP for use somewhere else, you must shut down that HS1CP (see Chapter 8). Shutting down one HS1CP forces failover to occur, which preserves access to your entire configuration through the surviving HS1CP.

5.14.3.2 After Failover Occurs

After one HS1CP in your dual-redundant pair has failed or been shut down, the surviving HS1CP services your entire configuration. To take the HS1CPs out of the failover configuration, enter the following command:

```
CLI> SET NOFAILOVER
```

You must consider the following before entering the SET NOFAILOVER command or removing one HS1CP:

- You do not need to SET NOFAILOVER unless you must make configuration changes before you can replace the failed HS1CP.
- If the surviving HS1CP and its cache module are functioning normally, you can SET NOFAILOVER without special preparation.
- If the surviving HS1CP is running with a low write-back cache battery (enter SHOW THIS_CONTROLLER to check the battery), you must replace/recharge the battery before entering SET NOFAILOVER and before removing the failed HS1CP's cache module.

Entering SET NOFAILOVER removes the HS1CP from the failover configuration (as well as the other HS1CP, if it is reachable) and causes the other HS1CP to shutdown. You now can make the configuration changes under the surviving HS1CP.

5.14.4 Failover Setup Mismatch

During failover mismatch, one HS1CP functions while the second HS1CP does not recognize any devices. Although it is rare, a failover mismatch can occur during the following scenarios:

- If the HS1CPs initialize at *exactly* the same time, one HS1CP can be set for failover while the other is not.
- If one HS1CP is running (operating normally) when the second HS1CP is initialized, mismatch can occur. For example, this can happen after one HS1CP has undergone maintenance.

To correct a failover mismatch, stop all processes on the devices for both HS1CPs. Then enter the following commands to determine which HS1CP has the desired, good configuration information:

```
CLI> SHOW UNITS
CLI> SHOW STORAGESETS
CLI> SHOW DEVICES
```

After deciding on one of the two configurations, use the SET FAILOVER command to copy the good information from one HS1CP to the other.

5.14.5 Preferred Path

Use the HS1CP preferred path command (SET *unit* PREFERRED_PATH) to balance units between HS1CPs for better performance in a dual-redundant configuration. Using preferred paths maintains unit-to-HS1CP assignments when dual-redundant pairs are restored from the failover condition.

Once established, preferred paths are activated automatically, and remain stored in HS1CP NVMEM.

5.15 Moving Devices Between HS1CPs

Moving devices from one HS1CP to another is supported under the following conditions:

- **Nontransportable devices**

  Under normal operation, the HS1CP makes a small portion of a disk inaccessible to the host and uses this area to store metadata. Metadata improves error detection and media defect management. Devices utilizing metadata are called **nontransportable**. Initializing a device that is set as nontransportable places or resets metadata on the device.

  When bringing other HS1CP (nontransportable) devices to an HS1CP subsystem, add the device to your configuration using the ADD command. Do not initialize the device or you will reset/destroy forced error information on the device.

  When adding devices, the HS1CP firmware verifies that metadata is present. If in doubt, try to add the device so that the HS1CP checks for metadata. If an error occurs stating that there is no metadata, initialize the device before adding it.

  A nontransportable device is interchangeable with an HSC™ K.scsi[1] module or another HS controller subsystem. Nontransportable devices are MSCP compliant and support forced error.

- **Transportable devices**

  A **transportable** feature is provided for transfer of devices between non-HS1CP controller systems and HS1CP controller arrays. Transportable devices do not have metadata on them, and initializing a device after setting it as transportable destroys metadata (if any) on the device.

  Before moving devices from an HS1CP subsystem to a non-HS controller system, delete the unit associated with the device and set the device as transportable. Then, initialize the device to remove any metadata.

  When bringing non-HS1CP devices to an HS1CP subsystem, initialize the device after setting it transportable, then copy the data on the device to another, nontransportable, unit. Then, reinitialize the device after setting it nontransportable (thereby putting metadata on the device). You *must* initialize these devices because they can contain intact metadata blocks, which can "fool" the controller into attempting to run with the device.

CAUTION

Do not keep any device set as transportable on an HS1CP controller subsystem. Doing so sacrifices forced error support on all units attached to the device. This is mandatory for HBVS and improving data integrity on the entire array.

[1] For purposes of setting transportable/nontransportable devices, the HSC K.scsi controllers are compatible with HS1CPs. However, the geometry of disk devices under these two subsystems is slightly different. To correct the geometry, always use the OpenVMS ANALYZE/DISK/REPAIR command on a disk after transfer between these two subsystems.

A transportable device is interchangeable with any SCSI interface that does not utilize the device metadata (for example, a VAX workstation, an SZ200, or a PC). Transportable devices are not MSCP compliant, cannot have write-back caching enabled, cannot be members of a shadowset or storageset, and do not support forced error. An HS1CP error (see Chapter 6) occurs if the operating system attempts to write forced error information to a transportable device.

Transportable/nontransportable device support is summarized in Table 5–2.

Table 5–2 Transportable and Nontransportable Devices

Media Format	VAX or Alpha Workstation	HSC K.scsi	HSD05	HS1CP
Transportable	Yes	No	Yes	Yes
Nontransportable	No	Yes	No	Yes

5.16 Moving Devices Under the Same HS1CP

CAUTION

Do not use a controller failure situation as an opportunity to move devices or otherwise reconfigure your subsystem. Doing so prevents the controller from communicating with its units once the fault is corrected.

You can physically relocate some or all of a storageset's member devices according to the following procedure. (This procedure also applies to reduced RAIDsets, but you must remember to add the RAIDset as REDUCED when you recreate it.)

CAUTION

Label any devices before you remove them. If you lose track of the storageset members at any point during this procedure, you must attempt to restore the storageset by guessing where its members are installed. There is currently no way to retrace your steps using the controller or HS1CP operating firmware.

1. Make note of all devices comprising the storageset. Digital recommends marking them after using the LOCATE command to find all storageset members. Note the order of the devices in the storageset—which disk is the first member, which disk is the second member, and so on.
2. Delete the unit associated with the storageset.
3. Delete the storageset.
4. Delete the devices to be moved.
5. Move the devices to the new port/target/LUN (PTL) locations.
6. Add the devices using the new PTL location.
7. Recreate the storageset, making sure to add the members in their original order. Make sure you create it from the exact, original set of devices.

CAUTION

Do *not* initialize the storageset. Doing so destroys its data.

8. Recreate the unit.

The following example shows the unit "D100" made of stripeset "STRIPE0." "STRIPE0" has member disks at PTLs 200 and 210. The member at PTL 210 can be relocated to PTL 300 as follows:

```
CLI> DELETE D100
CLI> DELETE STRIPE0
CLI> DELETE DISK210
```

(Move the disk to PTL 300.)

```
CLI> ADD DISK DISK300 3 0 0
CLI> ADD STRIPESET STRIPE0 DISK200 DISK300
CLI> ADD UNIT D100 STRIPE0
```

5.17 Command Disks

A special storage configuration is available for communication with generic SCSI devices. The concept centers around a virtual unit, called a **command disk**, that allows an MSCP-based host to control any physical SCSI device.

This feature uses the MSCP read and write operations as a transport for SCSI commands. The command disk appears to the host as another disk behind the controller. However, the command disk resides entirely in approximately 10 KB of controller memory. The command disk exchanges SCSI commands, SCSI status, and data with the SCSI device. The setup is necessary in order for MSCP to transfer the SCSI information required when performing anything more than simple device reading or writing.

5.17.1 Uses for Command Disks

Although a command disk allows you to perform a number of operations, including downloading of microcode into a SCSI device's processor memory, the example presented in this manual is that of a tape library.

A tape library contains a standard tape drive (or drives) plus a robotic mechanism for selecting and loading cartridges from a pool of tapes. (Such mechanisms are also known as "jukebox" style **loaders**, or medium changers.) The loader is a good application for a command disk because SCSI commands are needed to control the robotic movement within the device.

In some tape library devices, the actual drive has one logical unit number (LUN), and the loader has a different LUN.

5.17.2 Creating a Command Disk

You must configure a command disk at the controller and the host level.

5.17.2.1 Controller Setup

The following process explains what you must do at the controller level to create a command disk for a loader. Note that you can assign only one SCSI device to each command disk.

1. Add the tape device serviced by the loader robot.
2. Add the tape device as a unit visible to the host.
3. Add the loader mechanism using its PTL location.
4. Add a **passthrough** container, the bridge between the physical loader and the virtual unit.
5. Add the passthrough container as virtual unit visible to the host.

The following example shows the CLI commands necessary to add a tape library device to your configuration using a command disk. Note that the PTL addresses for devices in your storage subsystem are different.

```
CLI> ADD TAPE320 3 2 0 ❶
CLI> ADD UNIT T322 TAPE320 ❷
CLI> ADD LOADER LDR321 3 2 1 ❸
CLI> ADD PASSTHROUGH PASS1 LDR321 ❹
CLI> ADD UNIT D332 LDR321 ❺
```

The commands in the example do the following:

❶ Adds the tape drive at PTL 3 2 0 to the list of known tape drives with the name TAPE320.

❷ Creates a host-visible unit called T322 that consists of TAPE320.

❸ Adds the tape loader at PTL 3 2 1 to the list of known loaders with the name LDR321. Note that for this tape library device, both the drive from step 1 and the loader have the same Port and Target number; only the LUN is different.

❹ Creates a passthrough container (command disk) associated with LDR321 to allow the host direct access to the loader. The passthrough container is called PASS1.

❺ Creates a host-visible unit called D332 that consists of passthrough PASS1.

The host operating system uses *disk* unit 332 to exercise loader functions, and *tape* unit 332 for backups, and so forth.

5.17.2.2 Host Setup

To communicate with a loader through a command disk, you must have the following available at the host level:

- You must install the OpenVMS operating system JU driver. When you install JU driver, you must assign a JU device to the correct unit number (the unit number established through the CLI when you created the command disk):

```
Beginning installation of JUDRIVER V1.4 at 15:54
    .
    .
    .
For each command disk on your system, enter the DU device
designation.  Enter Ctrl/Z when there are no more command disks.
    .
    .
    .
* Enter the first DU device designation: $6$DUA622 ❶
* Enter the next DU device or Ctrl/Z if done:  Exit
%JUDRIVER-I-DEFINE, defining device JUA0 ❷
%JUDRIVER-I-LINK, Linking JUDRIVER ...
%JUDRIVER-I-LINK, Linking JUDRIVER_INIT ....
                                        .
                                        .
```

❶ In this case, the loader mechanism resides under unit number D622.

❷ The command disk is linked to JU device "JUA0."

- You must have special host-level application software[2] for controlling the particular loader. The application software can vary and is not provided with the HS1CP or firmware. The application software uses the JU driver to manage the command disk protocol, controlling the loader in conjunction with the drives served by it.

Note

If the virtual unit in your configuration is not linked to a JU device, you receive the following message when you use your application software to output SCSI commands for the loader:

```
%SYSTEM-W-NOSUCHDEV, no such device available
```

5.17.3 Communicating with a Command Disk

The OpenVMS host's DU driver recognizes command disks as normal disk devices. However, because the true purpose of command disks is to provide a conduit for passing raw SCSI command descriptor blocks (CDBs) through to the underlying loader, communication operates as follows:

1. The host application that is the source of the SCSI commands issues CDBs to the JU driver similar to when communicating with a native SCSI device.
2. The JU driver packages these SCSI commands using the command disk protocol, and queues them through the DU driver to the controller.
3. Controller firmware processes and issues the commands to the loader's SCSI bus, completing the illusion that the host is directly attached to the loader's bus.

[2] Your application software and JU driver probably are packaged together.

5.17.4 Performance

Command disks are most useful for loaders and other low-speed or infrequent uses, and are not optimized for speed. Implementing significant drive data movement through a command disk is not recommended because multiple MSCP operations are necessary for each SCSI command. Large data movement through command disks requires partitioning into smaller transfers because it is undesirable to allocate large buffers for the amount of time typically required by command disk applications.

5.17.5 Maintenance

Because command disks are essentially conduits for host commands, there is no controller-based diagnostic, test, or error notification capability for a command disk. However, the host application should be able to test the integrity of the command path with its own echo testing, if necessary.

5.18 Server Processor Configuration and Setup

This section provides configuration and setup information for the server processor components.

5.18.1 System Firmware

The console firmware provides the data structures and callbacks available to booted programs defined in both the SRM and and ARC standards.

SRM Command Line Interface

Systems running OpenVMS access the system using the Command Line Interface (CLI). The CLI is a UNIX® style shell that provides a set of commands and operators, as well as a scripting facility. The CLI allows you to configure and test the system, examine and alter system state, and boot the operating system. (See Appendix E for more specific CLI information).

The SRM console prompt is >>>.

Several system management tasks can be performed only from the SRM console CLI:

- All console test and reporting commands are run from the SRM console.
- Certain environment variables are changed using the SRM SET command. For example:

  er*0_protocols
  ew*0_mode
  ew*0_protocols
  ocp_text
  pk*0_fast
  pk*0_host_id

To run the ECU, you must enter the ECU command. This command boots the ARC firmware and the ECU software.

ARC Menu Interface

- You must run the ECU whenever you add, remove, or move an EISA option in your system. The ECU is run from diskette. Two diskettes are shipped with your system.

5.18.2 Switching Between Interfaces

For a few procedures it is necessary to switch from one console interface to the other.

- The TEST command is run from the SRM interface.
- The ECU and the RCU are run from the ARC interface.

Switching from SRM to ARC

Two SRM console commands are used to switch to the ARC console:

- The ARC command loads the ARC firmware and switches to the ARC menu interface.
- The ECU command loads the ARC firmware and then boots the ECU diskette.

Switching from ARC to SRM

Switch from the ARC console to the SRM console as follows:

1. From the Boot menu, select the "Supplementary menu."
2. From the Supplementary menu, select "Set up the system."
3. From the Setup menu, select "Switch to OpenVMS console."
4. Select your operating system, then press Enter on Setup menu.
5. When the "Power-cycle the system to implement the change" message is displayed, press the Reset button. Once the console firmware is loaded and the system is initialized, the SRM console prompt >>> is displayed.

5.18.3 Verifying Configuration: SRM Console Commands for OpenVMS

Use the following SRM console commands to verify system configuration on OpenVMS systems:

- SHOW CONFIG—Displays the buses on the system and the devices found on those buses.
- SHOW DEVICE—Displays the devices and HS1CPs in the system.
- SHOW MEM—Displays main memory configuration.
- SET and SHOW—Sets and displays environment variable settings.

These commands are described in the following sections.

5.18.3.1 The SHOW CONFIG Command

The SHOW CONFIG command displays all devices found on the system bus, and EISA bus. You can use the information in the display to identify target devices for commands such as BOOT and TEST, as well as to verify that the system sees all the devices that are installed.

The configuration display includes the following:

- Firmware:

  The version numbers for the firmware code, PALcode, SROM chip, and CPU are displayed.

- Memory:

  The memory size and configuration for each bank of memory is displayed.

- **EISA Bus:**

  Slot numbers correspond to EISA card slots (1–8). For storage HS1CPs, the devices off the HS1CP also are displayed.

For more information on device names, refer to Figure 5–5.

Synopsis:

show config

Example:

```
>>> SHOW CONFIG
Firmware
SRM Console:    V1.1-1
ARC Console:    3.5-14
PALcode:        VMS PALcode X5.55, OSF PALcode X1.35-53
Serial Rom:      1.1

Processor
DECchip (tm) 21064-2

MEMORY
     128 Meg of System Memory
     Bank 0 = 64 Mbytes(16 MB Per Simm) Starting at 0x00000000
     Bank 1 = 64 Mbytes(16 MB Per Simm) Starting at 0x01000000
     Bank 2 = No Memory Detected
     Bank 3 = No Memory Detected

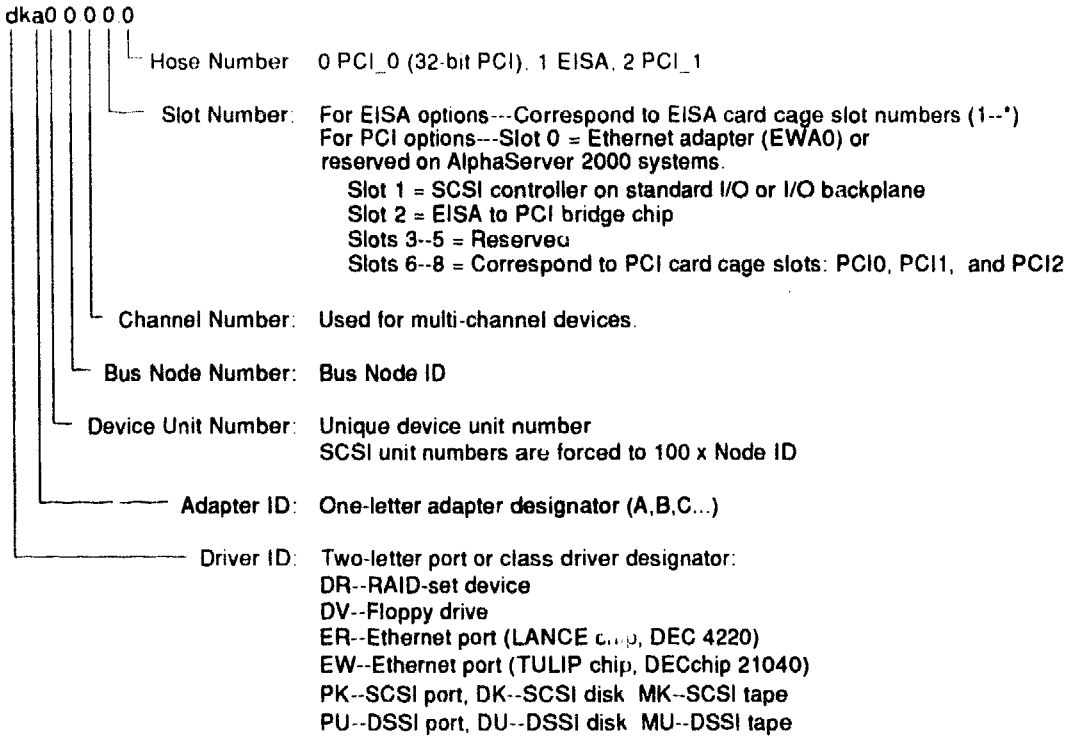
PCI Bus
     Bus 00  Slot 06:  NCR   810 SCSI  Controller
     Bus 00  Slot 07:  Intel  8275EB PCI to EISA Bridge
     Bus 00  Slot 11:  Digital DE435 Network Controller

EISA Bus
     Slot 2  DEC4220              era0.0.0.2.1       08-00-2B-BC-93-7A
>>>
```

5.18.3.2 The SHOW DEVICE Command

The SHOW DEVICE command displays the devices and controllers in the system. The device name convention is shown in Figure 5–5.

Figure 5–5 Device Name Convention



Synopsis:

show device [device_name]

Argument:

[device_name] The device name or device abbreviation. When abbreviations or wildcards are used, all devices that match the type are displayed.

Example:

```
>>> SHOW DEVICE
   ❶                    ❷                ❸      ❹
dka400.4.0.6.0          DKA400           RRD43    2893
dva0.0.0.0.1            DVA0
era0.0.0.2.1            ERA0             08-00-2B-BC-93-7A
pka0.7.0.6.0            PKA0                 SCSI Bus ID 7
>>>)
```

❶ Console device name

❷ Node name (alphanumeric, up to 6 characters)

❸ Device type

❹ Firmware version (if known)

5.18.3.3 The SHOW MEMORY Command

The SHOW MEMORY command displays information for each bank of memory in the system.

Synopsis:

```
show memory
```

Example:

```
>>> SHOW MEMORY
   128 Meg of System Memory
    Bank 0 = 64 Mbytes(16 MB Per Simm) Starting at 0x00000000
    Bank 1 = 64 Mbytes(16 MB Per Simm) Starting at 0x01000000
    Bank 2 = No Memory Detected
    Bank 3 = No Memory Detected
>>>
```

5.18.3.4 Setting and Showing Environment Variables

The environment variables described in Table 5–3 typically are set when you are configuring a system.

Synopsis:

set [-default] [-integer] -[string] envar value

Note

Whenever you use the SET command to reset an environment variable, you must initialize the system to put the new setting into effect. You initialize the system by entering the INIT command or pressing the Reset button.

show envar

Arguments:

envar	The name of the environment variable to be modified.
value	The value that is assigned to the environment variable. This can be an ASCII string.

Options:

-default	Restores variable to its default value.
-integer	Creates variable as an integer.
-string	Creates variable as a string (default).

Examples:

```
>>> SET BOOTDEF_DEV EZA0
>>> SHOW BOOTDEF_DRV
eza0
>>> SHOW AUTO_ACTION
boot
>>> SET BOOT_OSFLAGS 0,1
>>>
```

Table 5–3 Environment Variables Set During System Configuration

Variable	Attributes	Function
auto_action	NV,W	The action the console should take following an error halt or power failure. Defined values are: BOOT –Attempt bootstrap. HALT—Halt, enter console I/O mode. RESTART—Attempt restart. If restart fails, try boot. No other values are accepted. Other values result in an error message, and the variable remains unchanged.
bootdef_dev	NV	The device or device list from which booting is to be attempted, when no path is specified on the command line. Set at factory to disk with Factory Installed Software; otherwise NULL.
boot_file	NV,W	The default file name used for the primary bootstrap when no file name is specified by the BOOT command. The default value when the system is shipped is NULL.
boot_osflags	NV,W	Default additional parameters to be passed to system software during booting if none are specified by the BOOT command. On the OpenVMS Alpha operating system, these additional parameters are the root number and boot flags. The default value when the system is shipped is NULL. Common settings are a, autoboot; and Da, autoboot and create full dumps if the system crashes.
console	NV	Sets the device on which power-up output is displayed. GRAPHICS—Sets the power-up output to be displayed at a graphics terminal or device connected to the VGA module at the rear of the system. SERIAL—Sets the power-up output to be displayed on the device connected to the COM1 port at the rear of the system.

Key to variable attributes:

NV—Nonvolatile. The last value saved by system software or set by console commands is preserved across system initializations, cold bootstraps, and long power outages.
W—Warm nonvolatile. The last value set by system software is preserved across warm bootstraps and restarts.

(continued on next page)

Table 5–3 (Cont.) Environment Variables Set During System Configuration

Variable	Attributes	Function
ew*0_mode	NV	Sets the Ethernet controller to the default Ethernet device type. "aui"—Sets the default Ethernet device to AUI. "twisted"—Sets the default Ethernet device to 10BASE-T (twisted-pair). "auto"—Reads the device connected to the Ethernet port and sets the default to the appropriate Ethernet device type. (This option is not implemented at product introduction.)
er*0_protocols, ew*0_protocols	NV	Determines which network protocols are enabled for booting and other functions. "mop"—Sets the network protocol to mop: the setting typically used for systems using the OpenVMS operating system. "bootp"—Sets the network protocol to bootp. "bootp,mop"—When the settings are used in a list, the mop protocol is attempted first, followed by bootp.
os_type	NV	Sets the default operating system. "vms" or "osf"—Sets system to boot the SRM firmware. "nt"—Sets system to boot the ARC firmware.
pk*0_fast	NV	Enables Fast SCSI devices on a SCSI controller to perform in standard or fast mode. 0—Sets the default speed for devices on the controller to standard SCSI. If a controller has been set to standard SCSI mode, both standard and Fast SCSI devices perform in the standard mode. 1—Sets the default speed for devices on the controller to Fast SCSI mode. Devices on a controller that connect to both standard and Fast SCSI devices automatically perform at the appropriate rate for the device, either fast or standard mode.

Key to variable attributes:

NV—Nonvolatile. The last value saved by system software or set by console commands is preserved across system initializations, cold bootstraps, and long power outages.
W—Warm nonvolatile. The last value set by system software is preserved across warm bootstraps and restarts.

(continued on next page)

Table 5–3 (Cont.) Environment Variables Set During System Configuration

Variable	Attributes	Function
pk*0_host_id	NV	Sets the controller host bus node ID to a value between 0 and 7. 0 to 7—Assigns bus node ID for specified host adapter.

Key to variable attributes:

NV—Nonvolatile. The last value saved by system software or set by console commands is preserved across system initializations, cold bootstraps, and long power outages.
W—Warm nonvolatile. The last value set by system software is preserved across warm bootstraps and restarts.

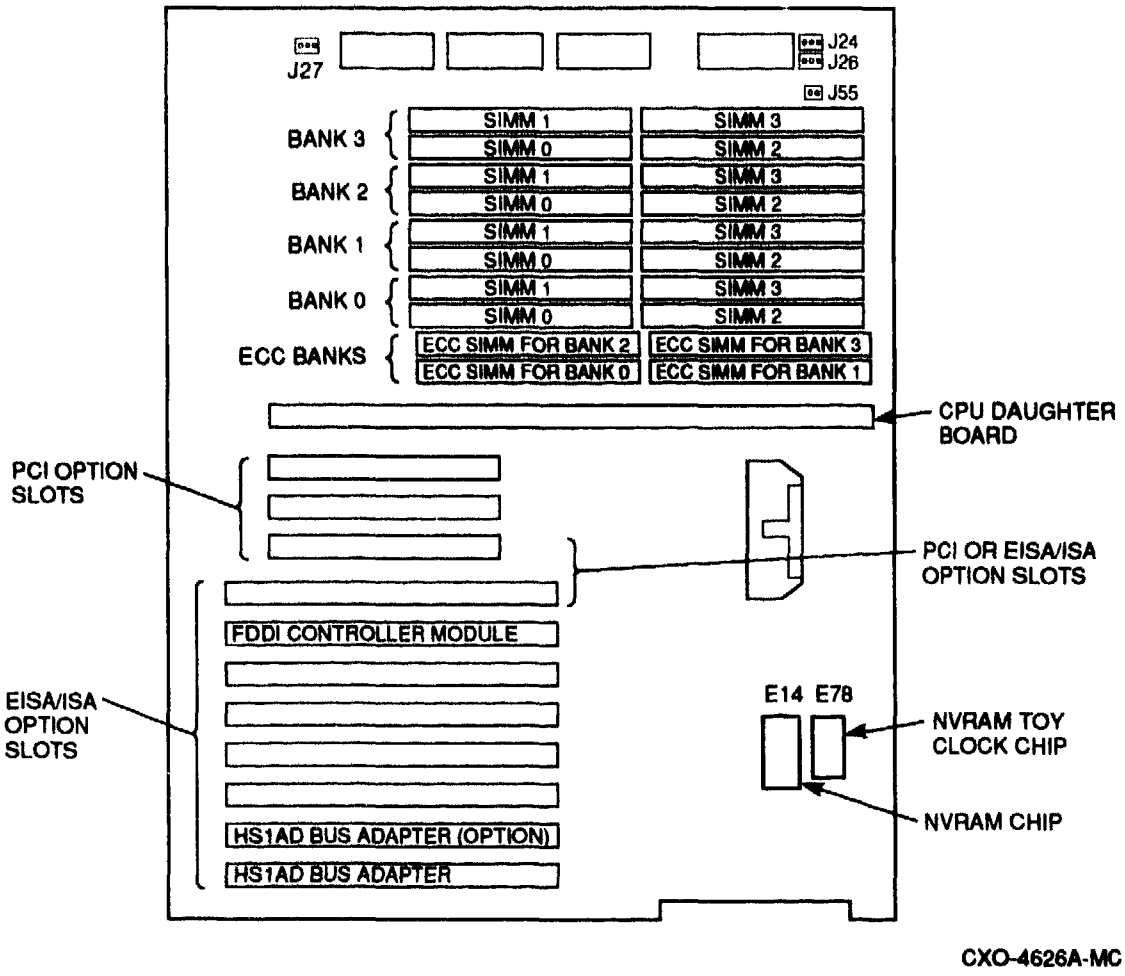
Note

Whenever you use the SET command to reset an environment variable, you must initialize the system to put the new setting into effect. Initialize the system by entering the INIT command or pressing the Reset button.

5.19 System Bus Options

The system bus interconnects the CPU and memory modules. Figure 5–6 shows the card slots and bus locations.

Figure 5–6 Card Slots and Bus Locations on the Motherboard



Note

If the top EISA connector is used (slot 8), the bottom PCI slot (slot 11) cannot be used. If the bottom PCI slot is used, the top EISA slot cannot be used.

5.19.1 CPU Daughterboard

FDDI Server systems use a CPU daughterboard. The daughterboard provides the following:

- The DECchip 21064 processor
- 2 MB of backup cache
- APECS chipset, which provides logic for external access to the cache for main memory control
- SROM code (SROM tests are controlled by jumper J6 on the CPU daughterboard)

5.19.2 Memory Modules

FDDI Server systems can support from 16 MB to 512 MB of memory. Memory options consist of five single in-line memory modules (SIMMs) and are available in the following variations:

- 64 MB (consisting of five 16 MB SIMMs)
- 128 MB (consisting of five 32 MB SIMMs)

The motherboard has 20 SIMM connectors. The SIMM connectors are grouped in four memory banks (0, 1, 2, and 3) and one bank for ECC (Error Correction Code) memory (Figure 5–6).

Memory Configuration Rules

Observe the following rules when configuring memory on FDDI Server systems:

- Bank 0 must contain a memory option (5 SIMMs–0, 1, 2, 3, and 1 ECC SIMM).
- A memory option consists of five SIMMs (0, 1, 2, 3 and 1 ECC SIMM for the bank).
- All SIMMs within a bank must be of the same capacity.
- OpenVMS requires 32 MB of memory as a minimum with 64 MB of memory recommended.

5.20 Motherboard

The motherboard provides a standard set of I/O functions:

- A Fast SCSI-2 controller chip that supports up to seven drives
- The firmware console subsystem on 1 MB of Flash ROM
- A floppy drive controller
- Two serial ports with full modem control and the parallel port
- The keyboard and mouse interface
- CIRRUS VGA controller
- The speaker interface
- Time-of-year (TOY) clock
- Connectors:
  1. EISA bus connectors (Slots 1-8)
  2. PCI bus connectors (Slots 11, 12, and 13)
     No PCI cards are used in the FDDI Server.
  3. Memory module connectors (20 SIMM connectors)
  4. CPU daughterboard connector

5.21 PCI Bus Options

PCI (Peripheral Component Interconnect) is an industry standard expansion I/O bus that is the preferred bus for high-performance I/O options. The server processor provides three slots for 32-bit PCI options.

PCI offers bus mastering and DMA capabilities. Up to three PCI cards can reside in the PCI portion of the motherboard. PCI slots can be filled in any order. Figure 5–6 shows the location of PCI slots on the motherboard. To access the slots, you must remove the rear top cover of the server processor enclosure. Be sure to replace the cover before attempting to start the server processor again.

CAUTION

Do not remove the server processor's front cover when it is running. If you remove the top cover without first properly turning off the system, the system shuts down, with potential loss of data.

For information about installing a specific option, refer to the documentation for that specific option.

WARNING

Before installing PCI bus options, turn off all power to the server processor.

5.22 EISA Bus Options

The EISA bus (Extended Industry Standard Architecture) is a 32-bit industry standard I/O bus. EISA is a superset of the well-established ISA bus. EISA was designed to accept newer 32-bit components while remaining compatible with older 8-bit and 16-bit cards.

EISA offers performance of up to 33 MB/sec for bus masters and DMA devices. Up to eight EISA modules can reside in the EISA bus portion of the card cage. All slots are bus master slots. EISA slots can be filled in any order.

WARNING

For protection against fir , only modules with current-limited outputs should be used.

5.23 EISA Configuration Utility

Whenever you add, remove, or move an EISA board in the system, you must run a utility called the EISA Configuration Utility (ECU). Each EISA board has a corresponding configuration (CFG) file, which describes the characteristics and the system resources required for that option. The ECU uses the CFG file to create a conflict-free configuration. The ECU is a menu-based utility that provides online help to guide you through the configuration process. The ECU is run from the ARC menu interface.

The ECU is supplied on the two System Configuration Diskettes shipped with the system. Make a backup copy of the system configuration diskette and keep the original in a safe place. Use the backup copy when you are configuring the system. The system configuration diskette must have the volume label "SYSTEMCFG."

Note

The CFG files supplied with the option you want to install may not work on this system if the particular option is not supported. Before you install an option, check to insure that the system supports the option.

5.23.1 Before You Run the ECU

Before running the ECU, do the following:

1. Install EISA options.

   For information about installing a specific option, refer to the documentation for that option.

2. Familiarize yourself with the utility.

   You can find more information about the ECU by reading the ECU online help. To read the online help, start the ECU (refer to Section 5.23.2). Online help for the ECU is located under Step 1, "Important EISA Configuration Information."

3. Familiarize yourself with the configuration procedure for the system (see Table 5–4).

4. Locate the ECU diskette for your operating system. The ECU diskette is shipped in the accessories box with the system. Do not use the original diskette. Use the backup copy for configuring options.

5.23.2 How to Start the ECU

Complete the following steps to run the ECU:

1. Invoke the console firmware.
   - Shut down the operating system and press the Halt button or power up with the Halt button set to the "in" position. The console prompt >>> is displayed.
   - Set the Halt button to the "out" position.
2. Start the ECU as follows:

   Note

   Make sure the ECU diskette is not write-protected.

   - Insert the ECU diskette for OpenVMS into the diskette drive.
   - Enter the ECU command.

     The console displays "Loading ARC firmware." Loading the ARC firmware takes approximately 2 minutes. When the firmware is loaded, the ECU program is booted.
3. Complete the ECU procedure according to the guidelines provided in the following sections.
4. After you have saved configuration information and exited from the ECU, remove the ECU diskette from the diskette drive.
5. Return to the SRM console firmware as follows:
   a. From the Boot menu, select "Supplementary menu."
   b. From the Supplementary menu, select "Set up the system." The Setup menu is displayed.
   c. From the Setup menu, select "Switch to OpenVMS console."
   d. Select your operating system console, then press ENTER on the "Setup menu."
   e. When the "Power-cycle the system to implement the change" message is displayed, press the Reset button. Once the console firmware is loaded and device drivers are initialized, you can boot the operating system.
6. Verify that the new options are configured correctly.

5.23.3 Configuring EISA Options

EISA boards are recognized and configured automatically. Review Section 5.23.1 then run the ECU as described in Section 5.23.2.

Note

It is not necessary to run Step 2 of the ECU, "Add or remove boards." (EISA boards are recognized and configured automatically.)

Table 5–4 Summary of Procedure for Configuring an EISA Bus Option

Step	Explanation
Install EISA option.	Use the instructions provided with the EISA option.
Power up and run ECU.	If the ECU locates the required CFG configuration files, it displays the main menu. The CFG file for the option resides either on a configuration diskette that is packaged with the option or on the system configuration diskette. Note: It is not necessary to run Step 2 of the ECU, "Add or remove boards." EISA boards are recognized and configured automatically.
View or Edit Details (optional).	The "View or Edit Details" ECU option is used to change user-selectable settings or to change the resources allocated for these functions (IRQs, DMA channels, I/O ports, and so on). This step is not required when using the board's default settings.
Save your configuration and restart the system.	The "Save and Exit" ECU option saves your configuration information to the system's nonvolatile memory.
Return to the SRM console and restart the system.	Refer to step 4 of Section 5.23.2 for operating system-specific instructions.

5.23.4 SCSI Buses

A fast SCSI-2 adapter on the system motherboard provides a single-ended SCSI bus for the FDDI Server.

The fast SCSI-2 adaptor on the motherboard supports one 3.5-inch and one 5.25-inch, internal half-height removable-media devices.

This bus also supports the internal 3.5-inch hard disk drive.

5.24 Console Port Configurations

Power-up information typically is displayed on the system's console terminal. The console terminal can be either a graphics terminal or a serial terminal (connected through the COM1 serial port). The setting of the console environment variable determines where the system displays power-up output. Set this environment variable according to the console terminal that you are using.

Synopsis:

set console output_device

Arguments:

graphics	Displays the power-up output to a graphics terminal or device connected to the VGA module at the rear of the system.
serial	Displays the power-up output to a device connected to the COM1 port at the rear of the system.

Note

Whenever you use the SET command to reset an environment variable, you must initialize the system to put the new setting into effect. Initialize the system by entering the INIT command or pressing the Reset button.

Example:

```
>>> SET CONSOLE SERIAL
>>>
```

VT*xxx* Console Terminal Setting for Running ECU

To run the EISA configuration utility (ECU) from the serial console port, the terminal must be set for 8-bit controls, the keyboard must be set so that the tilde (~) key sends the escape (ESC) signal, and the console environment variable must be set to serial.

Graphics Terminal Needed for Running RCU

A graphics terminal is needed to run the RCU. To enable the on-board VGA logic, the VGA enable jumper (J27) on the motherboard must be enabled. The console environment variable must be set to graphics.

Using a VGA Controller Other than the Standard On-Board VGA

When the system is configured to use an EISA-based VGA controller instead of the standard on-board VGA (CIRRUS), consider the following:

- The on-board CIRRUS VGA options must be set to disabled through the ECU.
- The VGA jumper (J27) on the upper-left corner of the motherboard must then be set to disable (off).
- The console environment variable must be set to graphics.
- If there are multiple VGA controllers, the system directs console output to the first controller it finds.

6

Error Analysis and Fault Isolation

This chapter describes the errors, faults, and significant events that can occur during initialization and normal operation of the HS1CP and the server processor. A translation of the events, and in most cases how to respond to a specific event, also is given.

The error and event descriptions isolate failures to the field replaceable unit (FRU). However, in most instances, additional information for diagnosis beyond the FRU is given. This information helps increase your knowledge of server processor and HS1CP functions and assists with your report to repair personnel.

CAUTION

If you attempt to replace or repair components within FRUs, equipment damage may result. Use the controller fault indications and error logs to isolate FRU-level failures.

This chapter discusses the following:

- Special considerations for the HS1CP
- HS1CP troubleshooting basics
- Types of HS1CP error reporting
- Host error logs

6.1 Special Considerations for the HS1CP

Some or all of the following situations could apply when your HS1CP detects a fault:

- **Nonredundant Configurations**

  When an HS1CP (or its cache module, or both) fails in a nonredundant configuration, a short period of system down time is needed to remove the faulty unit and to install a replacement. The devices attached to that HS1CP are offline for the duration of the remove and replace cycle.

- **Dual-Redundant Configurations**

  When an HS1CP fails in a dual-redundant configuration, fault isolation and corrective actions are similar to a nonredundant configuration. However, failover occurs, so the surviving HS1CP takes over service to the failed HS1CP's ports and devices.

- **Cache Module Failures**

  If a cache module fails, its HS1CP still functions using on-board cache; however, Digital recommends that you replace the cache module as soon as possible.

- **Write-Back Cache Battery Failures**

  When a write-back cache module's batteries fail or are low, the HS1CP alters its operations until the batteries are fully charged:

  - Stripesets and disk-based units with write-back caching enabled are accessed in write-through (read cache) mode.
  - If the CACHE_POLICY on the HS1CP is set to A and the batteries are low when the HS1CP initializes, any RAIDset or mirrorset that does not have access to good batteries is made inoperative.
  - If the CACHE_POLICY on the HS1CP is set to B and the batteries are low when the HS1CP initializes, any RAIDset or mirrorset that does not have access to good batteries is accessed in write-through (read cache) mode.
  - If the batteries go low after the HS1CP initialized, any RAIDset or mirrorset that does not have access to good batteries is made inoperative, regardless of the cache policy.
  - If the batteries are bad or missing, RAIDsets and mirrorsets are made inoperative.
  - Non-RAIDset, disk-based units with write-back caching enabled are accessed in write-through (read cache) mode until the cache batteries are replaced or fully recharged. Once good batteries are in place, write-back caching automatically resumes.
  - RAIDsets fail over to the companion cache in a dual-redundant configuration, provided the other cache's batteries are fully charged. Any RAIDset that does not have access to good batteries is not served by the HS1CP.

6.2 HS1CP Troubleshooting Basics

When an error occurs, use the following steps as top-level guidelines for fault isolation:

1. Make a note of all visual indicators (front panel indicators or error messages) available to you.
2. Extract and read host error logs (see Section 6.4).
3. For surviving HS1CPs in dual-redundant pairs, try entering the RESTART OTHER_CONTROLLER command. The surviving HS1CP could be keeping its companion from operating.

4. Errors can be intermittent; reseat the HS1CP to see if the error clears.[1]
5. See if the error indication changes after reseating the HS1CP. If the error remains the same, look up information for that error. If the error indication changes, look up information for the newer error.
6. Always consider reseating the HS1CP and/or cache module when troubleshooting. Poor connections between module and backplane can cause a variety of errors.

   See Sections 6.3.1 through 6.3.3 for detailed information about errors and repair actions.

CAUTION

Do not use an HS1CP failure situation as an opportunity to move devices or otherwise reconfigure your subsystem. Doing so prevents the HS1CP from communicating with its units once the fault is corrected.

6.3 Types of HS1CP Error Reporting

The HS1CP can notify you of an error through one or more of the following means:

- The front panel indicators
- Error messages at a host virtual terminal, or

  Error messages at a maintenance terminal (if attached)
- Host error logs

6.3.1 HS1CP Front Panel Indicators

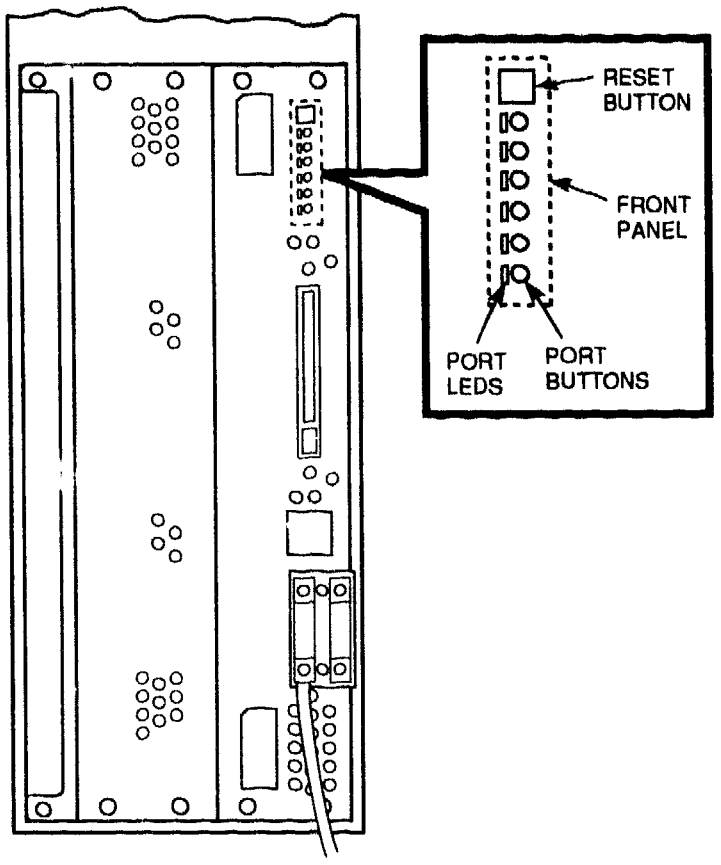
The front panel includes the following indicators:

- One reset button with an embedded green LED
- One button per SCSI port
- Six amber LEDs

Figure 6–1 shows an example of an HS1CP and the front panel indicators. The buttons and LEDs serve different functions with respect to controlling the SCSI ports and/or reporting fault and normal conditions. The button and LED functions are discussed in the following sections.

[1] Record which devices have lit/flashing fault LEDs before reseating, as reseating the HS1CP can clear the LED temporarily, even though the fault remains.

Figure 6–1 HS1CP Layout



CXO-4461B-MC

6.3.1.1 HS1CP Normal Operation

The green LED button reflects the state of the HS1CP and the host interface. Once HS1CP initialization completes and its firmware is functioning, the green button flashes continuously at 1 Hz (once per second). Pressing the green button during normal operation resets the HS1CP.

Under normal operation, the amber LEDs indicate the state of the respective SCSI–2 device buses attached to the HS1CP. When the devices on the buses are functioning correctly, the amber LEDs do not light or flash.

Pressing one of the port buttons during normal operation lights its corresponding amber LED and momentarily quizzes its SCSI–2 port. You must quiesce a port to remove or warm swap a device on the SCSI–2 bus for that port. (Use a small pointed object to press the port button.) See Chapter 8 for a detailed description of removing and replacing devices.

6.3.1.2 Fault Notification

The front panel LEDs display information when the HS1CP encounters a problem with a device configuration, a device, or the HS1CP itself.

If a configuration mismatch or a device fault occurs, the amber LED for the affected device's bus lights continuously.

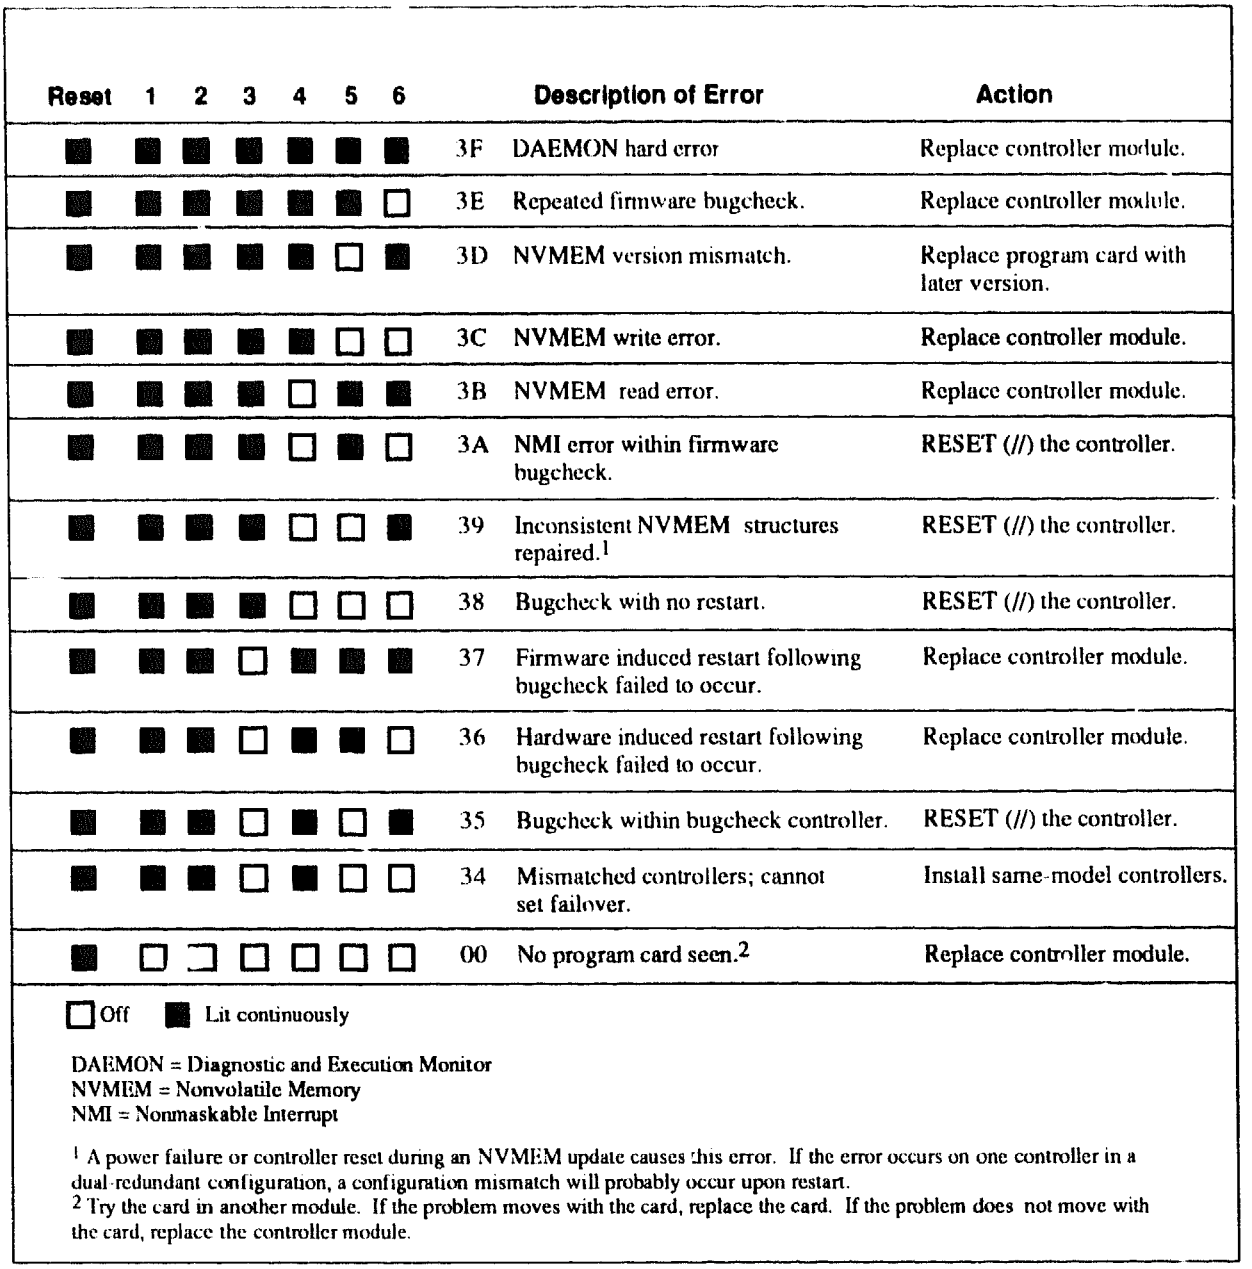
For HS1CP problems, LED codes determined by internal diagnostics and operating firmware indicate either HS1CP faults or HS1CP operating firmware program card faults. In either case, the single green reset LED lights continuously when an error is detected. The remaining amber LEDs display the error codes in two different ways:

- The error code lights continuously for faults detected by internal diagnostic and initialization routines. See Figure 6–2 to determine what these codes mean.
- The error code flashes at 3 Hz representing faults that occur during normal controller operation. See Figure 6–3 to determine what these codes mean.

Note

If the controller finds errors during initialization or self-test diagnostics, any flashing codes momentarily stop flashing while the error messages are posted to the console. Normal flashing resumes once the display completes. Do not attempt to look up the codes while this information is printing at the terminal.

Figure 6–2 Solid Front Panel Codes



Reset	1	2	3	4	5	6		Description of Error	Action
■	■	■	■	■	■	■	3F	DAEMON hard error	Replace controller module.
■	■	■	■	■	■	□	3E	Repeated firmware bugcheck.	Replace controller module.
■	■	■	■	■	□	■	3D	NVMEM version mismatch.	Replace program card with later version.
■	■	■	■	■	□	□	3C	NVMEM write error.	Replace controller module.
■	■	■	■	□	■	■	3B	NVMEM read error.	Replace controller module.
■	■	■	■	□	■	□	3A	NMI error within firmware bugcheck.	RESET (//) the controller.
■	■	■	■	□	□	■	39	Inconsistent NVMEM structures repaired.[1]	RESET (//) the controller.
■	■	■	■	□	□	□	38	Bugcheck with no restart.	RESET (//) the controller.
■	■	■	□	■	■	■	37	Firmware induced restart following bugcheck failed to occur.	Replace controller module.
■	■	■	□	■	■	□	36	Hardware induced restart following bugcheck failed to occur.	Replace controller module.
■	■	■	□	■	□	■	35	Bugcheck within bugcheck controller.	RESET (//) the controller.
■	■	■	□	■	□	□	34	Mismatched controllers; cannot set failover.	Install same-model controllers.
■	□	□	□	□	□	□	00	No program card seen.[2]	Replace controller module.

□ Off ■ Lit continuously

DAEMON = Diagnostic and Execution Monitor
NVMEM = Nonvolatile Memory
NMI = Nonmaskable Interrupt

[1] A power failure or controller reset during an NVMEM update causes this error. If the error occurs on one controller in a dual-redundant configuration, a configuration mismatch will probably occur upon restart.
[2] Try the card in another module. If the problem moves with the card, replace the card. If the problem does not move with the card, replace the controller module.

Figure 6–3 Flashing Front Panel Codes

Reset	1	2	3	4	5	6		Description of Error	Action
■	□	□	□	□	□	◪	01	Program card EDC error.	Replace program card.
■	□	□	□	◪	□	□	04	Timer zero in the timer chip will run when disabled.	Replace controller module.
■	□	□	□	◪	□	◪	05	Timer zero in the timer chip decrements incorrectly.	Replace controller module.
■	□	□	□	◪	◪	□	06	Timer zero in the timer chip did not interrupt the processor when requested.	Replace controller module.
■	□	□	□	◪	◪	◪	07	Timer one in the timer chip decrements incorrectly.	Replace controller module.
■	□	□	◪	□	□	□	08	Timer one in the timer chip did not interrupt the processor when requested.	Replace controller module.
■	□	□	◪	□	□	◪	09	Timer two in the timer chip decrements incorrectly.	Replace controller module.
■	□	□	◪	□	◪	□	0A	Timer two in the timer chip did not interrupt the processor when requested.	Replace controller module.
■	□	□	◪	□	◪	◪	0B	Memory failure in the I/D cache.	Replace controller module.
■	□	□	◪	◪	□	□	0C	No hit or miss to the I/D cache when expected.	Replace controller module.
■	□	□	◪	◪	□	◪	0D	One or more bits in the diagnostic registers did not match the expected reset value.	Replace controller module.
■	□	□	◪	◪	◪	□	0E	Memory error in the nonvolatile journal SRAM.	Replace controller module.
■	□	□	◪	◪	◪	◪	0F	Wrong image seen on program card.	Replace program card.
■	□	◪	□	□	□	□	10	At least one register in the controller DRAB does not read as written.	Replace controller module.
■	□	◪	□	□	□	◪	11	Main memory is fragmented into too many sections for the number of entries in the good memory list.	Replace controller module.
■	□	◪	□	□	◪	□	12	The controller DRAB chip does not arbitrate correctly.	Replace controller module.

□ Off ■ Lit continuously ◪ Flashing

I/D = Instruction/Data (cache on the controller module)
DRAB = Dynamic RAM Controller and Arbitration Engine (operates controller shared memory)
ECC = Error Correction Code
EDC = Error Detection Code
SRAM = Static RAM
NXM = Nonexistent Memory

(continued on next page)

Figure 6–3 (Cont.) Flashing Front Panel Codes

Reset	1	2	3	4	5	6		Description of Error	Action
■	□	◪	□	□	◪	◪	13	The controller DRAB chip failed to detect forced parity, or detected parity when not forced.	Replace controller module.
■	□	◪	□	◪	□	□	14	The controller DRAB chip failed to verify the EDC correctly.	Replace controller module.
■	□	◪	□	◪	□	◪	15	The controller DRAB chip failed to report forced ECC.	Replace controller module.
■	□	◪	□	◪	◪	□	16	The controller DRAB chip failed some operation in the reporting, validating, and testing of the multibit ECC memory error.	Replace controller module.
■	□	◪	□	◪	◪	◪	17	The controller DRAB chip failed some operation in the reporting, validating, and testing of the multiple single-bit ECC memory error.	Replace controller module.
■	□	◪	◪	□	□	□	18	The controller main memory did not write correctly in one or more sized memory transfers.	Replace controller module.
■	□	◪	◪	□	□	◪	19	The controller did not cause an I-to-N bus timeout when accessing a "reset" host port chip.	Replace controller module.
■	□	◪	◪	□	◪	□	1A	The controller DRAB chip did not report an I-to-N bus timeout when accessing a "reset" host port chip.	Replace controller module.
■	□	◪	◪	□	◪	◪	1B	The controller DRAB did not interrupt the controller processor when expected.	Replace controller module.
■	□	◪	◪	◪	□	□	1C	The controller DRAB did not report an NXM error when nonexistent memory was accessed.	Replace controller module.
■	□	◪	◪	◪	□	◪	1D	The controller DRAB did not report an address parity error when one was forced.	Replace controller module.
■	□	◪	◪	◪	◪	□	1E	There was an unexpected nonmaskable interrupt from the controller DRAB during the DRAB memory test.	Replace controller module.
■	□	◪	◪	◪	◪	◪	1F	Diagnostic register indicates there is no cache module, but an interrupt exists from the nonexistent cache module.	Replace controller shelf backplane.
■	◪	□	□	□	□	□	20	The required amount of memory available for the code image to be loaded from the program card is insufficient.	Replace controller module.

□ Off ■ Lit continuously ◪ Flashing

I/D = Instruction/Data (cache on the controller module)
DRAB = Dynamic RAM Controller and Arbitration Engine (operates controller shared memory)
ECC = Error Correction Code
EDC = Error Detection Code
SRAM = Static RAM
NXM = Nonexistent Memory

(continued on next page)

Figure 6–3 (Cont.) Flashing Front Panel Codes

Reset	1	2	3	4	5	6		Description of Error	Action
■	◪	□	□	□	□	◪	21	The required amount of memory available in the pool area is insufficient for the controller to run.	Replace controller module.
■	◪	□	□	□	◪	◪	23	The required amount of memory available in the buffer area is insufficient for the controller to run.	Replace controller module.
■	◪	□	□	◪	□	□	24	The code image was not the same as the image on the card after the contents were copied to memory.	Replace ontroller module.
■	◪	□	□	◪	□	◪	25	Diagnostic register indicates that the cache module exists, but access to that cache module caused an error.	Replace controller shelf backplane.
■	◪	□	□	◪	◪	□	26	Diagnostic register indicates that the cache module does not exist, but access to that cache module did not cause an error	Replace controller shelf backplane.
■	◪	◪	□	□	□	□	30	The journal SRAM battery is bad.	Replace controller module.
■	◪	◪	◪	□	◪	□	3A	There was an unexpected interrrupt from a read cache or the present and lock bits are not working correctly.	Replace controller module.
■	◪	◪	◪	□	◪	◪	3B	There is an interrupt pending to the co. oller's policy processor when there should be n e.	Replace controller module.
■	◪	◪	◪	◪	□	□	3C	There was an unexpected fault during initialization.	Replace controller module.
■	◪	◪	◪	◪	□	◪	3D	There was an unexpected maskable interrupt received during initialization.	Replace controller module.
■	◪	◪	◪	◪	◪	□	3E	There was an unexpected nonmaskable interrupt received during initialization.	Replace controller module.
■	◪	◪	◪	◪	◪	◪	3F	An illegal process was activated during initialization.	Replace controller module.

□ Off ■ Lit continuously ◪ Flashing

I/D = Instruction/Data (cache on the controller module)
DRAB = Dynamic RAM Controller and Arbitration Engine (operates controller shared memory)
ECC = Error Correction Code
EDC = Error Detection Code
SRAM = Static RAM
NXM = Nonexistent Memory

6.3.2 Device LEDs

The storage devices (SBBs) and their power supplies have LEDs to indicate power and status. You can use these LEDs in conjunction with the front panel indicators to isolate certain faults, as discussed in the following sections.

6.3.2.1 Storage SBB Status

Device shelves monitor the status of the storage SBBs. When a fault occurs, the fault and the SBB device address (SCSI target ID) are reported to the HS1CP for processing. The SBB internal fault/identity bus controls the fault (lower) LED.

As shown in Figures 6–4 and 6–5, each storage SBB has two LED indicators that display the SBB's status. These LEDs have three states: on, off, and flashing.

- The upper LED (green) is the device activity LED and is on or flashing when the SBB is active.

  CAUTION

  Do not remove a storage SBB when the upper LED is on or flashing. This can cause the loss or corruption of data.

- The lower LED (amber) is the storage SBB fault LED and indicates an error condition when it is either on or flashing. When this LED indicates a fault, the HS1CP's amber LED for the device's port is lit continuously as well.

  CAUTION

  You should record which devices have lit/flashing fault LEDs before resetting the HS1CP; resetting it can clear this LED temporarily, even though the fault remains.

Figure 6–4 3.5-Inch Storage SBB LEDs

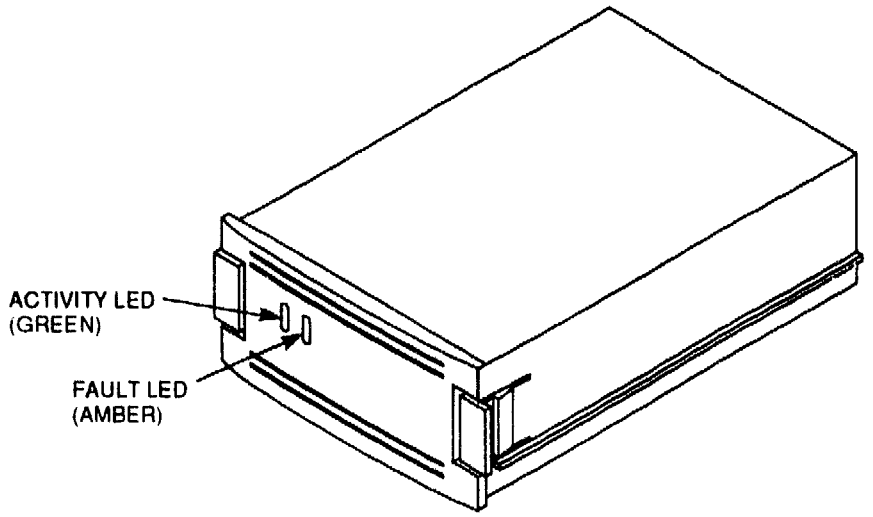


Figure 6–5 5.25-Inch Storage SBB LEDs

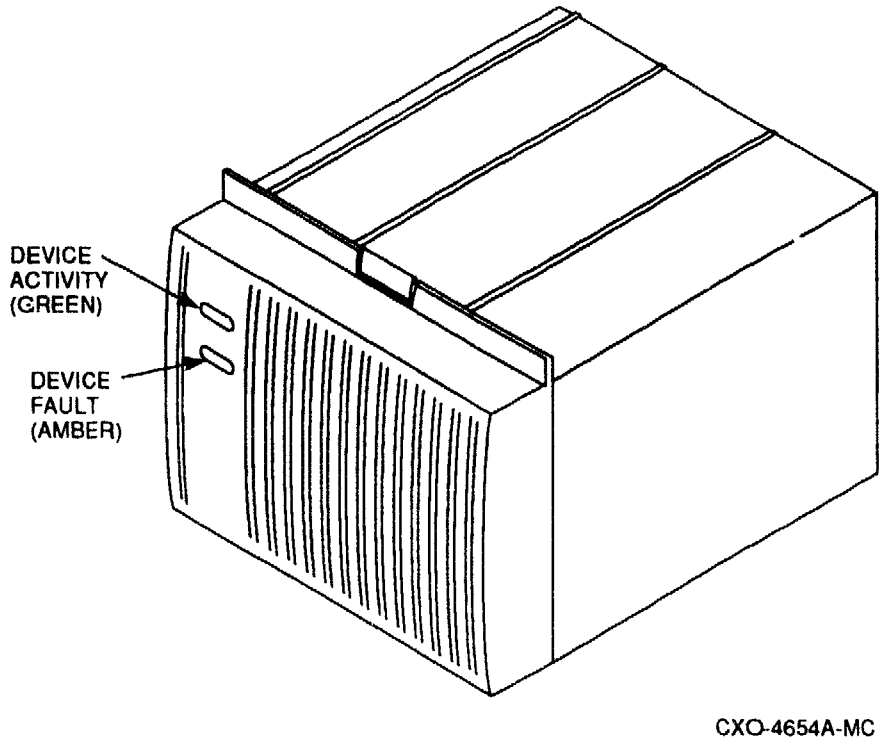


Table 6–1 explains the SBB LED codes.

Table 6–1 Storage SBB Status LEDs

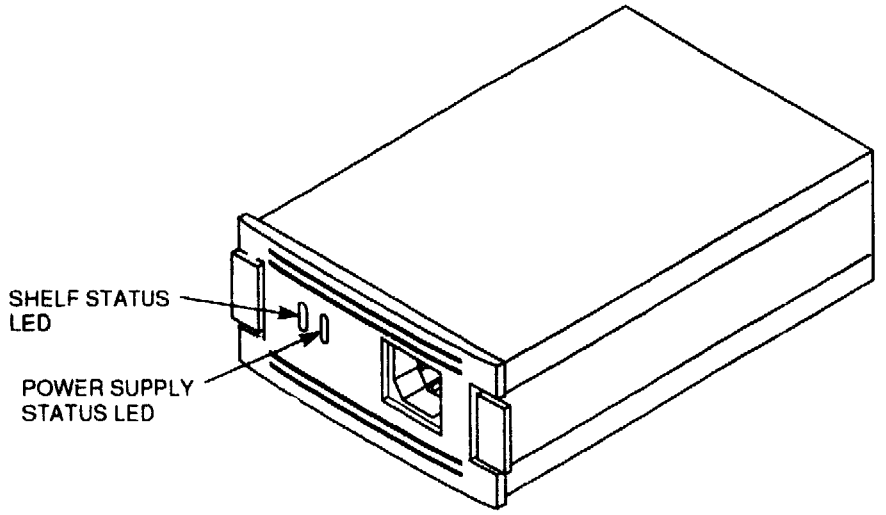
LED	Status	Indication
Device activity Device fault	On Off	SBB is operating normally.
Device activity Device fault	Flashing Off	SBB is operating normally.
Device activity Device fault	Off Off	SBB is operating normally. The SBB is inactive, and there is no fault.
Device activity Device fault	On On	Fault status. SBB is probably not responding to control signals. Digital recommends that you replace the SBB.
Device activity Device fault	Off On	Fault status. SBB is inactive and spun down. Digital recommends that you replace the SBB.
Device activity Device fault	On Flashing	Fault status. SBB is active and is spinning down because of the fault.

6.3.2.2 Device Shelf Status and Power Supply Status

The status of both the device shelf blowers and power supplies is displayed on the power supply LEDs, as shown in Figure 6–6. The upper LED displays the shelf status and the lower LED displays the power supply status.

- When the upper LED is on, both the shelf blowers and the power supplies are functioning properly.
- When the upper LED is off, either a shelf blower or a power supply is not functioning properly.
- When the lower LED is off, either there is an input power problem or the power supply is not functioning.

Figure 6–6 Power Supply LEDs



For a detailed explanation of the power supply LED codes, see 6–2 and 6–3.

Table 6–2 Shelf and Single Power Supply Status LEDs

Status LED	State	Indication
Shelf (upper) PS (lower)	On On	System is operating normally.
Shelf (upper) PS (lower)	Off On	Fault status. There is a shelf fault; there is no power supply fault. Replace blower as described in Chapter 8
Shelf (upper) PS (lower)	Off Off	Fault status. Shelf and power supply fault. Replace power supply as described in Chapter 8.

Note

When a shelf has two power supplies, you must observe the LEDs on *both* power supplies to determine the status (see Table 6–3).

Table 6–3 Shelf and Dual Power Supply Status LEDs

Status LED	PS1†	PS2‡	Indication
Shelf (upper) PS (lower)	On On	On On	Normal status. System is operating normally.
Shelf (upper) PS (lower)	Off On	Off On	Fault status. There is a shelf fault; there is no power supply fault. Replace blower as described in Chapter 8.
Shelf (upper) PS (lower)	Off On	Off Off	Fault status. PS1 is operational. Replace PS2 as described in Chapter 8.
Shelf (upper) PS (lower)	Off Off	Off On	Fault status. PS2 is operational. Replace PS1 as described in Chapter 8.
Shelf (upper) PS (lower)	Off Off	Off Off	Fault status. Possible PS1 and PS2 fault or input power problem.

† Shelf power supply installed in slot 7.

‡ Redundant power supply installed in slot 6.

6.3.3 Error Messages

The HS1CP operating firmware is designed to send both spontaneous and stored messages to your terminal session under certain fault conditions.

Spontaneous messages are part of a command line interpreter (CLI) error report (CER) system, which causes the message text to appear on a maintenance terminal along with the CLI prompt, as shown in the following example. (Spontaneous messages do not appear on virtual terminals.)

```
%CER -- 13-JUL-1994 13:28:45 -- SWAP signal cleared - all SWAP interrupts rambled
CLI>
%CER -- 13-JUL-1994 13:29:11 -- Other controller restarted
CLI>
```

6.3.3.1 Conditions under Which CER Messages Appear

The spontaneous CER messages appear only under the following conditions:

- You have a maintenance terminal connected for receiving messages.
- The subsystem has finished initializing.
- You currently are not running a utility program on the maintenance terminal.
- A maintenance terminal is not actively displaying input from another source, such as event logging or last failure logging.
- No CLI commands are in progress on a maintenance terminal.

If any one of the previous conditions is not met, the spontaneous CER message display does not occur. In this case, the HS1CP operating firmware stores messages for you. You need only connect a virtual or maintenance terminal (if one is not already connected) and press the Return key from the CLI prompt to review the 15 most recently received error messages:

```
SWAP signal cleared - all SWAP interrupts rambled
Other controller restarted
CLI>
```

6.3.3.2 Clearing the Terminal of Error Messages

Often, message review continues to occur each time you press the Return key. To clear the terminal of the errors, enter the CLEAR_ERRORS CLI command. (You may want to make a note of the errors before clearing them, because they cannot be recalled afterwards.)

Note

Because the severity of errors varies, the HS1CP may not initialize or operate, or both, even though an error message appears.

For example, if all of the SCSI ports, or the host port and local terminal port fail diagnostics, the HS1CP will not operate. However, if the cache module fails during normal operation, the HS1CP continues to operate.

6.3.3.3 Types of Error Messages

The following sections list CER messages you may encounter. The HS1CP sends these messages when the specific fault is detected, regardless of whether you are interactively viewing or using the virtual or maintenance terminal. These messages differ in this respect from the ones listed in Appendix E, which appear based on your interactive use of the CLI. The CER messages are presented in the following error categories:

- Configuration and CLI
- Diagnostic and initialization
- Failover
- NVPM
- Shelf
- Write-back cache
- Failure indications
- System information structure and system information page

Consult your firmware release notes for updates to the list of error messages.

Configuration and CLI Messages

This section describes error messages that can display if a fault occurs related to configuration or CLI. See Appendix D for information on configuration rules. See Appendix E for information on the command line interpreter.

Configuration information deleted due to internal inconsistencies

Explanation: This message is displayed if a test of nonvolatile memory shows corruption. The configuration information for the HS1CP is deleted when this message is displayed.

Controllers misconfigured. Type SHOW THIS_CONTROLLER

Explanation: If this message appears, examine the SHOW THIS_CONTROLLER display to determine the source of the misconfiguration.

Device and/or storageset names changed to avoid conflicts

Explanation: Digital adds new CLI keywords at each new HS1CP operating firmware release that can conflict with existing device and/or storageset names. When this occurs, HS1CP operating firmware changes your device and/or storageset names and sends this message. The functional operation of your configuration is not changed when this message appears.

Disk mirroring option has been turned on without a valid license key

Explanation: Disk mirroring is a licensed feature. This message indicates that it has been turned on without entering a valid license key in the Firmware Licensing System (FLS).

Licensing different between the two controllers

Explanation: The licensing features are set differently on two controllers of a dual-redundant pair.

RAID5 option has been turned on without a valid license key

Explanation: RAID5 is a licensed feature. This message indicates that it has been turned on without entering a valid license key in the Firmware Licensing System (FLS).

Restart of the controller required to apply new patch

Explanation: You used the Code Patch utility to enter a firmware patch, but the patch is not applied until you restart the controller.

Restart of the other controller required

Explanation: When changing some parameters, you must reinitialize the companion controller in a dual-redundant pair to have the parameter take effect.

Restart of this controller required

Explanation: A changed parameter requires reinitialization of this controller to take effect.

Serial number initialized due to format error

Explanation: An invalid serial number was entered for the second controller of a dual-redundant pair.

Taken out of failover due to serial number format error

Explanation: An invalid serial number format was entered for the second controller of a dual-redundant pair.

When you have a configuration mismatch (both controllers not in dual-redundant mode, failover not correctly set, controller serial numbers or logical configurations do not match, and so on), the following message appears:

```
Two controllers present but not in dual-redundant configuration.
```

To correct the problem, do the following:

1. Enter the CLI SHOW THIS_CONTROLLER command at the CLI> prompt, and review this controller's configuration.
2. Enter the CLI SHOW OTHER_CONTROLLER command at the CLI> prompt, and review the other controller's configuration.
3. Decide which controller has the most current information.
4. Enter the CLI SET FAILOVER COPY= command based on where the most current information is.

   For example, if the most current information resides on this controller, enter SET FAILOVER COPY=THIS_CONTROLLER.

If the unit numbers are displayed in an obsolete format, the following message is displayed:

```
Units renamed to new Dn and Tn formats.
```

In the previous message, the variable *n* is the unit number from the old format.

Write-back cache option has been turned on without a valid license key

Explanation: Write-back cache is a licensed feature. This message indicates that it has been turned on without entering a valid license key in the Firmware Licensing System (FLS).

Diagnostic and Initialization Messages

This section describes error messages that can display if a fault occurs during initialization or self-test diagnostics. See Chapter 3 for more information on diagnostics.

HOST port FAILED Diagnostics

Explanation: The host port of the controller has failed diagnostics.

Local Terminal Port FAILED Diagnostics

Explanation: The maintenance (EIA–423) terminal port has failed diagnostics.

SCSI port *n* FAILED Diagnostics

Explanation: A SCSI–2 port has failed diagnostics. This message can appear even if you do not have a host connection. The variable *n* indicates which port failed.

The following Firmware Licensing Service component elements were initialized to default settings: [*n* ...

Explanation: Licensing information for licensed features has been lost. To correct this problem, reenter your license keys through FLS.

Failover Messages

This section contains error messages that can display if a failover between dual-redundant controllers occurs. See Section 5.14 for more information on failover.

Both HS1CP controllers are using SCSI address 6

Explanation: There is a hardware problem with the BA350–M shelf. This problem probably involves the shelf backplane.

Both HS1CP controllers are using SCSI address 7

Explanation: There is a hardware problem with the BA350–M shelf. This problem probably involves the shelf backplane.

Other controller not responding - RESET signal asserted

Explanation: One controller in a dual-redundant configuration is locked up, not responding, or the kill line to it is asserted.

Other controller restarted

Explanation: The other controller in a dual-redundant pair has restarted successfully after failing or undergoing a bugcheck. See Section 6.4 for more information on this message.

Received LAST GASP message from other controller

Explanation: One controller in a dual-redundant configuration is attempting an automatic restart after failing or undergoing a bugcheck. See Section 6.4 for more information on this message.

SCSI Device and HS1CP controller both configured at SCSI address 6

Explanation: This message appears when a device is accidentally configured as SCSI ID 6, and two controllers (SCSI IDs 6 and 7) are in a dual-redundant configuration.

NVPM Messages

This section contains error messages that can display because of a problem or fault associated with the nonvolatile parameter memory (**NVPM**).

Note

Some NVPM messages read "NVPM *component-name* component initialized to default settings." For some of these initialization cases, corrective action may clear the error message only until the next time the controller is reset, because the error could be caused by a fault in NVPM itself. If the error persists, replace the controller module.

NVPM *component-name* component initialized to default settings

Explanation: The NVPM component named by *component-name* has been reinitialized. Performing a controller reset may clear this error message only until the next time the controller is reset because the error could be caused by a fault in NVPM itself. If this error persists, replace the controller module.

NVPM Failover Information component initialized to default settings

Explanation: The identity of the other controller in a dual-redundant pair has been lost. Enter the SET FAILOVER COPY=OTHER_CONTROLLER command to correct this problem. If the error persists, replace the controller module.

NVPM FMU Parameters component initialized to default settings

Explanation: FMU settings have been lost. Default FMU options remain in place until you run FMU to change them. If the error persists, replace the controller module.

NVPM Host Interconnect Parameters component initialized to default settings

Explanation: The SCS node name, CI node number, or Path A, or Path B enable settings for this controller have been lost.

To correct this problem, enter the SHOW THIS_CONTROLLER and SHOW OTHER_CONTROLLER commands to determine the current controller settings. Use the SET THIS_CONTROLLER and SET OTHER_CONTROLLER commands to restore settings. If the error persists, replace the controller module.

NVPM Host Protocol Parameters component initialized to default settings

Explanation: The tape and disk MSCP allocation class settings for this controller have been lost.

To correct this problem, enter the SHOW THIS_CONTROLLER and SHOW OTHER_CONTROLLER commands to determine the current controller settings. Use the SET THIS_CONTROLLER and SET OTHER_CONTROLLER commands to restore settings. If the error persists, replace the controller module.

NVPM Product Information component initialized to default settings

Explanation: The controller product identification setting has been reset. If the controller "name" is not what it used to be, the controller module must be replaced immediately.

NVPM Revision level updated from *n* to *N*

Explanation: The format of the NVPM has changed as a result of installing a newer program card (containing updated firmware). However, all subsystem configuration information has been retained.

NVPM User Interface Parameters component initialized to default settings

Explanation: Terminal setting information has been lost.

To correct this problem, enter the SHOW THIS_CONTROLLER and SHOW OTHER_CONTROLLER commands to determine the current terminal settings. Compare the terminal settings with the CONFIGURATION.INFO output information, and use the SET THIS_CONTROLLER and SET OTHER_CONTROLLER commands to restore terminal settings. If the error persists, replace the controller module.

The following NVPM Configuration Information component elements were initialized to default settings: |*n* ...

Explanation: The settings given by *n* have been initialized in connection with another NVPM error. To clear this error, perform the following procedure:

1. Enter the following commands:

```
CLI> SHOW DEVICES
CLI> SHOW UNITS
CLI> SHOW STORAGESETS
```

2. Compare the information displayed with a printout of the CONFIGURATION.INFO file or with a copy of the most current configuration.
3. Reconfigure the necessary devices, units, or storagesets. (See the CLI commands described in Appendix E.)

If the error persists, replace the controller module.

CAUTION

Replace the controller immediately if any of the following NVPM messages occur. Do not continue to use the controller.

- All NVPM components initialized to their default settings
- Controller Characteristics component reformat failed during NVPM Revision Level 1 to 2 reformat
- Host Access Disabled
- NVPM Cache information component initialized to default settings
- NVPM Controller Characteristics component initialized to default settings

- NVPM Recursive Bugcheck Information component initialized to default settings
- NVPM System Information Page component initialized to default settings
- NVPM Volume Serial Number component initialized to default settings
- The following NVPM Manufacturing Failure Information component elements were initialized to default settings: */...list of component elements/*
- Unknown NVPM Revision Level
- Unknown reformat stage encountered during NVPM Revision Level 1 to 2 reformat
- Nonvolatile parameter/Write Journal Memory backup battery has failed

---

Shelf Messages

This section contains error messages that can display when a shelf problem occurs.

Controller shelf fixed

Explanation: The controller shelf has been correctly repaired.

Controller shelf has a bad power supply or fan

Explanation: Troubleshoot the system to isolate and replace the failed component.

Shelf *xx* fixed

Explanation: Shelf number *xx* has been correctly repaired.

Shelf *xx* has a bad power supply or fan

Explanation: Troubleshoot the system to isolate and replace the failed component.

SWAP signal cleared - all SWAP interrupts re-enabled

Explanation: This message indicates that the swap signal is now cleared.

Unable to clear SWAP signal on shelf *xx* - all SWAP interrupts disabled

Explanation: The subsystem is unable to clear the SWAP signal for a swapped device, where *xx* is the shelf number. This could indicate an unsupported SBB or no power to the device shelf.

Write-Back Cache Messages

This section contains error messages that can display if a write-back cache related fault occurs.

Cache battery charge is low

Explanation: The write-back cache battery is partially discharged. Any unwritten cache data on the cache is flushed. Non-RAIDset units are accessed in read cache mode. RAIDsets are unavailable unless they have access to good batteries on a dual-redundant companion cache. Replace the cache battery.

Cache battery failed diagnostic testing

Explanation: The write-back cache battery has failed the diagnostic tests described in Chapter 7. Any unwritten cache data on the cache is flushed. Non-RAIDset units are accessed in read cache mode. RAIDsets are unavailable unless they have access to good batteries on a dual-redundant companion cache. Replace the cache battery.

Cache module failed diagnostic testing

Explanation: The cache has failed the diagnostic tests described in Chapter 7. Unwritten cache data in the cache module is lost. Replace the write-back cache module.

Cache module failed diagnostic testing - half not accessible

Explanation: Up to 50 percent of the cache memory has failed the diagnostic tests described in Chapter 7. Unwritten cache data in the cache module is lost. Replace the write-back cache module.

Cache module has metadata incompatible with this firmware

Explanation: The subsystem was not properly run down before changing firmware versions. There could be unwritten cache data that cannot be recovered, because the cache metadata format has changed along with the firmware. Correct the problem in one of two ways:

- Restore the previous firmware version, and properly run down the subsystem.
- Enter the CLI command CLEAR_ERRORS INVALID_CACHE, which allows you to use the cache module (although you lose the unwritten cache data).

Cache modules are misconfigured

Explanation: This message is generated in dual-redundant configurations under the following circumstances:

- The companion controller's cache module is not a write-back cache.
- The companion controller's write-back cache is not the same size as this controller's write-back cache.

Cache failover of unwritten cache data is not performed if this message occurs. Correct the problem by replacing/adding cache to make sure both are compatible.

Invalid cache—CLI command set reduced until INVALID_CACHE error is cleared

Explanation: Data in the cache does not match the controller image of cache data. Use the CLEAR_ERRORS INVALID_CACHE command to delete the data.

This controller has a missing cache module

Explanation: The cache module is missing or is not seated properly. (Controller NVMEM indicates that a cache module is expected because there could be unwritten cache data.) You can either find/reseat the module, or enter the CLI command CLEAR_ERRORS INVALID_CACHE, which allows you to run (in write-through mode) without a cache but without accessing RAIDsets.

This controller has an invalid cache module

Explanation: The wrong cache module is present. This means the serial number stored in controller NVMEM and in the cache do not match, and unwritten cache data exists. (This message also can occur for a new, uninitialized module.) Correct the problem in one of two ways:

- Replace this cache with the correct one for this controller.
- Enter the CLI command CLEAR_ERRORS INVALID_CACHE, which allows you to use the random module.

Failure Indications

This section contains error messages that can display with various system failures.

EVL control block allocation failed

Explanation: The controller was unable to allocate enough memory to run the EVL utility. Report this error as a subsystem problem.

FMU control block allocation failed

Explanation: The controller was unable to allocate enough memory to run the FMU utility. Report this error as a subsystem problem.

Last Failure Entry: *n* reset–power-on time reset

Explanation: The power-on time contained in the Last Failure Entry identified by *n* was greater that the current power on time. That condition renders the identified Last Failure Entry unusable, resulting in its being reset to default settings. Reset of Last Failure Entry should be reported as a subsystem problem.

Last Failure Entry: *n* reset–predecessor invalid

Explanation: The Last Failure Entry that immediately precedes the Last Failure Entry identified by *n* was found to be invalid. That condition renders the identified Last Failure Entry unusable, resulting in its being reset to default settings. Reset of Last Failure Entry should be reported as a subsystem problem.

Last Failure Entry: *n* reset–predecessor reset

Explanation: One or more of the Last Failure Entries identified by *n* was reset to default settings. That condition renders the identified Last Failure Entry unusable, resulting in its being reset to default settings. Reset of Last Failure Entry should be reported as a subsystem problem.

Last Failure Entry: *n* reset–previously invalidated

Explanation: During termination of controller operation processing, the Last Failure Entry identified by *n* was reset to default settings for one or more of the reasons listed above. In this case, the reason for resetting the entry is not retained. Reset of Last Failure Entry should be reported as a subsystem problem.

Last Failure Entry: *n* reset–revision mismatch

Explanation: The revision level contained in the Last Failure Entry identified by *n* did not match the revision level understood by the current HSOF firmware. The identified Last Failure Entry therefore is not translatable and was reset to default settings. Reset of Last Failure Entry should be reported as a subsystem problem.

PCMCIA ECC invoked successfully. Replacement of PCMCIA card is recommended

Explanation: During controller initialization, the error correction code (ECC) detected and corrected an error in the firmware on the program card. The errors indicate that the card is becoming unreliable and should be replaced as soon as possible.

System Information Structure and System Information Page (SIP)

This section contains error messages that can display information about system information structure and pages.

System Information Page and Last Failure entries reset to default settings.

Explanation: This message is reported in either of two cases:

- All structures contained in the SIP and the Last Failure entries have been reset to their default settings. This is a normal occurrence for the first time boot following manufacture of the controller and during the transition from one firmware version to another if, and only if, the format of the SIP is different between the two versions.
- All structures contained in the SIP and the Last Failure entries have been reset to their default settings as a result of certain manufacturing configuration activities.

6.3.4 Interactive CLI Messages

The previous sections detailed automatic messages you could encounter. For a list of other messages you could see during interactive use of the CLI, see Appendix E.

6.3.5 Last Failure Logging Messages

If you used the Fault Management Utility (FMU) to enable spontaneous last failure logging (LFL) displays, you could see maintenance terminal messages that begin with "%LFL." With LFL display enabled, the controller spontaneously displays information relevant to the sudden termination of executing firmware, as in the following example (maintenance terminal only):

```
%LFL--HS1CP1> --24-AUG-1994 14:47:11-- Last Failure Code: 018B2580
 Power On Time: 0. Years, 22. Days, 5. Hours, 22. Minutes, 24. Seconds
 Controller Model: HS1CP
 Serial Number: ZG30355555 Hardware Version: 0000(00)
 Controller Identifier:
  Unique Device Number: 000130355555 Model: 40.(28) Class: 1.(01)
 Firmware Version: E28J(FE)
 Node Name: "HS1CP1" INTERNAL Node Number: 7.(07)
 Instance Code: 01010302 Description:
  An unrecoverable hardware detected fault occurred.
 Reporting Component: 1.(01) Description:
  Executive Services
    .
    .
    .
```

When an automatic hardware reset occurs (such as power failure, pressing the Reset button, and so on), the last failure log display is inhibited, because automatic reset does not allow sufficient time to complete the display.

See Chapter 7 for more information on the FMU and last failure logging.

6.3.6 Event Logging Messages

If you used FMU to enable spontaneous event logging (EVL) display, you could see maintenance terminal messages that begin with "%EVL." With EVL display enabled, the controller spontaneously displays error information packet (EIP) information during your maintenance terminal session, as in the following example:

```
%EVL--HS1CP1> --24-AUG-1994 14:57:36-- Instance Code: 4007640A
 Template: 50.(32) Description: INTERNAL Port/Port Driver Event
 Occurred on 24-AUG-1994 at 14:57:36
 Controller Model: HS1CP
 Serial Number: ZG30355555 Hardware Version: 0000(00)
 Controller Identifier:
  Unique Device Number: 000130355555 Model: 40.(28) Class: 1.(01)
 Firmware Version: E28J(FE)
 Node Name: "HS1CP1" INTERNAL Node Number: 7.(07)
 Command Reference Number: 00000000 Sequence Number: 0000000B
 HIS Status: 000E0009 Description:
  VC closed due to NAK DDL retry transmit failure on Path B.
 Error ID: 200A5A40
 Source Node Number: 7.(07) Destination Node Number: 27.(1B)
 INTERNAL Opcode: 0.(00) Description: Reserved
 Virtual Circuit State: 133.(85) Description: VC_CLOSING, ID polling completed
 CI Port/Port Driver (PPD) Opcode: 0.(0000) Description: START
 Instance Code: 4007640A Description:
  INTERNAL Port detected error upon attempting to transmit a packet. This resulted
  in the closure of the Virtual Circuit.
 Reporting Component: 64.(40) Description:
  Host Interconnect Services
    .
    .
    .
```

Event log displays are inhibited during the execution of both CLI commands and utilities invoked from a maintenance terminal. Events that are reported while a maintenance terminal is in use do not appear when the terminal again becomes available. (The %EVL display is lost.)

See Chapter 7 for more information on the FMU and event logging.

6.4 Host Error Logs

Events related to controller and device operation are recorded in the host error log. If the front panel, device LEDs, FMU, or error messages cannot help you determine the cause of a problem, review the host error logs. They provide the greatest level of detail about the controller and connected devices.

6.4.1 Translation Utilities

Translation utilities for log information vary depending on your operating system. The following are the translation tools currently available:

ERF

OpenVMS-based systems have the Errorlog Report Formatter (ERF) to aid in error log translation. ERF provides bit-to-text translation of the (binary) log, and provides the operator with more information about what the log means with respect to controller operation and repair.

The OpenVMS DCL command ANALYZE/ERROR_LOG invokes ERF. For a description of the VMS Analyze Error Log Utility, including more information about this command and its qualifiers, refer to the *VMS Error Log Utility Reference Manual*, or call Digital Multivendor Customer Services.

DECevent

OpenVMS Alpha-based systems include, in addition to ERF, the DECevent™ translation utility. DECevent provides another interface between you and the operating system event logger. Like ERF, DECevent allows you to translate events into ASCII reports (bit-to-text translation). However, DECevent reports are, in general, easier to read and provide more information than ERF error reports. The OpenVMS DIAGNOSE/TRANSLATE command invokes DECevent.

For more information about using DECevent, see the *DECevent Translation and Reporting Utility for OpenVMS User and Reference Guide*.

SWEAT

OpenVMS-based systems have the StorageWorks Errorlog Analyzer Tool (SWEAT) to translate VAX/VMS error log entries from HS1CPs that pertain to storage subsystems. SWEAT has improved summary output, including instance and last fail codes. New qualifiers of /INSTANCE and /LAST_FAIL allow the user to get decodes for specific instance and last fail codes. SWEAT was developed to provide quick, accurate problem diagnosis. For more information on SWEAT, see Appendix F.

6.4.2 Fault Management Utility

FMU allows you to do the following:

- Control the spontaneous LFL and EVL displays.
- Display controller last failure and memory system failure information.

FMU can provide a convenient way to review some error log information during your terminal session. See Chapter 7 for information on running FMU.

6.4.3 Host Error Log Translation

For error log examples, see Appendix C.

Note

Host error log translations are correct as of the date of publication of this manual. However, log information can change with firmware updates. Refer to your firmware release notes for error log information updates.

7

Diagnostics, Exercisers, and Utilities

This chapter discusses the automatic and manual programs available to assist operation and diagnosis of the HS1CP controller subsystem, including the following:

- Initialization and self-test routines
- Disk inline exerciser
- Tape inline exerciser
- VTDPY utility
- CONFIG utility
- Configuration menu
- Code Patch utility
- Firmware licensing system
- Fault Management utility
- CLONE utility

7.1 Initialization and Self-Test Routines

The controller initializes after any of the following conditions:

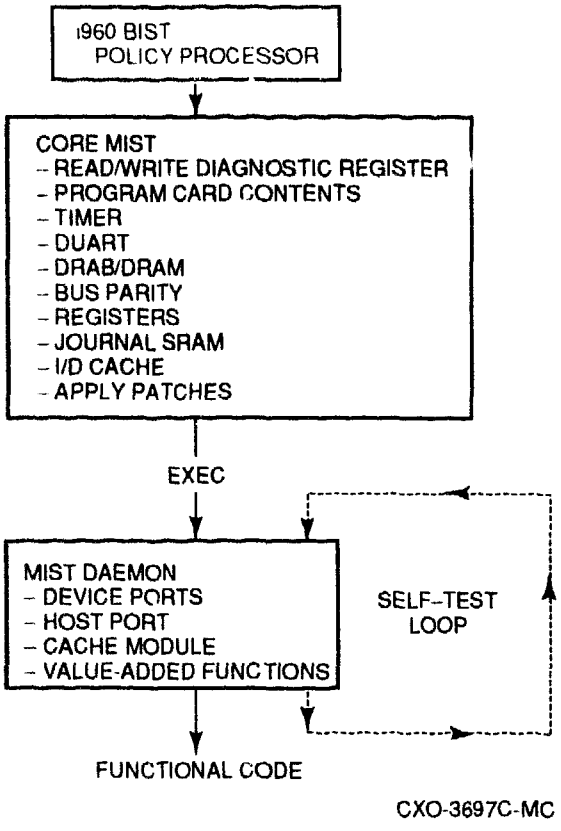
- Power is turned on.
- The firmware resets the controller.
- The operator presses the green Reset button.
- The host clears the controller.

Whenever the controller initializes, it steps through a three-phase series of tests designed to detect any hardware or firmware faults. The three test areas are as follows:

- Built-in self-test
- Core module integrity self-test
- Module integrity self-test DAEMON

Initialization time varies depending on your model of controller and what size and type of cache module, if any, you are running. However, initialization always completes in under 1 minute. Figure 7–1 shows the initialization process.

Figure 7–1 Controller Initialization



7.1.1 Built-In Self-Test

The controller begins initialization by executing its policy processor's internal built-in self-test (**BIST**). BIST always executes upon initialization, because it is an integral part of the Intel® i960 chip microcode. BIST runs entirely from the i960 chip and a small portion of the firmware program card.

Successful completion of BIST means the i960 chip is functioning properly. If BIST fails, the controller shows no activity, and all port indicators on the front panel are off. (The green reset LED is lit continuously.) BIST fails if an incorrect program card is present.

7.1.2 Core Module Integrity Self-Test

After BIST completes successfully, initialization routines and diagnostics expand to testing of the controller module itself. The tests are part of the program card firmware and are known as core module integrity self-test (**MIST**).

Just before beginning core MIST, the controller reads the initial boot record (**IBR**) to determine the address of hardware setup parameters and process control information. After reading the IBR, the firmware within the program card is initialized to the IBR parameters. Program card firmware then executes core MIST as follows:

1. MIST checks the initial state of the read/write diagnostic register.

2. The test validates program card contents by reading each memory location and computing an error detection code (EDC). The test then compares the computed EDC with a predetermined EDC. The program card contents are valid if both EDCs match.
3. Core MIST then tests and/or checks module hardware attached to the buses:
   - Timer operation.
   - DUART operation.
   - DRAB/DRAM (shared memory) operation.
     - The test writes to and reads all legal addresses. Then, boundaries are checked by attempting to access nonexistent addresses. To pass this test, the first two megabytes of memory must test good. If bad segments are found, the bad segments may divide total memory into no more than 16 good, continuous sections.
     - The test selects a device, then checks whether the bus has selected that device.
     - The test verifies that each allowable memory transfer size works and that illegal transfer sizes do not.
   - Bus parity.
   - Registers. The test checks registers for frozen bits.
   - Journal SRAM. The test writes to and reads all journal SRAM addresses.
   - I/D cache.
4. After core MIST successfully tests the program card and bus hardware, the initialization routine loads the firmware into the first two megabytes of controller shared memory. The initialization routine then uses the EDC method to compare the memory contents with the program card to make sure of a successful download.
5. The initialization routine applies any memory-resident code patches for this firmware version, and updates the free memory list accordingly.
6. The i960 chip is initialized to the new parameters (the ones read from the IBR). At this time, control of initialization passes to the firmware executive (EXEC). EXEC runs from controller shared memory.

   If, at any time, a fault occurs during core MIST, the front panel displays a code (refer to Chapter 6). In addition, register contents are displayed at the terminal. The register contents should be reported to Digital Multivendor Customer Services personnel to help with module troubleshooting and repair.

   The following is an example of MIST diagnostic error output:

```
MIST Diag-err Error Code  0000000F
MIST Diag-err Parameter 1 01234567
MIST Diag-err Parameter 2 89ABCDEF
MIST Diag-err Parameter 3 00000001
```

```
MIST Diag-err Master/Cache DRAB register contents:
MIST Diag-err   DSR  2D170CBC, 2C17003D
MIST Diag-err   CSR  00000000, 00000000
MIST Diag-err   DCSR 00000000, 00003402
MIST Diag-err   DER  00001C00, 00001C76
MIST Diag-err   EAR  00800000, 00000000
MIST Diag-err   EDR  00800000, FFFFFFFF
MIST Diag-err   ERR  0000002F, 00000000
MIST Diag-err   RSR  09805432, 09805432
MIST Diag-err   CHC  00000000, 00000000
MIST Diag-err   CMC  00078006, 00000000

MIST Diag-err Read diagnostic register 0  E7FFFFFC
MIST Diag-err Read diagnostic register 1  FFFFFEF0
MIST Diag-err Write diagnostic register 0 C3FFF80F
MIST Diag-err Write diagnostic register 1 0002FF04

MIST Diag-err Halting.
```

7.1.3 Module Integrity Self-Test DAEMON

Once initialization control is passed to EXEC, EXEC calls the diagnostic and execution monitor (**DAEMON**). DAEMON tests the device port hardware, host port hardware, cache module, and value-added functions.

- To test the device ports, DAEMON checks each NCR 53C710 SCSI processor chip. Initialization continues unless *all* SCSI device ports fail testing. In other words, it is possible for the controller to run with only one functioning device port.
- DAEMON tests the host port hardware by exercising the NCR 53C720, the host processor chip. Initialization continues even if the host port tests fail. However, DAEMON stops initialization if the dual universal asynchronous receiver transmitter (**DUART**) test (from core MIST) *and* the host port tests fail.
- DAEMON tests the cache module as follows:

  Note

  The controller still functions if the cache module fails its testing. In this case, the controller uses its on-board shared memory for caching operations.

  – Write-back cache

    DAEMON works in close conjunction with the cache manager, a program that runs in the background of controller firmware. DAEMON first asks the cache manager if unwritten cache data exists on the write-back cache.

    If unwritten cache data does not exist, DAEMON tests the DRAB (memory controller) on the cache module, followed by testing the write-back cache batteries. After DAEMON completes, and functional code takes control of the firmware, the cache manager tests the memory on the cache. At least the first megabyte of the memory must test good, or the cache is declared bad.

    If unwritten cache data exists, DAEMON tests only the batteries, and postpones other tests. Later, during functional code, the cache manager flushes the unwritten cache data after testing the DRAB and memory.

If cache is locked by the other controller (dual-redundant configurations), then all cache DAEMON diagnostics are postponed. During functional code, when the cache manager determines that the cache is unlocked, the cache manager tests the DRAB, batteries, and memory.

- Read cache

  DAEMON tests the DRAB (memory controller) on the read cache module. After DAEMON completes, and functional code takes control of the firmware, the cache manager tests the memory on the cache. At least the first megabyte of the memory must test good, or the cache is declared bad.

- The tests run by DAEMON and the cache manager are summarized in Table 7–1.

Table 7–1 Cache Module Testing

Test	DAEMON	Cache Manager
DRAB	• All memory is initialized. • Full address test.	• No memory is initialized. • Address test on diagnostic pages only.
Memory	• Never invoked.	• Always invokes all memory tests. • Read only or read/write.
Battery	• If battery is charged, full battery test. Otherwise, battery interrupt is checked.	• If battery is charged, full battery test except non-DRAB DRAM refresh circuitry is not tested. Otherwise, battery interrupt is checked.

DAEMON handles all interrupts and errors received during cache module testing. If DAEMON receives any interrupt, it stops initialization. DAEMON and/or the cache manager display any errors as a code on the front panel (refer to Chapter 6).

- To make sure of proper RAID functions, DAEMON tests the XOR operations of the value-added hardware and the associated buffer memory. Initialization stops if this test fails.

After successful test completion, DAEMON releases control. At this time, initialization is finished, and functional controller firmware takes over.

Self-Test

Self-test is a special function of DAEMON, when you set DAEMON to run in a continuous loop. Self-test allows you to diagnose intermittent hardware failures, because the loop continues until an error is detected. [1]

In addition, self-test checks the controller hardware without affecting devices on any ports. Digital recommends you run self-test from the maintenance terminal because the host port disconnects once the controller begins self-test.

For self-test to properly execute, you must have a valid configuration and enable the host paths.

To run self-test, enter one of the following commands (the command you need depends on your configuration, which controller the terminal is connected to, and which controller you wish to test.)

[1] The DAEMON battery tests run only once per self-test. They do not repeat or loop.

CAUTION

Do not use the OVERRIDE_ONLINE qualifier for the SELFTEST command, as customer data may be overwritten.

```
CLI> SELFTEST THIS_CONTROLLER
CLI> SELFTEST OTHER_CONTROLLER
```

See Appendix E for more information on the command and its qualifiers.

When you run self-test, all outstanding I/O operations complete. The controller also attempts to flush the cache. However, even if self-test fails to flush the cache, the program continues to execute.

Self-test halts if it detects a fault. Otherwise, the self-test loop continues until you press the Reset button or cycle the controller power off and on, after which the controller reinitializes.

7.2 Disk Inline Exerciser

The disk inline exerciser (DILX) is a diagnostic tool used to exercise the data transfer capabilities of selected disks connected to an HSICP. DILX exercises disks in a way that simulates a high level of user activity. Using DILX, you can read and write to all customer-available data areas. DILX also runs on CD–ROM *read-only* mode. Thus, DILX is used to determine the health of a controller and the disks connected to it and to acquire performance statistics. You can run DILX from a maintenance terminal, virtual terminal, or VCS.

DILX allows for autoconfiguring of drives. This allows for quick configuring and testing of all units at once. Please be aware that *customer data is lost* by running this test. Digital recommends using only the Auto-Configure option during initial installations.

DILX tests logical units that can consist of storagesets of multiple physical devices. Error reports identify the logical units, not the physical devices. Therefore, if errors occur while running against a unit, its storageset must be reconfigured as individual devices, and then DILX run again, against the individual devices.

There are no limitations on the number of units DILX can test at one time. However, Digital recommends using DILX only when no host activity is present. If you must run DILX during a live host connection, you should limit your testing to no more than half of any controller's units at one time. This conserves controller resources and minimizes performance degradation on the live units you are not testing.

DILX and the tape inline exerciser (TILX) can run concurrently with one initiated from a maintenance terminal and the other from a virtual terminal connection. Digital recommends, however, that the exercisers *not* be run while normal I/O operations are in progress, as system performance degrades due to the heavy load the exercisers impose on the controller.

7.2.1 Invoking DILX

Note

Before running DILX, be sure that all units that you wish to test have been dismounted from the host.

The following describes how to invoke DILX from a maintenance terminal at the CLI> prompt or from a VCS, or from a virtual terminal through a DUP connection:

- To invoke DILX from a maintenance terminal, enter the following command at the CLI> prompt:

  ```
  CLI> RUN DILX
  ```

- To invoke DILX from a maintenance terminal using a VCS, enter the following command at the CLI> prompt:

  ```
  CLI> VCS CONNECT node-name
  ```

  where *node-name* is the controller's SCS node name.

Consult the *VAXcluster Console System User's Guide* for complete details on using a VCS.

Note

The node name *must* be specified for a VCS.

- To invoke DILX from a virtual terminal using a DUP connection, enter the following command (for the OpenVMS operating system):

  ```
  $ SET HOST/DUP/SERVER=MSCP$DUP/TASK=DILX SCS_nodename
  ```

  Specify the contro''er's *SCS node-name* to indicate where DILX executes.

7.2.2 Interrupting DILX Execution

Use the following guidelines to interrupt DILX execution:

Note

The symbol "^" is equivalent to the Ctrl key. You must press and hold the Ctrl key and type the character key given.

Note

Do not use Ctrl/G from a VCS, because it causes VCS to terminate. VCS acts on the sequence, and the sequence never is sent to DILX. Use Ctrl/T when invoking DILX from a VCS.

- Ctrl/G causes DILX to produce a performance summary. DILX continues normal execution without affecting the runtime parameters.
- Ctrl/C causes DILX to produce a performance summary, stop testing, and ask the "reuse parameters" question.
- Ctrl/Y causes DILX to abort. The "reuse parameters" question is not asked.
- Ctrl/T causes DILX to produce a performance summary. DILX then continues executing normally without affecting any of the runtime parameters.

7.2.3 DILX Tests

There are two DILX tests, as follow:

- The Basic Function test
- The User-Defined test

7.2.3.1 Basic Function Test—DILX

The Basic Function test for DILX executes in three or four phases. The four phases are as follow:

- **Initial Write Pass**—Is the only optional phase and always is executed first (if selected). The initial write pass writes the selected data patterns to the entire specified data space or until the DILX execution time limit is reached. Once the initial write pass has completed, it is not reexecuted no matter how long the DILX execution time is set. The other phases are reexecuted on a 10-minute cycle.

- **Random I/O**—Simulates typical I/O activity with random transfers from one byte to the maximum size I/O possible with the memory constraints under which DILX runs. Note that the length of all I/Os is in bytes and is evenly divisible by the sector size (512 bytes). Read, write, access, and erase commands are issued using random logical block numbers (LBNs).

  In the read/write mode, DILX issues the read and write commands in the ratio specified previously under read/write ratio, and issues access and erase commands in the ratio specified previously under access/erase ratio.

  When read-only mode is chosen, only read and access commands are issued.

  If compare operations are enabled, they are performed on write and read commands using the data compare modifier and DILX internal checks. The percentage of compares to perform can be specified. This phase is executed 60 percent of the time. It is the first phase executed after the initial write pass has completed. It is reexecuted at 10-minute intervals with each cycle lasting approximately 6 minutes.

  Intervals are broken down into different cycles. The interval is repeated until the user-selected time interval expires.

```
<-----------------------10 min---------------------------------------->
<------6 min Random I/O----->2 min Data Inten-><--2 min Seek Inten--->
```

- **Data Intensive**—Designed to test disk throughput by selecting a starting LBN and repeating transfers to the next sequential LBN that has not been written to by the previous I/O. The transfer size of each I/O equals the maximum sized I/O that is possible with the memory constraints DILX must run under. This phase continues performing spiraling I/O to sequential tracks.

  Read and write commands are issued in read/write mode. This phase is executed 20 percent of the time after the initial write pass has completed. This phase always executes after the random I/O phase. It is reexecuted at 10-minute intervals with each cycle approximately 2 minutes.

- **Seek Intensive**—Designed to stimulate head motion on the selected disk units. Single sector erase and access commands are issued if the test is write enabled. Each I/O uses a different track on each subsequent transfer. The access and erase commands are issued in the ratio that you selected using the access/erase ratio parameter. This phase is executed 20 percent of the time after the initial write pass has completed. This phase always executes after the data intensive I/O phase. It is reexecuted at 10-minute intervals with each cycle approximately 2 minutes.

7.2.3.2 User-Defined Test—DILX

CAUTION

The User-Defined test must be run *only* by very knowledgeable personnel. Otherwise, customer data can be destroyed.

When this test is selected, DILX prompts you for input to define a specific test. In the DILX User-Defined test, a total of 20 or fewer I/O commands can be defined. Once all of the commands are issued, DILX issues the commands again in the same sequence. This is repeated until the selected time limit is reached. As you build the test, DILX collects the following information from you for each command:

- The I/O command name (write, read, access, erase, or quit). Note that quit is not a command; instead it indicates to DILX that you have finished defining the test.
- The starting LBN.
- The size of the I/O in 512 byte blocks.
- The MSCP command modifiers.

7.2.4 DILX Test Definition Questions

The following text is displayed when running DILX. The text includes questions that are listed in the approximate order that they are displayed on your terminal. These questions prompt you to define the runtime parameters for DILX.

Note

Defaults for each question are given inside [ ]. If you press the Return key as a response to a question, the default is used as the response.

After DILX has been started, the following message describing the Auto-Configure option is displayed:

```
The Auto-Configure option will automatically select, for testing, half
or all of the disk units configured. It will perform a very thorough
test  with *WRITES* enabled. The user will only be able to select the
run time and performance summary options and whether to test a half or
full  configuration. The user will not be able to specify specific
units to test. The Auto-Configure option is only recommended for
initial installations. It is the first question asked.
```

Do you wish to perform an Auto-Configure (y/n) [n] ?

Explanation: Enter "Y" if you wish to invoke the Auto-Configure option.

DILX next displays the following information:

```
If you want to test a dual redundant subsystem, it is recommended that
you pick option 2 on the first controller and then option 2 on the
other controller. Auto-Configure options are:

1. Configure all disk units for testing. This is recommended for a
   single controller subsystem.

2. Configure half of all disk units for testing. This is recommended
   for a dual controller subsystem.

3. Exit Auto-Configure and DILX.
```

Enter Auto-Configure option (1:3) [3] ?

Explanation: This is self-explanatory.

After you enter the desired Auto-Configure option, DILX displays the following caution statement:

```
 **CAUTION**
All data on the Auto-Configured disks will be destroyed. You *MUST* be sure
of yourself.
```

Are you sure you want to continue (y/n) [n] ?

Explanation: This question is asked only if the Auto-Configure option was selected and if the user selected Auto-Configure option 1 or 2 as described in the last question above.

Use All Defaults and Run in Read Only Mode (y/n)[y]?

Explanation: Enter "Y" to use the defaults for DILX and run in read-only mode. Most of the other DILX questions are not asked. Enter "N" and the defaults are not used. You then must answer each question as it is displayed. The following defaults are assumed for all units selected for testing:

- Execution time limit is 10 minutes.
- Performance summary interval is 10 minutes.
- Displaying hard or soft error Error Information Packets (**EIPs**) and end messages is disabled.
- The hard error limit is 65535. Testing stops if the limit is reached.
- A hex dump of the extended error log information is disabled.
- The I/O queue depth is 4. A maximum of 4 I/Os is outstanding at any time.
- The selected test is identical to the Basic Function Test.
- Mode is read-only.
- All user available LBNs are available for testing.
- Data compares are disabled.

Enter the execution time limit in minutes (1:65535)[10]?

Explanation: Enter the desired time you want DILX to run. The default run time is 10 minutes.

Enter performance summary interval in minutes (1:65535)[10]?

Explanation: Enter a value to set the interval for which a performance summary is displayed. The default is 10 minutes.

Include performance statistics in performance summary (y/n)[n]?

Explanation: Enter "Y" to see a performance summary that includes the performance statistics that include the total count of read, write, access, and erase I/O requests and the kilobytes transferred for each command. Enter "N" and no performance statistics are displayed.

Display hard/soft errors (y/n)[n]?

Explanation: Enter "Y" to enable error reporting, including end messages and EIPs. Enter "N" to disable error reporting, including end messages and EIPs. The default is disabled error reporting.

Display hex dump of Event Information Packet requester specific information (y/n)[n]?

Explanation: Enter "Y" to enable the hex dump display of the requester specific information contained in the EIP. Enter "N" to disable the hex dump.

When the hard error limit is reached, the unit will be dropped from testing.
Enter hard error limit (1:65535) [65535] ?

Explanation: Enter a value to specify the hard error limit for all units to test. This question is used to obtain the hard error limit for *all* units under test. If the hard error limit is reached, DILX discontinues testing the unit that reaches the hard error limit. If other units are currently being tested by DILX, testing continues for those units.

When the soft error limit is reached, soft errors no longer display
but testing will continue for the unit. Enter soft error limit (1:65535) [32] ?

Explanation: Enter a value to specify the soft error limit for *all* units under test. When the soft error limit is reached, soft errors are no longer displayed, but testing continues for the unit.

Enter IO queue depth (1:12) [4]?

Explanation: Enter the maximum number of outstanding I/Os for each unit selected for testing. The default is 4.

Enter unit number to be tested?

Explanation: Enter the unit number for the unit to be tested.

Note

When DILX asks for the unit number, it requires the number designator for the disk, where D117 would be specified as unit number 117.

Unit *x* will be write enabled.
Do you still wish to add this unit (y/n) [n]?

Explanation: This is a reminder of the consequences of testing a unit while it is write enabled. This is the last chance to back out of testing the displayed unit. Enter "Y" to write enable the unit. Enter "N" to back out of testing that unit.

Select another unit (y/n) [n]?

Explanation: Enter "Y" to select another unit for testing. Enter "N" to begin testing the units already selected. The system displays the following test selections:

```
***Available tests are:
1.   Basic Function
2.   User Defined Test

Use the Basic Function 99.9% of the time.  The User Defined
test is for special problems only.
```

Enter test number (1:2) [1]?

Explanation: Enter "1" for the Basic Function test or "2" for the User-Defined test. After selecting a test, the system displays the following message:

```
* IMPORTANT *  If you answer yes to the next question, user data
WILL BE destroyed.
```

Write enable disk unit (y/n) [n] ?

Explanation: Enter "Y" to write enable the unit. Write commands are enabled for the currently selected test. Data within your selected LBN range will be destroyed. *Be sure of your actions before answering this question.* This question applies to all DILX tests. Enter "N" to enable read only mode, in which read and access commands are the only commands enabled.

Perform initial write (y/n) [n] ?

Explanation: Enter "Y" to write to the entire user-selected LBN range with the user-selected data patterns. Enter "N" for no initial write pass.

If you respond with "Y," the system performs write operations starting at the lowest user-selected LBN and issues spiral I/Os with the largest byte count possible. This continues until the specified LBN range has been completely written. Upon completion of the initial write pass, normal functions of the Random I/O phase start. The advantage of selecting the initial write pass is that compare host data commands can then be issued and the data previously written to the media can be verified for accuracy. It makes sure that all LBNs within the selected range are accessed by DILX.

The disadvantage of using the initial write pass is that it may take a long time to complete, because a large LBN range was specified. You can bypass this by selecting a smaller LBN range, but this creates another disadvantage in that the entire disk space is not tested. The initial write pass applies only to the Basic Function test.

The write percentage will be set automatically.
Enter read percentage for random IO and data intensive phase (0:100) [67] ?

Explanation: This question is displayed if read/write mode is selected. It allows you to select the read/write ratio to use in the Random I/O and Data Intensive phases. The default read/write ratio is similar to the I/O ratio generated by a typical OpenVMS system.

Enter data pattern number 0=all, 19=user_defined, (0:19) [0] ?

Explanation: The DILX data patterns are used in write commands. This question is displayed when write operations are enabled for the Basic Function or User-Defined tests. There are 18 unique data patterns to select from. These patterns were carefully selected as worst case or most likely to produce errors for disks connected to the controller. (See Table 7–2 for a list of data patterns.) The default uses all 18 patterns in a random method. This question also allows you to create a unique data pattern of your own choice.

Enter the 8-digit hexadecimal user defined data pattern [ ] ?

Explanation: This question is displayed only if you choose to use a user-defined data pattern for write commands. The data pattern is represented in a longword and can be specified with eight hexadecimal digits.

Enter start block number (0:highest_lbn_on_the_disk) [0] ?

Explanation: Enter the starting block number of the area on the disk you wish DILX to test. Zero is the default.

Enter end block number (starting_lbn:highest_lbn_on_the_disk)
[highest_lbn_on_the_disk] ?

Explanation: Enter the highest block number of the area on the disk you wish DILX to test. The highest block number (of that type of disk) is the default.

Perform data compare (y/n) [n] ?

Explanation: Enter "Y" to enable the use of the compare modifier bit with read and write commands. Enter "N" and no data compare operations are done.

This question applies only to the Basic Function test. If the compare modifier is set on write commands, the data is written to the disk. The data is then read from the disk and compared against the corresponding DILX buffers. On read commands, the data is read from the disk into the DILX buffers, read again, then compared against the corresponding DILX buffers. If a discrepancy is found, an error is reported. If the initial write was chosen for the Basic Function test and you enter "Y" to this question, compare host data commands are then enabled and data previously written to the media are verified for accuracy.

Enter compare percentage (1:100) [5] ?

Explanation: This question is displayed only if you choose to perform data compares. This question allows you to change the percentage of read and write commands that have a data compare operation performed. Enter a value indicating the compare percentage. The default is 5.

The erase percentage will be set automatically.
Enter access percentage for Seek Intensive Phase (0:100) [90] ?

Explanation: This question applies only to the Seek Intensive phase if writes are enabled. It allows you to select the percentage of access and erase commands to be issued. Enter a value indicating the access percentage.

Enter command number x (read, write, access, erase, quit) [ ] ?

Explanation: This question applies only to the User-Defined test. It allows you to define command *x* as a read, write, access, or erase command. Enter quit to finish defining the test.

Enter starting LBN for this command (0:highest_lbn_on_the_disk) [ ] ?

Explanation: This question applies only to the User-Defined test. It allows you to set the starting LBN for the command currently being defined. Enter the starting LBN for this command.

Enter the IO size in 512 byte blocks for this command (1:size_in_blocks) [ ] ?

Explanation: This question applies only to the User-Defined test. It allows you to set the I/O size in 512-byte blocks for the command currently being defined. Enter values indicating the I/O size for this command.

Enter in HEX, the MSCP Command Modifiers[0] ?

Explanation: This question applies only to the User-Defined test. It allows you to specify the MSCP command modifiers. You must understand the meaning of the MSCP command modifiers before you enter any value other than the default.

Reuse parameters (stop, continue, restart, change_unit) [stop] ?

Explanation: This question is displayed after the DILX execution time limit expires, after the hard error limit is reached for every unit under test, or after you enter Ctrl/C. These options are as follows:

- **Stop**—DILX terminates normally.
- **Continue**—DILX resumes execution without resetting the remaining DILX execution time or any performance statistics. If the DILX execution time limit has expired, or all units have reached their hard error limit, DILX terminates.
- **Restart**—DILX resets all performance statistics and restarts execution so that the test performs exactly as the one that just completed. However, there is one exception. If the previous test was the Basic Function test with the initial write pass and the initial write pass completed, the initial write pass is not performed when the test is restarted.
- **Change_unit**—DILX allows you to drop or add units to testing. For each unit dropped, another unit must be added, until all units in the configuration have been tested. The unit chosen will be tested with the same parameters that were used for the unit that was dropped from testing. When you have completed dropping and adding units, all performance statistics are initialized and DILX execution resumes with the same parameters as the last run.

Drop unit #x (y/n) [n] ?

Explanation: This question is displayed if you choose to change a unit as an answer to the reuse parameters (previous) question. Enter the unit number that you wish to drop from testing.

The new unit will be write enabled. Do you wish to continue (y/n) [n] ?

Explanation: This question is displayed if you choose to change a unit as an answer to the reuse parameters question. It is asked only if the unit being dropped was write enabled. This question gives you the chance to terminate DILX testing if you do not want data destroyed on the new unit. Enter "N" to terminate DILX.

7.2.5 DILX Output Messages

The following message is displayed when DILX is started:

```
Disk Inline Exerciser - version 2.0
```

This message identifies the internal program as DILX and gives the DILX software version number.

Change Unit is not a legal option if Auto-Configure was chosen.

Explanation: This message is displayed if the user selects the Auto-Configure option and selects the "change unit response" to the "reuse parameters" question. You cannot drop a unit and add a unit if all units were selected for testing.

DILX - Normal Termination.

Explanation: This message is displayed when DILX terminates under normal conditions.

Insufficient resources.

Explanation: Following this line is a second line that gives more information about the problem, which could be one of the following messages:

- Unable to allocate memory.

  DILX was unable to allocate the memory it needed to perform DILX tests. You should run DILX again but choose a lower queue depth and/or choose fewer units to test.

- Cannot perform tests.

  DILX was unable to allocate all of the resources needed to perform DILX tests. You should run DILX again but choose a lower queue depth and/or choose fewer units to test.

- Unable to change operation mode to maintenance.

  DILX tried to change the operation mode from normal to maintenance using the SYSAP$CHANGE_STATE( ) routine but was not successful due to insufficient resources. This problem should not occur. If it does occur, submit a CLD (error report), then reset the controller.

Disk unit *x* does not exist.

Explanation: An attempt was made to allocate a unit for testing that does not exist on the controller.

Unit *x* successfully allocated for testing.

Explanation: All processes that DILX performs to allocate a unit for testing, have been completed. The unit is ready for DILX testing.

Unable to allocate unit.

Explanation: This message should be preceded by a reason the unit could not be allocated for DILX testing.

DILX detected error, code *x*.

Explanation: The "normal" way DILX recognizes an error on a unit is through the reception of an EIP. This loosely corresponds to an MSCP error log. However, the following are some errors that DILX detects without the reception of an EIP:

- Illegal Data Pattern Number found in data pattern header. Unit *x*

  This is code 1. DILX read data from the disk and found that the data was not in a pattern that DILX previously wrote to the disk.

- No write buffers correspond to data pattern Unit *x*.

  This is code 2. DILX read a legal data pattern from the disk at a place where DILX wrote to the disk, but DILX does not have any write buffers that correspond to the data pattern. Thus, the data have been corrupted.

- Read data do not match what DILX thought was written to the media. Unit *x*.

  This is code 3. DILX writes data to the disk and then reads it and compares it against what was written to the disk. This indicates a compare failure. More information is displayed to indicate where in

the data buffer the compare operation failed and what the data was and should have been.

- Compare Host Data should have reported a compare error but did not. Unit *x*

  This is code 4. A compare host data compare was issued in a way that DILX expected to receive a compare error but no error was received.

DILX terminated. A termination, a print summary, or a reuse parameters request was received but DILX is currently not testing any units.

Explanation: The user entered Ctrl/Y (termination request), Ctrl/G (print summary request), or Ctrl/C (reuse parameters request) before DILX had started to test units. DILX cannot satisfy the second two requests so DILX treats all of these requests as a termination request.

DILX will not change the state of a unit if it is not NORMAL.

Explanation: DILX cannot allocate the unit for testing, because it is already in Maintenance mode. (Maintenance mode can be invoked only by the firmware. If another DILX session is in use, the unit is considered in Maintenance mode.)

Unit is not available - if you dismount the unit from the host, it may correct this problem.

Explanation: The unit has been placed online by another user (or host) or the media is not present. The most common reason for this message is that the unit is mounted on the host.

Unit *x* is online to other controller.

Explanation: The unit has been placed online by another user (or host) that is accessing it through the other controller. The most common reason for this message is that the unit is mounted on the host, or DILX is running on the other controller and has the unit allocated for testing.

Soft error reporting disabled. Unit *x*.

Explanation: This message indicates that the soft error limit has been reached and therefore no more soft errors will be displayed for this unit.

Hard error limit reached, unit *x* dropped from testing.

Explanation: This message indicates that the hard error limit has been reached and the unit must be dropped from testing.

Soft error reporting disabled for controller errors.

Explanation: This message indicates that the soft error limit has been reached for controller errors. Thus, controller soft error reporting is disabled.

Hard error limit reached for controller errors. All units dropped from testing.

Explanation: This message is self-explanatory.

Unit is already allocated for testing.

Explanation: This message is self-explanatory.

No drives selected.

Explanation: DILX parameter collection was exited without choosing any units to test.

Maximum number of units are now configured.

Explanation: This message is self-explanatory. (Testing starts after this message is displayed.)

Unit is write protected.

Explanation: The user wants to test a unit with a write commands, or erase commands, or both enabled but the unit is write protected.

The unit status and/or the unit device type has changed unexpectedly.
Unit *x* dropped from testing.

Explanation: The unit status may change if the unit experienced hard errors or if the unit is disconnected. Either way, DILX cannot continue testing the unit.

Last Failure Information follows. This error was NOT produced by running DILX. It represents the reason why the controller crashed on the previous

Explanation: This message may be displayed while allocating a unit for testing. It does not indicate any reason the unit is or is not successfully allocated, but rather represents the reason the controller went down in the previous run. The information that follows this message is the contents of an EIP.

Disk unit numbers on this controller include:

Explanation: After this message is displayed, a list of disk unit numbers on the controller is displayed.

IO to unit *x* has timed out. DILX aborting.

Explanation: One of the DILX I/Os to this unit did not complete within the command timeout interval and when examined, was found not progressing. This indicates a failing controller.

DILX terminated prematurely by user request.

Explanation: Ctrl/Y was entered. DILX interprets this as a request to terminate. This message is displayed and DILX terminates.

Unit is owned by another sysap.

Explanation: DILX could not allocate the unit specified, because the unit is currently allocated by another system application. Terminate the other system application or reset the controller.

Exclusive access is declared for this unit.

Explanation: The unit could not be allocated for testing, because exclusive access has been declared for the unit.

The other controller has exclusive access declared for this unit.

Explanation: This message is self-explanatory.

This unit is marked inoperative.

Explanation: The unit could not be allocated for testing, because the controller internal tables have the unit marked as inoperative.

The unit does not have any media present.

Explanation: The unit could not be allocated for testing, because no media is present.

The RUNSTOP_SWITCH is set to RUN_DISABLED.

Explanation: The unit could not be allocated for testing, because the RUNSTOP_SWITCH is set to RUN_DISABLED. This is enabled and disabled through the Command Line Interpreter (CLI).

Unable to continue, run time expired.

Explanation: A continue response was given to the "reuse parameters" question. This is not a valid response if the run time has expired. Reinvoke DILX.

Unit *x* is not operational.

Explanation: This message is self-explanatory.

When DILX starts to exercise the disk units, the following message is displayed with the current time of day:

```
DILX testing started at: xx:xx:xx
    Test will run for x minutes
    Type ^T(if running DILX through a VCS) or ^G(in all other cases)
      to get a current performance summary
    Type ^C to terminate the DILX test prematurely
    Type ^Y to terminate DILX prematurely
```

7.2.6 DILX End Message Display

To interpret the end message fields correctly, you must contact Digital Multivendor Customer Services. Example 7–1 is an example of a DILX end message display.

Example 7–1 DILX End Message Display

```
Bad Value Added Completion Status for unit x, End message in hex

      Event Code                x
      Op Code                   x
      Cmd Ref Number            x
      Byte Count                x
      Error Byte Count          x
      Sequence Number           x
      Flags                     x
```

7.2.7 DILX Error Information Packet Displays

A DILX EIP display may not include a hex dump of the Requestor Specific Data. This is an option you can select as a DILX parameter.

The EIP will be in one of the following formats that corresponds to MSCP error log formats:

- Controller Error
- Memory Error
- Disk Transfer Error
- Bad Block Replacement (**BBR**) Attempt Error

Examples 7–2 through 7–5 are examples of each display. Each display includes the optional requestor specific information. In all cases, the **instance code**, **template** type, and all requestor specific information correspond to event (error) log device dependent parameters, while everything else has a one-to-one correspondence to error log fields. See Appendix C for a translation of these codes.

Example 7–2 Controller Error

```
Error Information Packet in hex

      Cmd Ref Number          x
      Unit Number             x
      Log Sequence            x
      Format                  x
      Flags                   x
      Event Code              x
      Controller ID           x
      Controller SW ver       x
      Controller HW ver       x
      Multi Unit Code         x
      Instance                      x
      Template Type                 x
      Requestor Information Size    x
      Requestor Specific Data bytes 0 7    xx xx xx xx xx xx xx xx
      Requestor Specific Data bytes 8 15    xx xx xx xx xx xx xx xx
        :
      Requestor Specific Data bytes xx xx    xx xx xx xx xx xx xx xx
```

Example 7–3 Memory Error

```
Error Information Packet in hex

      Cmd Ref Number          x
      Unit Number             x
      Log Sequence            x
      Format                  x
      Flags                   x
      Event Code              x
      Controller ID           x
      Controller SW ver       x
      Controller HW ver       x
      Multi Unit Code         x
      Memory Address          x
      Instance                      x
      Template Type                 x
      Requestor Information Size    x
      Requestor Specific Data bytes 0 7    xx xx xx xx xx xx xx xx
      Requestor Specific Data bytes 8 15    xx xx xx xx xx xx xx xx
        :
      Requestor Specific Data bytes xx xx    xx xx xx xx xx xx xx xx
```

Example 7–4 Disk Transfer Error

```
Error Information Packet in hex

      Cmd Ref Number          x
      Unit Number             x
      Log Sequence            x
      Format                  x
      Flags                   x
      Event Code              x
      Controller ID           x
      Controller SW ver       x
      Controller HW ver       x
      Multi Unit Code         x
      Unit ID[0]              x
      Unit ID[1]              x
      Unit Software Rev       x
      Unit Hardware Rev       x
      Recovery Level          x
      Retry Count             x
      Serial Number           x
      Header Code             x
      Instance                      x
      Template Type                 x
      Requestor Information Size    x
      Requestor Specific Data bytes 0 7    xx xx xx xx xx xx xx xx
      Requestor Specific Data bytes 8 15    xx xx xx xx xx xx xx xx
        :
      Requestor Specific Data bytes xx xx    xx xx xx xx xx xx xx xx
```

Example 7–5 Bad Block Replacement Attempt Error

```
Error Information Packet in hex

      Cmd Ref Number          x
      Unit Number             x
      Log Sequence            x
      Format                  x
      Flags                   x
      Event Code              x
      Controller ID           x
      Controller SW ver       x
      Controller HW ver       x
      Multi Unit Code         x
      Unit ID[0]              x
      Unit ID[1]              x
      Unit Software Rev       x
      Unit Hardware Rev       x
      Replace Flags           x
      Serial Number           x
      Bad LBN                 x
      Old RBN                 x
      New RBN                 x
      Cause                   x
      Instance                      x
      Template Type                 x
      Requestor Information Size    x
      Requestor Specific Data bytes 0 7    xx xx xx xx xx xx xx xx
      Requestor Specific Data bytes 8 15    xx xx xx xx xx xx xx xx
        :
      Requestor Specific Data bytes xx xx    xx xx xx xx xx xx xx xx
```

7.2.8 DILX Data Patterns

Table 7–2 defines the data patterns used with the DILX Basic Function or User-Defined tests. There are 18 unique data patterns. These data patterns were selected as worst case, or the ones most likely to produce errors on disks connected to the controller.

Table 7–2 DILX Data Patterns

Pattern Number	Pattern in Hexadecimal Numbers
1	0000
2	8B8B
3	3333
4	3091
5, shifting 1s	0001, 0003, 0007, 000F, 001F, 003F, 007F, 00FF, 01FF, 03FF, 07FF, 0FFF, 1FFF, 3FFF, 7FFF
6, shifting 0s	FIE, FFFC, FFFC, FFFC, FFE0, FFE0, FFE0, FFE0, FE00, FC00, F800, F000, F000, C000, 8000, 0000
7, alternating 1s, 0s	0000, 0000, 0000, FFFF, FFFF, FFFF, 0000, 0000, FFFF, FFFF, 0000, FFFF, 0000, FFFF, 0000, FFFF
8	B6D9
9	5555, 5555, 5555, AAAA, AAAA, AAAA, 5555, 5555, AAAA, AAAA, 5555, AAAA, 5555, AAAA, 5555, AAAA, 5555
10	DB6C
11	2D2D, 2D2D, 2D2D, D2D2, D2D2, D2D2, 2D2D, 2D2D, D2D2, D2D2, 2D2D, D2D2, 2D2D, D2D2, 2D2D, D2D2
12	6DB6
13, ripple 1	0001, 0002, 0004, 0008, 0010, 0020, 0040, 0080, 0100, 0200, 0400, 0800, 1000, 2000, 4000, 8000
14, ripple 0	FIE, FFFD, FFFB, FFF7, FFEF, FFDF, FFBF, FF7F, FEFF, FDFF, FBFF, F7FF, EFFF, BFFF, DFFF, 7FFF
15	DB6D, B6DB, 6DB6, DB6D, B6DB, 6DB6, DB6D, B6DB, 6DB6, DB6D, B6DB, 6DB6, DB6D
16	3333, 3333, 3333, 1999, 9999, 9999, B6D9, B6D9, B6D9, B6D9, FFFF, FFFF, 0000, 0000, DB6C, DB6C
17	9999, 1999, 699C, E99C, 9921, 9921, 1921, 699C, 699C, 0747, 0747, 0747, 699C, E99C, 9999, 9999
18	FFFF
Default—Use all of the above patterns in a random method	

7.2.9 DILX Examples

This section provides DILX examples using different options.

7.2.9.1 DILX Example—Using All Defaults

In Example 7–6, DILX is run using all defaults. DILX is executed in read-only mode. No data on the units under test are destroyed. The entire user-available LBN range on each disk is accessible for DILX testing. DILX was invoked from a maintenance terminal.

Example 7–6 Using All Defaults—DILX

```
CLI> SHOW DISK
Name          Type          Port Targ  LUN                    Used by
------------------------------------------------------------------------------

DISK100       disk             1    0    0                    D10
DISK120       disk             1    2    0                    D12
DISK140       disk             1    4    0                    D14
DISK210       disk             2    1    0                    D21
DISK230       disk             2    3    0                    D23
DISK610       disk             6    1    0                    D61
DISK630       disk             6    3    0                    D63

CLI> RUN DILX

Disk Inline Exerciser - version 2.0

The Auto-Configure option automatically selects, for testing, half or
all of the disk units configured. It performs a very thorough test with
*WRITES* enabled. The user is able to select only the run time and
performance summary options and whether to test a half or full
configuration. The user will not be able to specify specific units to test.
The Auto-Configure option is recommended only for initial installations.

Do you wish to perform an Auto-Configure (y/n) [n] ? N

Use all defaults and run in read only mode (y/n) [y] ? Y
Disk unit numbers on this controller include:
    10
    12
    14
    21
    23
    61
    63
Enter unit number to be tested ? 10
Unit 10 successfully allocated for testing
Select another unit (y/n) [n] ? Y
Enter unit number to be tested ? 12
Unit 12 successfully allocated for testing
Select another unit (y/n) [n] ? N

   DILX testing started at: 13-JAN-1993 04:47:57
    Test will run for 10 minutes
    Type ^T(if running DILX through VCS) or ^G(in all other cases)
      to get a current performance summary
    Type ^C to terminate the DILX test prematurely
    Type ^Y to terminate DILX prematurely

   DILX Summary at 13-JAN-1993 04:49:14
   Test minutes remaining: 9, expired: 1

Unit 10     Total IO Requests 4530
  No errors detected
Unit 12     Total IO Requests 2930
  No errors detected
Reuse Parameters (stop, continue, restart, change_unit) [stop] ?

DILX - Normal Termination
CLI>
```

7.2.9.2 DILX Example—Using All Functions

In Example 7–7, all functions are chosen for DILX. DILX was invoked from the virtual terminal using the DUP connection from an OpenVMS system. This is an extensive (long) run because the initial write pass was chosen, and because there

was enough time for the initial write pass to complete and for normal testing to continue for a reasonable length of time after the initial write pass.

CAUTION

This test writes to disks. All user data will be destroyed.

Example 7–7 All Functions—DILX

```
$ SHOW CLUSTER/CONTINUOUS
View of Cluster from system ID 9038  node: ENGHRN          7-APR-1993 14:54:01
      SYSTEMS        MEMBERS
  NODE    SOFTWARE    STATUS
 ENGHRN  VMS V5.5  MEMBER
 FORCE   HSC V700
 WODWND  VMS V5.5  MEMBER
 CYMBAL  VMS V5.5  MEMBER
 LUTE    VMS V5.5  MEMBER
 MASS2   HS1CP HS1CP2
 MASS1   HS1CP HS1CP1
(Entered Ctrl/C here.)
DUP> SET HOST/DUP/SERVER=MSCP$DUP MASS1/TASK=DILX
%HSCPAD-I-LOCPROGEXE, Local program executing - type ^\ to exit
Disk Inline Exerciser - version 2.0
The Auto-Configure option automatically selects, for testing, half or
all of the disk units configured. It performs a very thorough test with
*WRITES* enabled. The user is able to select only the run time and
performance summary options and whether to test a half or full
configuration. The user is not able to specify specific units to test.
The Auto-Configure option is only recommended for initial installations.
Do you wish to perform an Auto-Configure (y/n) [n] ?
Use all defaults and run in read only mode (y/n) [y] ? N
Enter execution time limit in minutes (1:65535) [10] ? 45
Enter performance summary interval in minutes (1:65535) [10] ? 45
Include performance statistics in performance summary (y/n) [n] ? Y
Display hard/soft errors (y/n) [n] ? Y
Display hex dump of Error Information Packet requester specific
information (y/n) [n] ? Y
When the hard error limit is reached, the unit will be dropped from testing.
Enter hard error limit (1:65535) [65535] ?
When the soft error limit is reached, soft errors no longer will
be displayed but testing will continue for the unit.
Enter soft error limit (1:65535) [32] ?
Enter IO queue depth (1:20) [4] ? 10
  *** Available tests are:
    1. Basic Function
    2. User Defined
Use the Basic Function test 99.9% of the time. The User Defined
test is for special problems only.
Enter test number (1:2) [1] ? 1
 **CAUTION**
If you answer yes to the next question, user data WILL BE destroyed.
```

(continued on next page)

Example 7–7 (Cont.) All Functions—DILX

```
Write enable disk unit(s) to be tested (y/n) [n] ? Y
The write percentage will be set automatically.
Enter read percentage for Random IO and Data Intensive phase (0:100) [67] ?
Enter data pattern number 0=ALL, 19=USER_DEFINED, (0:19) [0] ?
Perform initial write (y/n) [n] ? Y
The erase percentage will be set automatically.
Enter access percentage for Seek Intensive phase (0:100) [90] ?
Perform data compare (y/n) [n] ? Y
Enter compare percentage (1:100) [5] ?
Disk unit numbers on this controller include:
    10
    12
    14
    21
    23
    61
    63
Enter unit number to be tested ? 10
Unit 10 will be write enabled.
Do you still wish to add this unit (y/n) [n] ? Y
Enter start block number (0:1664214) [0] ?
Enter end block number (0:1664214) [1664214] ?
Unit 10 successfully allocated for testing
Select another unit (y/n) [n] ? Y
Enter unit number to be tested ? 12
Unit 12 will be write enabled.
Do you still wish to add this unit (y/n) [n] ? Y
Enter start block number (0:832316) [0] ?
Enter end block number (0:832316) [832316] ?
Unit 12 successfully allocated for testing
Select another unit (y/n) [n] ? N

   DILX testing started at: 13-JAN-1993 04:52:26
    Test will run for 45 minutes
    Type ^T(if running DILX through VCS) or ^G(in all other cases)
      to get a current performance summary
    Type ^C to terminate the DILX test prematurely
    Type ^Y to terminate DILX prematurely

   DILX Summary at 13-JAN-1993 04:56:20
   Test minutes remaining: 42, expired: 3

Unit 10     Total IO Requests 40794
  Read Count 0   Write Count 40793
  Access Count 0   Erase Count 0
  KB xfer   Read 0   Write 326344   Total 326344
  No errors detected
Unit 12     Total IO Requests 13282
  Read Count 0   Write Count 13281
  Access Count 0   Erase Count 0
  KB xfer   Read 0   Write 106248   Total 106248
  No errors detected
Reuse Parameters (stop, continue, restart, change_unit) [stop] ?

DILX - Normal Termination
CLI>
```

7.2.9.3 DILX Examples—Auto-Configure with All Units

In Example 7–8, DILX is run using the Auto-Configure option with the all units option.

Example 7–8 Auto-Configuration with All Units

```
CLI> RUN DILX

Disk Inline Exerciser - version 2.0

The Auto-Configure option automatically selects, for testing, half or
all of the disk units configured. It performs a very thorough test with
*WRITES* enabled. The user is able to select only the run time and
performance summary options and whether to test a half or full
configuration. The user is not able to specify specific units to test.
The Auto-Configure option is only recommended for initial installations.

Do you wish to perform an Auto-Configure (y/n) [n] ? Y

 If you want to test a dual redundant subsystem, it is recommended that
 you pick option 2 on the first controller and then option 2 on the
 other controller. Auto-Configure options are:

      1. Configure all disk units for testing. This is recommended for a
           single controller subsystem.
      2. Configure half of all disk units for testing, this is recommended
         for a dual controller subsystem.
      3. Exit Auto-Configure and DILX.

Enter Auto-Configure option (1:3) [3] ? 1

     **** C a u t i o n ****

All data on the Auto-Configured disks will be destroyed.
You *MUST* be sure of yourself.

Are you sure you want to continue (y/n) [n] ? Y
Enter execution time limit in minutes (1:65535) [60] ?
Enter performance summary interval in minutes (1:65535) [60] ?
Unit 10 successfully allocated for testing
Unit 12 successfully allocated for testing
Unit 14 successfully allocated for testing
Unit 21 successfully allocated for testing
Unit 23 successfully allocated for testing
Unit 61 successfully allocated for testing
Unit 63 successfully allocated for testing

   DILX testing started at: 13-JAN-1993 04:42:39
    Test will run for 60 minutes
    Type ^T(if running DILX through VCS) or ^G(in all other cases)
      to get a current performance summary
    Type ^C to terminate the DILX test prematurely
    Type ^Y to terminate DILX prematurely

   DILX Summary at 13-JAN-1993 04:44:11
   Test minutes remaining: 59, expired: 1
```

(continued on next page)

Example 7–8 (Cont.) Auto-Configuration with All Units

```
Unit 10     Total IO Requests 9595
  No errors detected
Unit 12     Total IO Requests 5228
  No errors detected
Unit 14     Total IO Requests 10098
  No errors detected
Unit 21     Total IO Requests 9731
  No errors detected
Unit 23     Total IO Requests 5230
  No errors detected
Unit 61     Total IO Requests 11283
  No errors detected
Unit 63     Total IO Requests 5232
  No errors detected
Reuse Parameters (stop, continue, restart, change_unit) [stop] ?

DILX - Normal Termination
CLI>
```

In Example 7–9, DILX is run using the Auto-Configure option with the half of all units option.

Example 7–9 Auto-Configuration with Half of All Units

```
CLI> RUN DILX

Disk Inline Exerciser - version 2.0

The Auto-Configure option automatically selects, for testing, half or
all of the disk units configured. It performs a very thorough test with
*WRITES* enabled. The user is able to select only the run time and
performance summary options and whether to test a half or full
configuration. The user is not able to specify specific units to test.
The Auto-Configure option is only recommended for initial installations.

Do you wish to perform an Auto-Configure (y/n) [n] ? Y

 If you want to test a dual redundant subsystem, it is recommended that
 you pick option 2 on the first controller and then option 2 on the
 other controller. Auto-Configure options are:

      1. Configure all disk units for testing. This is recommended for a
         single controller subsystem.
      2. Configure half of all disk units for testing, this is recommended
         for a dual controller subsystem.
      3. Exit Auto-Configure and DILX.

Enter Auto-Configure option (1:3) [3] ? 2

     **** C a u t i o n ****

All data on the Auto-Configured disks will be destroyed.
You *MUST* be sure of yourself.

Are you sure you want to continue (y/n) [n] ? Y
Enter execution time limit in minutes (1:65535) [60] ?
Enter performance summary interval in minutes (1:65535) [60] ?
Unit 12 successfully allocated for testing
Unit 21 successfully allocated for testing
Unit 61 successfully allocated for testing
```

(continued on next page)

Example 7–9 (Cont.) Auto-Configuration with Half of All Units

```
   DILX testing started at: 13-JAN-1993 04:39:20
    Test will run for 60 minutes
    Type ^T(if running DILX through VCS) or ^G(in all other cases)
      to get a current performance summary
    Type ^C to terminate the DILX test prematurely
    Type ^Y to terminate DILX prematurely

   DILX Summary at 13-JAN-1993 04:41:39
   Test minutes remaining: 58, expired: 2

Unit 12     Total IO Requests 8047
  No errors detected
Unit 21     Total IO Requests 15239
  No errors detected
Unit 61     Total IO Requests 19270
  No errors detected
Reuse Parameters (stop, continue, restart, change_unit) [stop] ?

DILX - Normal Termination
CLI>
```

7.2.10 Interpreting the DILX Performance Summaries

A DILX performance display is produced under the following conditions:

- When a specified performance summary interval elapses
- When DILX terminates for any conditions except an abort
- When Ctrl/G is entered (or Ctrl/T when running from a VCS)

The performance display has different formats depending on whether performance statistics are requested in the user-specified parameters and if errors are detected.

The following is an example of a DILX performance display when performance statistics were not selected and when no errors were detected:

```
   DILX Summary at 18-JUN-1993 06:18:41
   Test minutes remaining: 0, expired: 6

Unit 1     Total IO Requests 482
  No errors detected
Unit 2     Total IO Requests 490
  No errors detected
```

The following is an example of a DILX performance display when performance statistics were selected and when no errors were detected:

```
   DILX Summary at 18-JUN-1993 06:18:41
   Test minutes remaining: 0, expired: 6

Unit 1     Total IO Requests 482
  Read Count   292   Write Count 168
  Access Count 21   Erase Count 0
  KB xfer   Read 7223   Write 4981   Total 12204
  No errors detected
```

The following is an example of a DILX performance display when performance statistics were not selected and when errors were detected on a unit under test:

```
     DILX Summary at 18-JUN-1993 06:18:41
     Test minutes remaining: 0, expired: 6

❶ Unit 10     Total IO Requests 153259
          No errors detected
❷ Unit 40     Total IO Requests 2161358
          Err in Hex: IC:031A4002  PTL:04/00/00  Key:04  ASC/Q:B0/00  HC:0  SC:1
          Total Errs   Hard Cnt 0   Soft Cnt 1
❸ Unit 55     Total IO Requests 2017193
          Err in Hex: IC:03094002  PTL:05/05/00  Key:01  ASC/Q:18/89  HC:0  SC:1
          Err in Hex: IC:03094002  PTL:05/05/00  Key:01  ASC/Q:18/86  HC:0  SC:1
❹         Total Errs   Hard Cnt 0   Soft Cnt 2
```

where:

❶ Represents the unit number and the total I/O requests to this unit.

❷ Represents the unit number and total I/O requests to this unit.

All values for the following codes are described in Appendix C. This also includes the following items associated with this error, and the total number of hard and soft errors for this unit:

- The controller Instance code (in hex)
- The port/target/LUN (PTL)
- The SCSI Sense Key
- The SCSI ASC and ASQ (ASC/Q) codes
- The total hard and soft count for this error

❸ Represents information about the first two unique errors.

All values for the following codes are described in Appendix C. This also includes the following items associated with this error, and the total number of hard and soft errors for this unit:

- The Controller Instance code (in hex)
- The port/target/LUN (PTL)
- The SCSI Sense (Key)
- The SCSI ASC and ASQ (ASC/Q) codes
- The total hard and soft count for this error

A line of this format may be displayed up to three times in a performance summary. There would be a line for each unique error reported to DILX for this unit, up to three errors.

❹ Represents the total hard and soft errors experienced for this unit.

The following is an example of a DILX performance display when performance statistics were not selected and when a controller error was detected:

```
   DILX Summary at 18-JUN-1993 06:18:41
   Test minutes remaining: 0, expired: 6

Cnt err in HEX  IC:07080064  Key:06  ASC/Q:A0/05  HC:1  SC:0
  Total Cntrl Errs   Hard Cnt 1   Soft Cnt 0
```

```
Unit 1     Total IO Requests 482
  No errors detec ed
Unit 2     Total IO Requests 490
  No errors detected
```

For the previous examples, the following definitions apply. These codes are translated in Appendix C.

- IC—The Controller Instance code
- ASC/Q—The SCSI ASC and ASCQ code associated with this error
- HC—The hard count of this error
- SC—The soft count of this error
- PTL—The location of the unit (port/target/LUN)

The performance displays contain error information for up to three unique errors. Hard errors always have precedence over soft errors. A soft error represented in one display may be replaced with information on a hard error in subsequent performance displays.

7.2.11 DILX Abort Codes

Table 7–3 lists the DILX abort codes and definitions.

Table 7–3 DILX Abort Codes and Definitions

Value	Definition
1	An I/O has timed out.
2	dcb_p->htb_used_count reflects an available HTB to test IOs but none could be found.
3	FAO returned either FAO_BAD_FORMAT or FAO_OVERFLOW.
4	TS$SEND_TERMINAL_DATA returned either an ABORTED or INVALID_BYTE_COUNT.
5	TS$READ_TERMINAL_DATA returned either an ABORTED or INVALID_BYTE_COUNT.
6	A timer is in an unexpected expired state that prevents it from being started.
7	The semaphore was set after a one-shot I/O was issued but nothing was found in the received HTB que.
8	A termination, a print summary, or a reuse parameters request was received when DILX was not testing any units.
9	User requested an abort via Ctrl/Y.

7.2.12 DILX Error Codes

Table 7–4 list the DILX error codes and definitions for DILX-detected errors.

Table 7–4 DILX Error Codes and Definitions

Value	Definition
1	Illegal Data Pattern Number found in data pattern header.
2	No write buffers correspond to data pattern.
3	Read data does not match write buffer.
4	Compare Host Data should have reported a compare error but did not.

7.3 Tape Inline Exerciser

TILX is a diagnostic tool used to exercise the data transfer capabilities of selected tape drives connected to an HS1CP controller. TILX exercises tape drives in a way that simulates a high level of user activity. Thus, TILX can be used to determine the health of the controller and the tape drives connected to it. You can run TILX from a maintenance terminal or from a virtual terminal.

DILX and TILX can run concurrently with one initiated from a maintenance terminal and the other from a virtual terminal connection. Digital recommends, however, that the exercisers *not* be run while normal I/O operations are in progress, as system performance degrades due to the heavy load the exercisers impose on the controller.

7.3.1 Invoking TILX

Note

Before running TILX, be sure that all units you wish to test have been dismounted from the host.

The following describes how to invoke TILX from a maintenance terminal at the CLI> prompt, from a VCS, or from a virtual terminal through the DUP connection.

- To invoke TILX from a maintenance terminal, enter the following command at the CLI> prompt:

  ```
  CLI> RUN TILX
  ```

- To invoke TILX from a maintenance terminal using a VCS, enter the following command at the CLI> prompt:

  ```
  CLI> VCS CONNECT node name
  ```

  where *node name* is the controller's SCS node name.

  Consult the *VAXcluster Console System User's Guide* for complete details on using a VCS.

  Note

  The node name *must* be specified for a VCS.

- To invoke TILX from a virtual terminal, enter the following command (for OpenVMS software):

  ```
  $ SET HOST/DUP/SERVER=MSCP$DUP/TASK=TILX SCS_nodename
  ```

  where *SCS_nodename* indicates where TILX will execute.

7.3.2 Interrupting TILX Execution

Use the following guidelines to interrupt TILX execution:

Note

The symbol "^" is equivalent to the Ctrl key. You must press and hold the Ctrl key and type the character key given.

Note

Do not use Ctrl/G from a VCS because it causes VCS to terminate. VCS acts on the sequence and the sequence is never sent to TILX. Use Ctrl/T when invoking TILX from a VCS.

- Ctrl/G causes TILX to produce a performance summary. TILX continues normal execution without affecting the runtime parameters.
- Ctrl/C causes TILX to produce a performance summary, stop testing, and asks the "reuse parameters" question.
- Ctrl/Y causes TILX to terminate. The "reuse parameters" question is not asked.
- Ctrl/T causes TILX to produce a performance summary. TILX then continues executing normally without affecting any of the runtime parameters.

7.3.3 TILX Tests

There are three TILX tests, as follow:

- The Basic Function test
- The User-Defined test
- The Read Only test

7.3.3.1 Basic Function Test—TILX

The Basic Function test executes a write pass followed by a read pass. The write pass executes in two phases, as follows:

- **Data Intensive**—The first one third of the records are written in this phase. All records written to the tape have a byte count of 16 kilobytes. With this high byte count and the default queue depth, this phase should test the streaming capability (if supported) of the tape unit.
- **Random**—This test is performed for the remaining two-thirds of the selected record count. It consists of writes with random byte counts. Intermixed is the sequence *write, reposition back one record, read*. This sequence performed three times in a row. Tape mark writing also is intermixed in the test.

The write pass is complete when the selected record count is reached, or if the end of tape (EOT) is reached. The tape is rewound and the read pass is started.

The read pass consists of the following three phases:

- **Data Intensive**—Consists of reads of fixed record sizes with a byte count equal to the expected tape record byte count. When tape marks are encountered, forward position commands are issued.

- **Random**—Begins at the point where random sized records were written to the tape. Most reads are issued with a byte count equal to the expected tape record byte count. Occasionally, reads are intermixed with a byte count less than or greater than the expected tape record byte count. When tape marks are encountered, forward position commands are issued.
- **Position Intensive**—Begins halfway down from the start of the area where random sized records are located. In the Position Intensive phase, reads and position commands are intermixed so that the test gradually proceeds toward the EOT. When tape marks are encountered, forward position commands are issued.

In all phases, if the EOT is detected, the tape is rewound to the beginning of tape (BOT), and the write pass is entered again.

7.3.3.2 User-Defined Test—TILX

CAUTION

The User-Defined test must be run *only* by very knowledgeable personnel. Otherwise, customer data can be destroyed.

When the TILX User-Defined test is selected, TILX prompts you for input to allow a specific test to be defined. In a User-Defined test, a total of 20, or fewer, I/O commands can be defined. Once all of the commands are issued, TILX issues the commands again in the same sequence. This is repeated until the selected time limit is reached. As you build the test, TILX collects the following information for each command:

- The I/O command operation (write, read, reposition record, reposition file, write tape mark, rewind, quit. Note that quit is not a command; instead it indicates to TILX that you have finished defining the test).
- The number of times to repeat the command. (Applies only to write, read, and write tape mark.)
- The number of records or file marks to reposition.
- The data pattern to use.
- The direction of reposition operation (toward EOT or BOT).
- The size of the I/O in bytes.
- The TMSCP command modifiers.

7.3.3.3 Read Only Test—TILX

The Read Only test should be used only to verify that a tape is readable. The Read Only test reads records until the EOT or the selected record count is reached. At that point, the tape is rewound and another read pass proceeds. Tape marks are ignored. This test most likely issues reads with incorrect record sizes. If there are record size mismatches, they are ignored. All other errors are recorded.

7.3.4 TILX Test Definition Questions

The following section lists the questions that TILX asks to collect the parameters needed to perform a TILX test. Each of the following sections discusses specific TILX questions. The test questions are listed in the approximate order that they are displayed on your terminal. These questions prompt you to define the runtime parameters for TILX.

Note

Defaults for each question are given inside [ ]. If you press the Return key as a response to a question, the default is used as the response.

Use all defaults (y/n) [y] ?

Explanation: Enter "Y" to use the defaults for TILX and most of the other TILX questions are not asked. Enter "N" and the defaults are not used. You must then answer each question as it is displayed. The following defaults are assumed for all units selected for testing:

- Execution time limit is 10 minutes.
- Performance summary interval is 10 minutes.
- Displaying performance statistics is disabled.

Note

This does not include total I/O requests.

- Displaying hard/soft EIPs and end messages is disabled.
- Hard error limit is 65535. Testing stops if the limit is reached.
- Hex dump of extended error log information is disabled.
- I/O queue depth is 4. A maximum of 4 I/Os will be outstanding at one time.
- The selected test is identical to the Basic Function Test.
- The record count is 4093.
- All data patterns are used.
- Data compares are disabled.

Enter execution time limit in minutes (1:65535) [10] ?

Explanation: Enter the desired time you want TILX to run. The default run time is 10 minutes.

Enter performance summary interval in minutes (1:65535) [10] ?

Explanation: Enter a value to set the interval for which a performance summary is displayed. The default is 10 minutes.

Include performance statistics in performance summary (y/n) [n] ?

Explanation: Enter "Y" to see a performance summary that includes the performance statistics that include the total count of read and write I/O requests and the kilobytes transferred for each command. Enter "N" and no performance statistics are displayed.

Display hard/soft errors (y/n) [n] ?

Explanation: Enter "Y" to enable error reporting, including end messages and EIPs. Enter "N" to disable error reporting, including end messages and EIPs. The default is disabled error reporting.

Display hex dump of Event Information Packet requester specific information(y/n) [n] ?

Explanation: Enter "Y" to enable the hex dump display of the requester specific information contained in the EIP. Enter "N" to disable the hex dump.

When the hard error limit is reached, the unit will be dropped from testing. Enter hard error limit (1:65535) [65535] ?

Explanation: Enter a value to specify the hard error limit for all units to test. This question is used to obtain the hard error limit for *all* units under test. If the hard error limit is reached, TILX discontinues testing the unit that reaches the hard error limit. If other units are currently being tested by TILX, testing continues for those units.

When the soft error limit is reached, soft errors will no longer be displayed but testing will continue for the unit. Enter soft error limit (1:65535) [32] ?

Explanation: Enter a value to specify the soft error limit for all units under test. If the soft error limit is reached for a unit under test, soft error reporting is disabled for that unit only. However, testing continues for that unit.

Enter IO queue depth (1:20) [4] ?

Explanation: Enter the maximum number of outstanding I/Os for each unit selected for testing. The default is 4.

Enter unit number to be tested ?

Explanation: Enter the unit number for the (tape drive) unit to be tested.

Note

When TILX asks for the unit number, it requires the actual number of the tape, where T177 would be specified as unit number 177.

Is a tape loaded and ready, answer Yes when ready ?

Explanation: This question is self-explanatory.

Select another unit (y/n) [n] ?

Explanation: Enter "Y" to select another unit to test. Enter "N" to begin testing the units selected. The system displays the following test selections:

```
*** Available tests are:
    1.  Basic Function
    2.  User Defined Test
    3.  Read Only

Use the Basic Function test 99.9% of the time. The User-Defined test
is for special problems only.
```

Enter test number (1:3) [1] ?

Explanation: This question allows you to pick which TILX test you want to run on all selected units. The following questions define the TILX tests.

Enter data pattern number 0=all, 19=user_defined, (0:19) [0] ?

Explanation: The TILX data patterns are used in write commands. This question is displayed for the Basic Function and User-Defined tests. There are 18 unique data patterns from which to select. These patterns were carefully selected as worst case or most likely to produce errors for tapes connected to the controller. (See Table 7–5 for a list of the data patterns.) The default uses all 18 patterns in a random method. This question also allows you to create a unique data pattern of your choice.

Enter record count (1:4294967295) [4096] ?

Explanation: Enter the number of records to write to the tape.

Note

The record count does not include tape marks that are intermixed with the records written to the tape in the Basic Function test.

Enter the 8-digit hexadecimal user defined data pattern [ ] ?

Explanation: This question is displayed only if you choose to use a User-Defined data pattern for write commands. The data pattern is represented in a longword and can be specified with eight hexadecimal digits.

Perform data compare (y/n) [n] ?

Explanation: Enter "Y" to enable the compare modifier bit with the read and write commands. This question applies only to the Basic Function test. If the compare modifier is set on write commands, the data is written to the tape. The data is then read from the tape and compared against the corresponding TILX buffers. On read commands, the data is read from the tape into the TILX buffers, read again, and then compared against the corresponding TILX buffers. If a discrepancy is found, an error is reported. Enter "N" and the compare modifier bit is disabled. The default is to have the bit disabled.

Enter compare percentage (1:100) [2] ?

Explanation: This question is displayed only if you choose to perform data compares. It allows you to enter the percentage of read and write commands that have a data compare operation performed.

Enter command number *x* (red, wrt, rew, wtm, rpr, rpf, quit) [ ] ?

Explanation: This question applies only to the User-Defined test. It allows you to define command *x* as a read, write, rewind, write tape mark, reposition records, or reposition file marks. Enter quit to finish defining the test.

Reposition towards EOT (y=EOT/n=BOT) [y] ?

Explanation: If you specify the reposition records or reposition file marks command in the User-Defined test, this question is displayed. Enter the direction of the reposition operation you want, either towards the end of tape (EOT) or at the beginning of tape (BOT).

Enter number of records to reposition (1:255) [1] ?

Explanation: If you specify the reposition records command in the User-Defined test, this question is displayed. The question is self-explanatory.

Enter number of tape marks to reposition (1:255) [1] ?

Explanation: If you specify the reposition file marks command in the User-Defined test, this question is displayed. The question is self-explanatory.

Enter IO size in bytes (1,65535) [ ] ?

Explanation: This question is asked only in the User-Defined test for read or write commands. The question is self-explanatory.

Enter in HEX, the TMSCP Command Modifiers [0] ?

Explanation: This question applies only to the User-Defined test. It allows you to specify the TMSCP command modifiers. You *must* understand the meaning of the TMSCP command modifiers before entering any value other than the default. Contact Digital Multivendor Customer Services if you wish to use other than default values.

Reuse Parameters (stop, continue, restart, change_unit) [stop] ?

Explanation: This question is displayed after the TILX execution time limit expires after the hard error limit is reached for every unit under test, or after you enter Ctrl/C. The options are as follow:

- **Stop**—TILX terminates normally.
- **Continue**—TILX resumes execution without resetting the remaining TILX execution time or any performance statistics. If the TILX execution time limit has expired, or all units have reached their hard error limit, TILX terminates.
- **Restart**—TILX resets all performance statistics and restarts execution so that the test performs exactly as the test that just completed.
- **Change_unit**—If you select this option, TILX allows you to drop a unit from testing and add a unit to testing. For each unit dropped, another unit must be added until all units in the configuration have been tested. The unit chosen is tested with the same parameters chosen for the unit that was dropped from testing. When you have completed adding and dropping units, all performance statistics are initialized and TILX execution resumes with the same parameters as the last run.

Drop unit #*x* (y/n) [n] ?

Explanation: This question is displayed if you choose to change a unit as an answer to the "reuse parameters" question. It is asked for every unit that was tested. After entering "Y," you are prompted for the unit number. Enter the unit number to drop from testing. Enter "N" if you do not wish to drop a unit from testing.

Note

For each unit dropped from testing, one must be added.

7.3.5 TILX Output Messages

The following message is displayed when TILX is started:

```
Tape Inline Exerciser - version 2.0
```

This message identifies the internal program as TILX and gives the TILX software version number.

TILX - Normal Termination.

Explanation: This message is displayed when TILX terminates under normal conditions.

Insufficient resources.

Explanation: Following this line is a second line that gives more information about the problem, which could be one of the following messages:

- Unable to allocate memory.

  TILX was unable to allocate the memory needed to perform TILX tests. You should run TILX again but choose a lower queue depth and/or choose fewer units to test.

- Cannot perform tests.

  TILX was unable to allocate all of the resources needed to perform TILX tests. You should run TILX again but choose a lower queue depth and/or choose fewer units to test.

- Unable to change operation mode to maintenance.

  TILX tried to change the operation mode from normal to maintenance using the SYSAP$CHANGE_STATE( ) routine, but was not successful due to insufficient resources. This problem should not occur. If it does occur, submit an error report. Then reset the controller.

Tape unit *x* does not exist.

Explanation: An attempt was made to allocate a unit for testing that does not exist on the controller.

Unit *x* successfully allocated for testing.

Explanation: All processes that TILX performs to allocate a unit for testing have been completed. The unit is ready for TILX testing.

Unable to allocate unit.

Explanation: This message should be preceded by a reason the unit could not be allocated for TILX testing.

Cannot enable eip notification.

Explanation: This message indicates that TILX was not successful in enabling EIP notification. This should occur only if another copy of TILX is running. Wait for the first copy to finish or terminate the second copy. If there are no copies of TILX running, submit a CLD (error report) and restart the controller.

TILX detected error, code *x*.

Explanation: The "normal" way TILX recognizes an error on a unit is through the reception of an EIP, which loosely corresponds to an error log. However, there are some errors that TILX detects without the reception of an EIP. These errors are as follow:

- Illegal Data Pattern Number found in data pattern header. Unit *x*.

  This is code 1. TILX read data from the tape unit and found that the data was not in a pattern that TILX previously wrote to the tape.

- No write buffers correspond to data pattern. Unit *x*.

  This is code 2. TILX read a legal data pattern from the tape at a place where TILX wrote to the tape, but TILX does not have any write buffers that correspond to the data pattern. Thus, the data have been corrupted.

- Read data do not match what TILX wrote to the media.

  This is code 3. TILX writes data to the tape and then reads it and compares it against what it wrote to the tape. This indicates a compare failure. More information is displayed to indicate where in the data buffer the compare operation failed and what the data was and should have been.

- TILX/Tape record size mismatch.

  This is code 4. This error is detected only on a read pass. Because TILX knows what was written to the tape, TILX expects to encounter the records (of different sizes), tape marks, and the EOT in exactly the same positions as previously written. This error most likely means that the tape unit has a positioning problem.

- A tape mark was detected in a place not expected by TILX.

  This is code 5. This error is detected only on a read pass. Because TILX knows what was written to the tape, TILX expects to encounter the records, tape marks, and the EOT in exactly the same positions as previously written. This error most likely means that the tape unit has a positioning problem.

- Record Data Truncated not generated.

  This is code 6. This error is detected only on a read pass. Occasionally, TILX issues a read with a byte count less than what TILX knows was written to the current tape record. Thus, TILX would expect to receive a Record Data Truncated status. If TILX does

not receive the Record Data Truncated status when expected, this TILX detected error is reported.

- EOT encountered in unexpected position.

  This is code 7. This error is detected only on a read pass. Because TILX knows what was written to the tape, TILX expects to encounter the records, tape marks, and the EOT in exactly the same positions as previously written. This error most likely means that the tape unit has a positioning problem.

TILX terminated. A termination, a print summary or a reuse parameters request was received but TILX is currently not testing any units.

Explanation: Ctrl/Y (termination request), Ctrl/G (print summary request), or Ctrl/C (reuse parameters request) was entered before TILX started to test units. TILX cannot satisfy the second two requests, so TILX treats all of these requests as a termination request.

TILX does not change the state of a unit if it is not NORMAL.

Explanation: TILX cannot allocate the unit for testing, because it is already in Maintenance mode. (Maintenance mode can be invoked only by the firmware. If another TILX session is in use, the unit is considered in Maintenance mode.)

Unit is not available - if you dismount the unit from the host, it may correct this problem.

Explanation: The unit has been placed online by another user (or host) or the media is not present.

Soft error reporting disabled. Unit *x*.

Explanation: This message indicates that the soft error limit has been reached and that no more soft errors are printed for this unit.

Hard error limit reached, unit *x* dropped from testing.

Explanation: This message indicates that the hard error limit has been reached and the unit must be dropped from testing.

Soft error reporting disabled for controller errors.

Explanation: This message indicates that the soft error limit has been reached for controller errors. Controller soft error reporting is disabled.

Hard error limit reached for controller errors.
All units dropped from testing.

Explanation: This message is self-explanatory.

Unit is already allocated for testing.

Explanation: This message is self-explanatory.

No drives selected.

Explanation: TILX parameter collection was exited without choosing any units to test.

Maximum number of units are now configured.

Explanation: This message is self-explanatory. (Testing starts after this message is displayed.)

Unit is write protected.

Explanation: The user wants to test a unit with write and/or erase commands enabled but the unit is write protected.

The unit status and/or the unit device type has changed unexpectedly.
Unit *x* dropped from testing.

Explanation: The unit status may change if the unit experienced hard errors or if the unit is disconnected. Either way, TILX cannot continue testing the unit.

Last Failure Information follows. This error was NOT produced by running TILX. It represents the reason the controller crashed on the previous controller run.

Explanation: This message may be displayed while allocating a unit for testing. It does not indicate any reason the unit is or is not successfully allocated, but rather represents the reason the controller went down in the previous run. The information that follows this message is the contents of an EIP.

Tape unit numbers on this controller include:

Explanation: After this message is displayed, a list of tape unit numbers on the controller is displayed.

IO to unit *x* has timed out. TILX aborting.

Explanation: One of the TILX I/Os to this unit did not complete within the command timeout interval and when examined, was found not progressing. This indicates a failing controller.

TILX terminated prematurely by user request.

Explanation: Ctrl/Y was entered. TILX interprets this as a request to terminate. This message is then displayed and TILX terminates.

Unit is owned by another sysap.

Explanation: TILX could not allocate the unit specified, because the unit is currently allocated by another system application. Terminate the other system application or reset the controller.

Exclusive access is declared for this unit.

Explanation: The unit could not be allocated for testing, because exclusive access has been declared for the unit.

The other controller has exclusive access declared for this unit.

Explanation: This message is self-explanatory.

This unit is marked inoperative.

Explanation: The unit could not be allocated for testing, because the controller internal tables have the unit marked as inoperative.

The unit does not have any media present.

Explanation: The unit could not be allocated for testing, because no media is present.

The RUNSTOP_SWITCH is set to RUN_DISABLED.

Explanation: The unit could not be allocated for testing, because the RUNSTOP_SWITCH is set to RUN_DISABLED. This is enabled and disabled through the Command Line Interpreter (CLI).

Unable to continue, run time expired.

Explanation: A continue response was given to the "reuse parameters" question. This is not a valid response if the run time has expired. Reinvoke TILX.

When TILX starts to exercise the tape units, the following is displayed with the current time of day:

```
TILX testing started at: xx:xx:xx
    Test will run for x minutes
    Type ^T(if running TILX through a VCS) or ^G(in all other cases)
      to get a current performance summary
    Type ^C to terminate the TILX test prematurely
    Type ^Y to terminate TILX prematurely
```

7.3.6 TILX End Message Display

The Value Added Status field corresponds to the TMSCP end message status. Example 7–10 is an example of a TILX end message display.

Example 7–10 TILX End Message Display

```
Bad Value Added Completion Status for unit x, End message in hex

      Event Code                x
      Op Code                   x
      Cmd Ref Number            x
      End Flags                 x
      Host Xfer Byte Count      x
      Tape Rec Byte Count       x
      Tape Position             x
      Sequence Number           x
```

7.3.7 TILX Error Information Packet Displays

Contact Digital Multivendor Customer Services for assistance in deciphering the EIP fields.

A TILX EIP display may not include a hex dump of the Requestor Specific Data. This is an option you can select for TILX selectable parameters.

The EIP is in one of the following formats that corresponds to MSCP error log formats:

- Controller Error
- Memory Error
- Tape Error

Examples 7–11 through 7–13 are samples of each display. Each display includes the optional requestor specific information. In all cases, the Instance code, template type, and all requestor specific information correspond to event (error) log device dependent parameters, while everything else has a one-to-one correspondence to error log fields. See Appendix C for a translation of these codes.

Example 7–11 Controller Error

```
Error Information Packet in hex

      Cmd Ref Number          x
      Unit Number             x
      Log Sequence            x
      Format                  x
      Flags                   x
      Event Code              x
      Controller ID           x
      Controller SW ver       x
      Controller HW ver       x
      Multi Unit Code         x
      Instance                      x
      Template Type                 x
      Requestor Information Size    x
      Requestor Specific Data bytes 0 7    xx xx xx xx xx xx xx xx
      Requestor Specific Data bytes 8 15    xx xx xx xx xx xx xx xx
        :
      Requestor Specific Data bytes xx xx    xx xx xx xx xx xx xx xx
```

Example 7–12 Memory Error

```
Error Information Packet in hex

      Cmd Ref Number          x
      Unit Number             x
      Log Sequence            x
      Format                  x
      Flags                   x
      Event Code              x
      Controller ID           x
      Controller SW ver       x
      Controller HW ver       x
      Multi Unit Code         x
      Memory Address          x
      Instance                      x
      Template Type                 x
      Requestor Information Size    x
      Requestor Specific Data bytes 0 7    xx xx xx xx xx xx xx xx
      Requestor Specific Data bytes 8 15    xx xx xx xx xx xx xx xx
        :
      Requestor Specific Data bytes xx xx    xx xx xx xx xx xx xx xx
```

Example 7–13 Tape Erro

(continued on next page)

Example 7–13 (Cont.) Tape Error

```
Error Information Packet in hex

      Cmd Ref Number        x
      Unit Number           x
      Log Sequence          x
      Format                x
      Flags                 x
      Event Code            x
      Controller ID         x
      Controller SW ver     x
      Controller HW ver     x
      Multi Unit Code       x
      Unit ID[0]            x
      Unit ID[1]            x
      Unit Software Rev     x
      Unit Hardware Rev     x
      Recovery Level        x
      Retry Count           x
      Position              x
      Formatter SW version  x
      Formatter HW version  x
      Instance                    x
      Template Type               x
      Requestor Information Size  x
      Requestor Specific Data bytes 0 7   xx xx xx xx xx xx xx xx
      Requestor Specific Data bytes 8 15   xx xx xx xx xx xx xx xx
        :
      Requestor Specific Data bytes xx xx   xx xx xx xx xx xx xx xx
```

7.3.8 TILX Data Patterns

Table 7–5 defines the data patterns used with the TILX Basic Function or User-Defined tests. There are 18 unique data patterns. These data patterns were selected as worst case, or the ones most likely to produce errors on tapes connected to the controlle

Table 7–5 TILX Data Pattern Definitions

Pattern Number	Pattern in Hexadecimal Numbers
1	0000
2	8B8B
3	3333
4	3091
5, shifting 1s	0001, 0003, 0007, 000F, 001F, 003F, 007F, 00FF, 01FF, 03FF, 07FF, 0FFF, 1FFF, 3FFF, 7FFF
6, shifting 0s	FIE, FFFC, FFFC, FFFC, FFE0, FFE0, FFE0, FFE0, FE00, FC00, F800, F000, E000, C000, 8000, 0000
7, alternating 1s, 0s	0000, 0000, 0000, FFFF, FFFF, FFFF, 0000, 0000, FFFF, FFFF, 0000, FFFF, 0000, FFFF, 0000, FFFF
8	B6D9
9	5555, 5555, 5555, AAAA, AAAA, AAAA, 5555, 5555, AAAA, AAAA, 5555, AAAA, 5555, AAAA, 5555, AAAA, 5555
10	DB6C
11	2D2D, 2D2D, 2D2D, D2D2, D2D2, D2D2, 2D2D, 2D2D, D2D2, D2D2, 2D2D, D2D2, 2D2D, D2D2, 2D2D, D2D2
12	6DB6
13, ripple 1	0001, 0002, 0004, 0008, 0010, 0020, 0040, 0080, 0100, 0200, 0400, 0800, 1000, 2000, 4000, 8000
14, ripple 0	FIE, FFFD, FFFB, FFF7, FFEF, FFDF, FFBF, FF7F, FEFF, FDFF, FBFF, F7FF, EFFF, BFFF, DFFF, 7FFF
15	DB6D, B6DB, 6DB6, DB6D, B6DB, 6DB6, DB6D, B6DB, 6DB6, DB6D, B6DB, 6DB6, DB6D
16	3333, 3333, 3333, 1999, 9999, 9999, B6D9, B6D9, B6D9, B6D9, FFFF, FFFF, 0000, 0000, DB6C, DB6C
17	9999, 1999, 699C, E99C, 9921, 9921, 1921, 699C, 699C, 0747, 0747, 0747, 699C, E99C, 9999, 9999
18	FFFF
Default–Use all of the above patterns in a random method	

7.3.9 TILX Examples

This section provides some TILX examples with different options chosen.

7.3.9.1 TILX Example—Using All Defaults

In Example 7–14, TILX is run using all defaults. This is a semiextensive test even though the test runs for only 10 minutes. The only function not performed is data compares. Data compares are a time consuming operation with tapes. TILX is invoked from a maintenance terminal.

CAUTION

TILX should only be run using scratch tapes. This test writes to the tape and destroys any data that exists on the tape.

Example 7–14 Using All Defaults—TILX

```
CLI> SHOW TAPE
Name          Type          Port Targ  LUN                    Used by
------------------------------------------------------------------------------

TAPE500       tape             5    0    0                    T50
TAPE520       tape             5    2    0                    T52

CLI> RUN TILX

Tape Inline Exerciser - version 2.0

Use all defaults (y/n) [y] ?
Tape unit numbers on this controller include:
    50
    52
Enter unit number to be tested ? 50
Is a tape loaded and ready, answer Yes when ready ? Y
Unit 50 successfully allocated for testing
Select another unit (y/n) [n] ? Y
Enter unit number to be tested ? 52
Is a tape loaded and ready, answer Yes when ready ? Y
Unit 52 successfully allocated for testing
Maximum number of units are now configured

   TILX testing started at: 13-JAN-1993 04:35:08
    Test will run for 10 minutes
    Type ^T(if running TILX through VCS) or ^G(in all other cases)
      to get a current performance summary
    Type ^C to terminate the TILX test prematurely
    Type ^Y to terminate TILX prematurely

   TILX Summary at 13-JAN-1993 04:36:24
   Test minutes remaining: 9, expired: 1

Unit 50   Total IO Requests 868
  No errors detected
Unit 52   Total IO Requests 860
  No errors detected
Reuse Parameters (stop, continue, restart, change_unit) [stop] ?

TILX - Normal Termination
CLI>
```

7.3.9.2 TILX Example—Using All Functions

In Example 7–15, TILX is run using all functions and using a longer run time and higher record count than the default. The performance statistics and a performance summary are displayed every 15 minutes. TILX is invoked from a maintenance terminal. This is an extensive test.

Example 7–15 Using All Functions—TILX

```
CLI> RUN TILX

Tape Inline Exerciser - version 2.0
Enter TILX hex debug flags (0:ffff) [0] ?
```

(continued on next page)

Example 7–15 (Cont.) Using All Functions—TILX

```
Use all defaults (y/n) [y] ? N
Enter execution time limit in minutes (10:65535) [10] ?
Enter performance summary interval in minutes (1:65535) [10] ?
Include performance statistics in performance summary (y/n) [n] ? Y
Display hard/soft errors (y/n) [n] ? Y
Display hex dump of Error Information Packet requester specific
information (y/n) [n] ? Y
When the hard error limit is reached, the unit will be dropped from testing.
Enter hard error limit (1:65535) [32] ?
When the soft error limit is reached, soft errors will no longer be
displayed but testing will continue for the unit.
Enter soft error limit (1:65535) [32] ?
Enter IO queue depth (1:20) [4] ? 6
Suppress caching (y,n) [n] ?
  *** Available tests are:
    1. Basic Function
    2. User Defined
    3. Read Only

Use the Basic Function test 99.9% of the time. The User Defined test
is for special problems only.
Enter test number (1:3) [1] ? 1
Enter data pattern number 0=ALL, 19=USER_DEFINED, (0:19) [0] ?
Enter record count (1:4294967295) [4096] ? 1000
Perform data compare (y/n) [n] ? Y
Enter compare percentage (1:100) [2] ? 1
Tape unit numbers on this controller include:
    50
    52
Enter unit number to be tested ? 50
Is a tape loaded and ready, answer Yes when ready ? Y
Unit 50 successfully allocated for testing
Select another unit (y/n) [n] ? Y
Enter unit number to be tested ? 52
Is a tape loaded and ready, answer Yes when ready ? Y
Unit 52 successfully allocated for testing
Maximum number of units are now configured

   TILX testing started at: 13-JAN-1993 04:38:15
    Test will run for 10 minutes
    Type ^T(if running TILX through VCS) or ^G(in all other cases)
      to get a current performance summary
    Type ^C to terminate the TILX test prematurely
    Type ^Y to terminate TILX prematurely

   TILX Summary at 13-JAN-1993 04:40:14
   Test minutes remaining: 9, expired: 1

Unit 50   Total IO Requests 724
  Read Count 3   Write Count 681   Reposition Count 3
  Total KB xfer 6718   Read 10   Write 6707
  No errors detected
Unit 52   Total IO Requests 731
  Read Count 3   Write Count 687   Reposition Count 3
  Total KB xfer 6743   Read 10   Write 6733
  No errors detected
Reuse Parameters (stop, continue, restart, change_unit) [stop] ?

TILX - Normal Termination
CLI>
```

7.3.10 Interpreting the TILX Performance Summaries

A TILX performance display is produced under the following conditions:

- When the user-selectable performance summary interval elapses
- When TILX terminates for any conditions except an abort
- When Ctrl/G is entered (or Ctrl/T when running from a VCS)

The performance display has different formats depending on whether performance statistics were requested in the user-specified parameters and if errors were detected.

The following is an example of a TILX performance display when performance statistics were not selected and when no errors were detected:

```
   TILX Summary at 18-JUN-1993 06:18:41
   Test minutes remaining: 0, expired: 6

Unit 1     Total IO Requests 482
  No errors detected
Unit 2     Total IO Requests 490
  No errors detected
```

The following is an example of a TILX performance display when performance statistics were selected and when no errors were detected:

```
   TILX Summary at 18-JUN-1993 06:18:41
   Test minutes remaining: 0, expired: 6

Unit 1     Total IO Requests 482
  Read Count   292   Write Count 168
  Access Count 21   Erase Count 0
  KB xfer   Read 7223   Write 4981   Total 12204
  No errors detected
```

The following is an example of a TILX performance display when performance statistics were not selected and when errors were detected:

```
     TILX Summary at 18-JUN-1993 06:18:41
     Test minutes remaining: 0, expired: 6

❶ Unit 10     Total IO Requests 153259
    No errors detected
❷ Unit 40     Total IO Requests 2161368
    Err in Hex: IC:031A4002  PTL:04/00/00  Key:04  ASC/Q:B0/00  HC:0  SC:1
    Total Errs   Hard Cnt 0   Soft Cnt 1
❸ Unit 55     Total IO Requests 2017193
    Err in Hex: IC:03094002  PTL:05/05/00  Key:01  ASC/Q:18/89  HC:0  SC:1
    Err in Hex: IC:03094002  PTL:05/05/00  Key:01  ASC/Q:18/86  HC:0  SC:1
    Total Errs   Hard Cnt 0   Soft Cnt 2
```

where:

❶ Represents the unit number and the total I/O requests to this unit.

❷ Represents the unit number and total I/O requests to this unit.

All values for the following codes are described in Appendix C. This also includes the items associated with this error and the total number of hard and soft errors for this unit:

- The Controller Instance code (in hex)
- The port/target/LUN (PTL)

- The SCSI Sense (Key)
- The SCSI ASC and ASQ (ASC/Q) codes
- The hard and soft count for this error

❸ Represents information about the first two unique errors for this unit.

All values for the following codes are described in Appendix C. This also includes the items associated with this error and the total number of hard and soft errors for this unit:

- The Controller Instance code (in hex)
- The port/target/LUN (PTL)
- The SCSI Sense (Key)
- The SCSI ASC and ASQ (ASC/Q) codes
- The hard and soft count for this error

A line of this format may be displayed up to three times in a performance summary. There would be a line for each unique error reported to TILX for this unit, up to three errors.

The following is an example of a TILX performance display when performance statistics were not selected and when a controller error error was detected:

```
   TILX Summary at 18-JUN-1993 06:18:41
   Test minutes remaining: 0, expired: 6

Cnt err in HEX  IC:07080064  Key:06  ASC/Q:A0/05  HC:1  SC:0
  Total Cntrl Errs   Hard Cnt 1   Soft Cnt 0

Unit 1     Serial Number 1   Total IO Requests 482
  No errors detected
Unit 2     Serial Number 2   Total IO Requests 490
  No errors detected
```

The performance displays contain error information on up to three unique errors. Note that hard errors always have precedence over soft errors. A soft error represented in one display may be replaced with information on a hard error in subsequent performance displays.

7.3.11 TILX Abort Codes

Table 7–6 list TILX abort codes and definitions.

Table 7–6 TILX Abort Codes and Definitions

Value	Definition
1	An I/O has timed out.
2	A HTB was not available to issue an IO when it should have been.
3	FAO returned either FAO_BAD_FORMAT or FAO_OVERFLOW.
4	TS$SEND_TERMINAL_DATA returned either an ABORTED or INVALID_BYTE_COUNT.
5	TS$READ_TERMINAL_DATA returned either an ABORTED or INVALID_BYTE_COUNT.
6	A timer is in an unexpected expired state that prevents it from being started.
7	The semaphore was set after a one-shot I/O was issued but nothing was found in the received HTB que.
8	A termination or a print summary or a reuse parameters request was received when TILX was not testing any units.
9	User requested abort via Ctrl/Y.

7.3.12 TILX Error Codes

Table 7–7 lists TILX defined error codes and definitions for TILX-detected errors.

Table 7–7 TILX Abort Codes and Definitions

Value	Definition
1	Illegal Data Pattern Number found in data pattern header.
2	No write buffers correspond to data pattern.
3	Read data do not match write buffer.
4	TILX/TAPE record size mismatch.
5	A tape mark was detected in a place where it was not expected.
7	EOT encountered in unexpected position.

7.4 VTDPY Utility

The VTDPY utility gathers and displays system state and performance information for the HS1CP controller. The information displayed includes processor utilization, host port activity and status, device state, logical unit state, and cache and I/O performance.

The VTDPY utility requires a video terminal that supports **ANSI** control sequences, such as a VT220, VT320, or VT420 terminal. A graphics display that provides emulation of an ANSI compatible video terminal also can be used. VTDPY can be run either on terminals directly connected to the HS1CP or on terminals connected through a host-based DUP connection.

Note

VCS can be used from a terminal or DECserver attached to the maintenance terminal port on the HS1CP.

VTDPY is conceptually based on the HSC utility of the same name. Though the information displayed differs from the HSC utility due to system implementation differences, a user familiar with the HSC utility should be able to understand this display terminology easily.

The following sections show how to use the VTDPY utility.

7.4.1 How to Run VTDPY

Only one VTDPY session can be run on each controller at one time. Prior to running VTDPY, be sure the terminal is set in NOWRAP mode. Otherwise, the top line of the display scrolls off of the screen.

To initiate VTDPY from a maintenance terminal at the CLI> prompt, enter the following command:

```
CLI> RUN VTDPY
```

To initiate VTDPY from a virtual terminal, refer to Appendix E.

7.4.2 Using the VTDPY Control Keys

Use the following control key sequences to work the VTDPY display:

Table 7–8 VTDPY Control Keys

Control Key Sequence	Function
Ctrl/C	Prompts for commands.
Ctrl/G	Updates the screen (same as Ctrl/Z).
Ctrl/O	Pauses or resumes screen updates.
Ctrl/R	Refreshes current screen display (same as Ctrl/W).
Ctrl/W	Refreshes current screen display (same as Ctrl/R).
Ctrl/Y	Terminates VTDPY and resets screen characteristics.
Ctrl/Z	Updates the screen (same as Ctrl/G).

Note

While VTDPY and a maintenance terminal interface support passing all of the listed control characters, some host-based terminal interfaces restrict passing some of the characters. All of the listed characters have equivalent text string commands.

7.4.3 Using the VTDPY Command Line

VTDPY contains a command line interpreter that is invoked by entering Ctrl/C any time after the program has begun execution. The command line interpreter is used to modify the characteristics of the VTDPY display. Commands also exist to duplicate the function of the control keys listed in Section 7.4.2.

Table 7–9 VTDPY Commands

Command String	Function
DISPLAY CACHE	Use 132 column unit caching statistics display.
DISPLAY DEFAULT	Use default 132 column system performance display.
DISPLAY DEVICE	Use 132 column device performance display.
DISPLAY STATUS	Use 80 column controller status display.
EXIT	Terminates program (same as QUIT).
INTERVAL <seconds>	Changes update interval.
HELP	Displays help message text.
REFRESH	Refreshes the current display.
QUIT	Terminates program (same as EXIT).
UPDATE	Updates screen display.

The keywords in the command strings can be abbreviated to the minimum number of characters that are necessary to uniquely identify the keyword. Entering a question mark (?) after a keyword causes the parser to provide a list of keywords or values that can follow the supplied keyword. The command line interpreter is not case sensitive, so keywords can be entered in uppercase, lowercase, or mixed case.

Upon successful execution of a command other than HELP, the command line interpreter is exited and the display is resumed. Entering a carriage return without a command also exits the command line interpreter and resumes the display. If an error occurs in the command, the user prompts for command expansion help, or the HELP command is entered, the command line interpreter prompts for an additional command instead of returning to the display.

7.4.4 How to Interpret the VTDPY Display Fields

This section describes the major fields in the VTDPY displays. Examples of the VTDPY screens are shown followed by explanations of the fields of the screens.

Example 7-16 DEFAULT VTDPY DISPLAY

```
HS1CP                    S/N: ZG43100001 SW: V25D HW: AX-01
                             88.6% Idle      2891 KB/S        46 Rq/S
Pr    Name Stk/Max Typ Sta CPU% Node HSD4     Port  4 Unit  ASWC KB/S Rd% Wr% Cm% HT% Unit  AS
 0     NULL  0/  0      Rn 88.6 SysId 420011042138   D1000 o^ r  997 100   0   0   9
 2    RECON 10/  1 FNC  Bl  0.0                      D1100 x^ r    0   0   0   0   0
 3      HPT 40/  3 FNC  Bl  6.3 DSSI    Pkts  Pkts/S D1200 o^ r  982 100   0   0   7
 5      CLI 10/  3 DUP  Bl  0.0 RCV     7193    127  D1300 x^ r    0   0   0   0   0
 6    VTDPY 10/  2 DUP  Rn  0.0 ACK     7198    127
16    MDATA 20/  3 FNC  Bl  0.0 NAK        0      0    Target
17   FMTHRD 10/  3 FNC  Bl  0.0 NOR       24      0   01234567
18    DS_HB 10/  2 FNC  Bl  0.0                      P1D D D hH
19      DUP  6/  1 FNC  Bl  0.0                      o2 D D DhH
20      SCS  6/  1 FNC  Bl  0.0                      r3DDDDTDhH
21     MSCP 20/  1 FNC  Bl  0.0                      t4      hH
22     RMGR 40/  2 FNC  Rn  0.3                       5      hH
23       VA 10/  2 FNC  Bl  0.0                       6      hH
24     DS_1 40/  3 FNC  Bl  3.5
25     DS_0 20/  1 FNC  Bl  1.0 Connections Path Status
26      HIS 10/  1 FNC  Bl  0.0  0123456789  0123456789
27 CLIMAIN 20/  8 FNC  Bl  0.0 0..V...MM   0..^...^^
28    NVFOC 10/  2 FNC  Bl  0.0 1           1
29   REMOTE 20/  2 FNC  Bl  0.0 2           2
30      FOC 20/  3 FNC  Bl  0.0 3           3
31    DUART 10/  1 FNC  Bl  0.0
```

Example 7–17 VTDPY DISPLAY STATUS

```
HS1CP                    S/N: ZG43100001 SW: V25D HW: AX-01   30-AUG-1995 10:44:54
 88.7% Idle     2889 KB/S       46 Rq/S                        Up:    0 23:12.21
Pr   Name  Stk/Max Typ Sta  CPU%    Target   Unit  ASWC  KB/S  Unit  ASWC  KB/S
 0    NULL   0/  0      Rn  88.7   01234567  D1000 o^ r   1014
 2   RECON  10/  1 FNC  Bl   0.0 P1D D D hH  D1100 x^ r      0
 5     CLI  10/  3 DUP  Bl   0.0 r3DDDDTDhH  D1300 x^ r      0
 6   VTDPY  10/  2 DUP  Rn   2.2 t4      hH  D1400 o^ r    934
16   MDATA  20/  3 FNC  Bl   0.0  5      hH
17  FMTHRD  10/  3 FNC  Bl   0.0  6      hH
18   DS_HB  10/  2 FNC  Bl   0.0
19     DUP   6/  1 FNC  Bl   0.0 Connections
20     SCS   6/  1 FNC  Bl   0.0  0123456789
21    MSCP  20/  1 FNC  Bl   0.0 0..V...MM
22    RMGR  40/  2 FNC  Rn   0.0 1
23      VA  10/  2 FNC  Bl   0.0 2
24    DS_1  40/  3 FNC  Bl   2.2 3
25    DS_0  20/  1 FNC  Rn   0.0
26     HIS  10/  1 FNC  Bl   0.0 Path Status
27 CLIMAIN  20/  8 FNC  Bl   0.0  0123456789
28   NVFOC  10/  2 FNC  Bl   0.0 0..^...^^
29  REMOTE  20/  2 FNC  Bl   0.0 1
30     FOC  20/  3 FNC  Bl   0.0 2
31   DUART  10/  1 FNC  Bl   0.0 3
```

Example 7–18 VTDPY DISPLAY DEVICE

```
HS1CP                 S/N: ZG43100001 SW: V25D HW: AX-01     30-AUG-1995 10:46:12
                       88.4% Idle     2701 KB/S      111 Rq/S   Up:     0 23:14.39
  01234567                          D100 A^      14     305       0   8   8  0  0  0
P1D D D hH                          D140 A^      13     295       0   8   8  0  0  0
o2 D D DhH                          D120 a^       0       0       0   0   0  0  0  0
r3DDDDTDhH                          D210 A^      13     289       0   7   7  0  0  0
t4       hH                         D230 a^       0       0       0   0   0  0  0  0
 5       hH                         D250 A^      13     276       0   7   7  0  0  0
 6       hH                         D300 A^      13     314       0  11  11  0  0  0
                                    D310 a^       0       0       0   0   0  0  0  0
                                    D320 A^      30     955       0  16  11  0  0  0
                                    D330 A^      13     287       0   7   7  0  0  0
                                    T340 Av       0       0       0   0   0  0  0  0
                                    D350 a^       0       0       0   0   0  0  0  0

Port  Rq/S  RdKB/S  WrKB/S   CR   BR   TR
 1     27     601       0    0    0    0
 2     26     565       0    0    0    0
 3     57    1557       0    0    0    0
 4      0       0       0    0    0    0
 5      0       0       0    0    0    0
 6      0       0       0    0    0    0
```

Display Header

```
    HSICP ❶ S/N: ZG43100001 ❷ SW: V25D ❸ HW:  AX-01 ❹
VTDPY Monitor Copyright © 1994, Digital Equipment Corp.❺
```

Description

This subdisplay provides title information for the display. For 132 column displays, this subdisplay is displayed across one line.

❶ Controller model

❷ Controller serial number

❸ Controller firmware version

❹ Controller hardware version

❺ Copyright notice

Date and Time

```
30-AUG-1995 10:44:54 ❶
  Up:    0  23:14:39 ❷
```

Description

This subdisplay provides time information for the display.

❶ System date and time.

❷ Time in days, hours, minutes, and seconds since the last controller boot.

Controller Performance Summary

88.7% Idle ❶ 2889 KB/S ❷46 Rq/S ❸

Description

This subdisplay provides total system performance information.

❶ Policy processor idle rate.

❷ Cumulative data transfer rate in kilobytes per second. When logical units are being displayed, this is the transfer rate between the host and the controller. When physical devices are being displayed, this is the transfer rate between the controller and the devices.

❸ Cumulative unit or device request rate per second.When logical units are being displayed, this is the request rate between the host and the controller. When physical devices are being displayed, this is the request rate between the controller and the devices.

Controller Threads Display

```
Pr❶ Name❷  Stk/Max❸ Typ❹ Sta❺ CPU%❻
 0    NULL   0/  0       Rn   88.7
 2   RECON  10/  1 FNC   Bl    0.0
 5     CLI  10/  3 DUP   Bl    0.0
 6   VTDPY  10/  2 DUP   Rn    2.2
16   MDATA  20/  3 FNC   Bl    0.0
17  FMTHRD  10/  3 FNC   Bl    0.0
18  DS_HB   10/  2 FNC   Bl    0.0
19     DUP   6/  1 FNC   Bl    0.0
20     SCS   6/  1 FNC   Bl    0.0
21    MSCP  20/  1 FNC   Bl    0.0
22    RMGR  40/  2 FNC   Rn    0.0
23      VA  10/  2 FNC   Bl    0.0
24    DS_1  40/  3 FNC   Bl    2.2
25    DS_0  20/  1 FNC   Rn    0.0
26     HIS  10/  1 FNC   Bl    0.0
27 CLIMAIN  20/  8 FNC   Bl    0.0
28   NVFOC  10/  2 FNC   Bl    0.0
29  REMOTE  20/  2 FNC   Bl    0.0
30     FOC  20/  3 FNC   Bl    0.0
31   DUART  10/  1 FNC   Bl    0.0
```

Description

This display shows the status and characteristics of the active threads in the controller. Threads that are not active, such as DUP Local Program threads, are not displayed until they become active. If the number of active threads exceeds the available space, not all of them are displayed.

❶ The **Pr** column lists the thread priority. The higher the number, the higher the priority.

❷ The **Name** column contains the thread name. For DUP Local Program threads, this is the name used to invoke the program.

❸ The **Stk** column lists the allocated stack size in 512 byte pages. The **Max** column lists the number of stack pages actually used.

❹ The **Typ** column lists the thread type. The following thread types may appear:

- **FNC**—Functional thread. Those threads that are started when the controller boots and never exit.
- **DUP**—DUP local program threads. These threads are active only when run either from a DUP connection or through the command line interpreter's **RUN** command.
- **NULL**—The NULL thread does not have a thread type because it is a special type of thread that executes only when no other thread is executable.

❺ The **Sta** column lists the current thread state. The following thread states may appear:

- **Bl**—The thread is blocked waiting for timer expiration, resources, or a synchronization event.
- **Io**—A DUP local program is blocked waiting for terminal I/O completion.

- **Rn**—The thread is currently executable.

❻ The **CPU%** column lists the percentage of execution time credited to each thread since the last screen update. The values may not add up to exactly 100 percent due to both rounding errors and the fact that there may not be enough room to display all of the threads. An unexpected amount of time may be credited to some threads because the controller's firmware architecture allows code from one thread to execute in the context of another thread without a context switch.

Table 7–10 describes the processes that may appear in the active thread display.

Note

It is possible that different versions of the controller firmware have different threads or different names for the threads.

Table 7–10 Thread Description

Thread Name	Description
CLI	A local program that provides an interface to the controller's command line interpreter thread.
CLIMAIN	The command line interpreter (CLI) thread.
DS_0	A device error recovery management thread.
DS_1	The thread that handles successful completion of physical device requests.
DS_HB	The thread that manages the device and controller error indicator lights and port Reset buttons.
DUART	The console terminal interface thread.
DUP	The DUP protocol server thread.
FMTHREAD	The thread that performs error log formatting and fault reporting for the controller.
FOC	The thread that manages communication between the controllers in a dual controller configuration.
HIS	The SCS protocol interface thread for HS1CP controllers.
HPT	The thread that handles interaction with the host port logic and PPD protocol for HS1CP controllers.
MDATA	The thread that processes metadata for non-transportable disks.
MSCP	The MSCP and TMSCP protocol server thread.
NULL	The process that is scheduled when no other process can be run.
NVFOC	The thread that initiates state change requests for the other controller in a dual controller configuration.
REMOTE	The thread that manages state changes initiated by the other controller in a dual controller configuration.
RMGR	The thread that manages the data buffer pool.
RECON	The thread that rebuilds the parity blocks on RAID 5 storagesets when needed.

(continued on next page)

Table 7–10 (Cont.) Thread Description

Thread Name	Description
SCS	The SCS directory thread.
VA	The thread that provides host protocol independent logical unit services.
VTDPY	A local program thread that provides a dynamic display of controller configuration and performance information.

Internal Bus Performance Display

```
DSSI    Pkts  Pkts/S
RCV     7193     127 ❶
ACK     7198     127 ❷
NAK        0       0 ❸
NOR       24       0 ❹
```

Description

This display indicates the number of packets sent and received through the Internal port and the packet rate.

❶ Packets received from a remote node

❷ Packets sent to a remote node that were ACKed

❸ Packets sent to a remote node that were NAKed

❹ Packets sent to a remote node for which no response was received

Connection Status

```
Connections
 0123456789 ❶
0..V. .MM   ❷
1
2
3
```

Description

This display shows the current status of any connections to a remote node.

❶ Each position in the data field represents one of the possible nodes to which the controller can communicate. To locate the connection status for a given node, use the column on the left to determine the high order digit of the node number, and use the second row to determine the low order digit of the node number. The number of nodes is fixed at 8.

❷ Each location in the grid contains a character to indicate the connection status:

- **C** indicates one connection to that node. In this example, node 12 shows one connection. This usually happens if a host has multiple adapters and it is using more than one adapter for load balancing.
- **M** indicates multiple connections to that node. Because each host system can make a separate connection to each of the disk, tape, and DUP servers, this field frequently shows multiple connections to a host system. In this example, nodes 8, 9, and 14 show multiple connections.
- **V** indicates that only a virtual circuit is open and no connection is present. This happens prior to establishing a connection. It also happens when there is another controller on the same network and when there are systems with multiple adapters connected to the same network. Node 15 demonstrates this principle.
- If a period (.) is in a position corresponding to a node, that node does not have any virtual circuits or connections to this controller.
- A space indicates the address is beyond the visible node range for this controller.

Host Path Status

```
Path Status
 0123456789 ❶
0........^^ ❷
1..^....
2.......X..
3..
```

Description

This display indicates the path status to any system for which a virtual circuit exists.

❶ Each position in the data field represents one of the possible nodes to which the controller can communicate. To locate the path status for a given node, use the column on the left to determine the high order digit of the node number, and use the second row to determine the low order digit of the node number. The number of nodes is fixed at 8.

❷ Each location in the grid contains a character to indicate the path status:

- A circumflex (^) indicates the single Internal Bus path is functioning properly. In this example, nodes 8, 9, and 15 demonstrate this.
- If a period (.) is in a position corresponding to a node, that node does not have any virtual circuits or connections to this controller so either the path status cannot be determined, or neither path is functioning properly.
- A space indicates the address is beyond the visible node range for this controller.

Host Path Status

```
Path Status
 0123456789 ❶
0........^^ ❷
1..^....
2.......X..
3..
```

Description

This display indicates the path status to any system for which a virtual circuit exists.

❶ Each position in the data field represents one of the possible nodes to which the controller can communicate. To locate the path status for a given node, use the column on the left to determine the high order digit of the node number, and use the second row to determine the low order digit of the node number. The number of nodes is fixed at 8.

❷ Each location in the grid contains a character to indicate the path status:

- A circumflex (^) indicates the single Internal Bus path is functioning properly. In this example, nodes 8, 9, and 15 demonstrate this.
- If a period (.) is in a position corresponding to a node, that node does not have any virtual circuits or connections to this controller so either the path status cannot be determined, or neither path is functioning properly.
- A space indicates the address is beyond the visible node range for this controller.

Unit Status (abbreviated)

Unit❶	ASWC❷	KB/S❸	Rd%❹	Wr%❺	Cm%❻	HT%❼
D1000	o^ r	997	100	0	0	9
D1100	x^ r	0	0	0	0	0
D1200	o^ r	982	100	0	0	7
D1300	x^ r	0	0	0	0	0

Description

This subdisplay shows the status of the logical units that are known to the controller firmware. It also indicates performance information for the units. Up to 42 units can be displayed in this subdisplay.

❶ The **Unit** column contains a letter indicating the type of unit followed by the unit number of the logical unit. The list is sorted by unit number. There may be duplication of unit numbers between devices of different types. If this happens, the order of these devices is arbitrary. The following device type letters may appear:

- **D** indicates a disk device.
- **T** indicates a tape device.
- **L** indicates a media loader.
- **C** indicates a CDROM device.
- **F** indicates a device type not listed above.
- **U** indicates the device type is unknown.

❷ The **ASWC** columns indicate respectively the availability, spindle state, write protect state, and cache state of the logical unit.

The availability state is indicated using the following letters:

- **a**—Available. Available to be mounted by a host system.
- **d**—Offline, Disabled by Digital Multivendor Customer Services. The unit has been disabled for service.
- **e**—Online, Exclusive Access. Unit has been mounted for exclusive access by a user.
- **f**—Offline, Media Format Error. The unit cannot be brought available due to a media format inconsistency.
- **i**—Offline, Inoperative. The unit is inoperative and cannot be brought available by the controller.
- **m**—Offline, Maintenance. The unit has been placed in maintenance mode for diagnostic or other purposes.
- **o**—Online. Mounted by at least one of the host systems.
- **r**—Offline, Rundown. The CLI SET NORUN command has been issued for this unit.
- **v**—Offline, No Volume Mounted. The device does not contain media.
- **x**—Online to other controller. Not available for use by this controller.
- A space in this column indicates the availability is unknown.

The spindle state is indicated using the following characters:

- ^—For disks, this symbol indicates the device is at speed. For tapes, it indicates the tape is loaded.
- >—For disks, this symbol indicates the device is spinning up. For tapes, it indicates the tape is loading.
- <—For disks, this symbol indicates the device is spinning down. For tapes, it indicates the tape is unloading.
- v—For disks, this symbol indicates the device is stopped. For tapes, it indicates the tape is unloaded.
- For other types of devices, this column is left blank.

For disks and tapes, a **w** in the write protect column indicates the unit is write protected. This column is left blank for other device types.

The data caching state is indicated using the following letters:

- **b**—Both Read caching and Write Back caching are enabled.
- **r**—Read caching is enabled.
- **w**—Write Back caching is enabled.
- A space in this column indicates caching is disabled.

❸ **KB/S**—This column indicates the average amount of kilobytes of data transferred to and from the unit in the previous screen update interval. This data is available only for disk and tape units.

❹ **Rd%**—This column indicates what percentage of data transferred between the host and the unit were read from the unit. This data is contained only in the **DEFAULT** display for disk and tape device types.

❺ **Wr%**—This column indicates what percentage of data transferred between the host and the unit were written to the unit. This data is contained only in the **DEFAULT** display for disk and tape device types.

❻ **Cm%**—This column indicates what percentage of data transferred between the host and the unit were compared. A compare operation can be accompanied by either a read or a write operation, so this column is not cumulative with read percentage and write percentage columns. This data is contained only in the **DEFAULT** display for disk and tape device types.

❼ **HT%**—This column indicates the cache hit percentage for data transferred between the host and the unit.

Help Example

```
VTDPY> HELP

Available VTDPY commands:

^C - Prompt for commands
^G or ^Z - Update screen
^O - Pause/Resume screen updates
^Y - Terminate program
^R or ^W - Refresh screen
DISPLAY CACHE - Use 132 column unit caching statistics display
DISPLAY DEFAULT - Use default 132 column system performance display
DISPLAY DEVICE - Use 132 column device performance display
DISPLAY STATUS - Use 80 column controller status display
EXIT - Terminate program (same as QUIT)
INTERVAL <seconds> - Change update interval
HELP - Display this help message
REFRESH - Refresh the current display
QUIT - Terminate program (same as EXIT)
UPDATE - Update screen display

VTDPY>
```

Description

This is the sample output from executing the HELP command.

7.5 CONFIG Utility

The CONFIG utility locates and adds devices to the controller. You should run the CONFIG utility whenever new devices are added to the controller.

7.5.1 CONFIG Conventions

The CONFIG utility searches all PTL device combinations to determine what devices exist on the subsystem. It adds all new devices that are found. The CONFIG utility does not initialize these devices, and it does not add units or storagesets.

If a device somewhere in the cluster already has the PTL that the CONFIG utility plans to assign, the program assigns an alpha character after the numbers. For example, if another device is already called DISK100, the program assigns the name DISK100A to the new device. (The program compares DISK100A to other PTLs in the cluster, and if DISK100A already has been used, the program increments to DISK100B and so forth.) This avoids the assignment of duplicate PTLs in the same cluster.

7.5.2 Running the CONFIG Utility

You can run the CONFIG utility either on a virtual terminal or on a maintenance terminal.

Before running the CONFIG utility, you can use the SHOW DEVICES command to verify the list of devices that are currently configured on the controller, as shown in the following example. The example shows the CONFIG utility as it is run on an HS1CP.

```
CLI> SHOW DEVICES
No devices
CLI> RUN CONFIG

Config Local Program Invoked

Config is building its tables and determining what devices exist
on the subsystem. Please be patient.

add disk DISK100   1 0 0
add disk DISK120   1 2 0
add disk DISK140   1 4 0
add disk DISK210   2 1 0
add disk DISK230   2 3 0
add disk DISK500   5 0 0
add disk DISK520   5 2 0
add tape TAPE600   6 0 0
add tape TAPE610   6 1 0

Config - Normal Termination
CLI>
```

```
CLI> SHOW DEVICES
Name            Type           Port Targ  LUN                    Used by
-------------------------------------------------------------------------------

DISK100         disk              1    0    0
DISK120         disk              1    2    0
DISK140         disk              1    4    0
DISK210         disk              2    1    0
DISK230         disk              2    3    0
DISK500         disk              5    0    0
DISK520         disk              5    2    0
TAPE600         tape              6    0    0
TAPE610         tape              6    1    0
CLI>
```

After you run the CONFIG utility, you may have to initialize your containers using the INITIALIZE command as described in Appendix E.

7.6 Configuration Menu

The configuration menu (CFMENU) allows you to quickly configure storage devices attached to the controller. CFMENU presents configuration commands normally entered at the CLI in a menu-driven format.

The CFMENU utility requires a video terminal that supports ANSI control sequences, such as a VT220, VT320, or VT420 terminal. You can use a graphics display to provide emulation of an ANSI compatible video terminal. CFMENU runs on maintenance and virtual terminals. Prior to running CFMENU, be sure the terminal is set in NOWRAP mode. Otherwise, the display may not correctly appear on the screen.

You should have a good understanding of the various CLI commands described in Appendix E before running CFMENU, because CFMENU asks you to choose options for devices, storagesets, and units based on the command qualifiers of the CLI.

This section presents an overview of one configuration, from start to finish. Only one example is described because of the many combinations of choices that can be made during any configuring situation.

7.6.1 Restrictions

The following restrictions apply to CFMENU:

- For dual-redundant configurations, do not run CFMENU on both controllers at the same time.
- There is currently no support for modifying switches (options) on existing devices, storagesets, or units. You must delete and recreate to modify options.

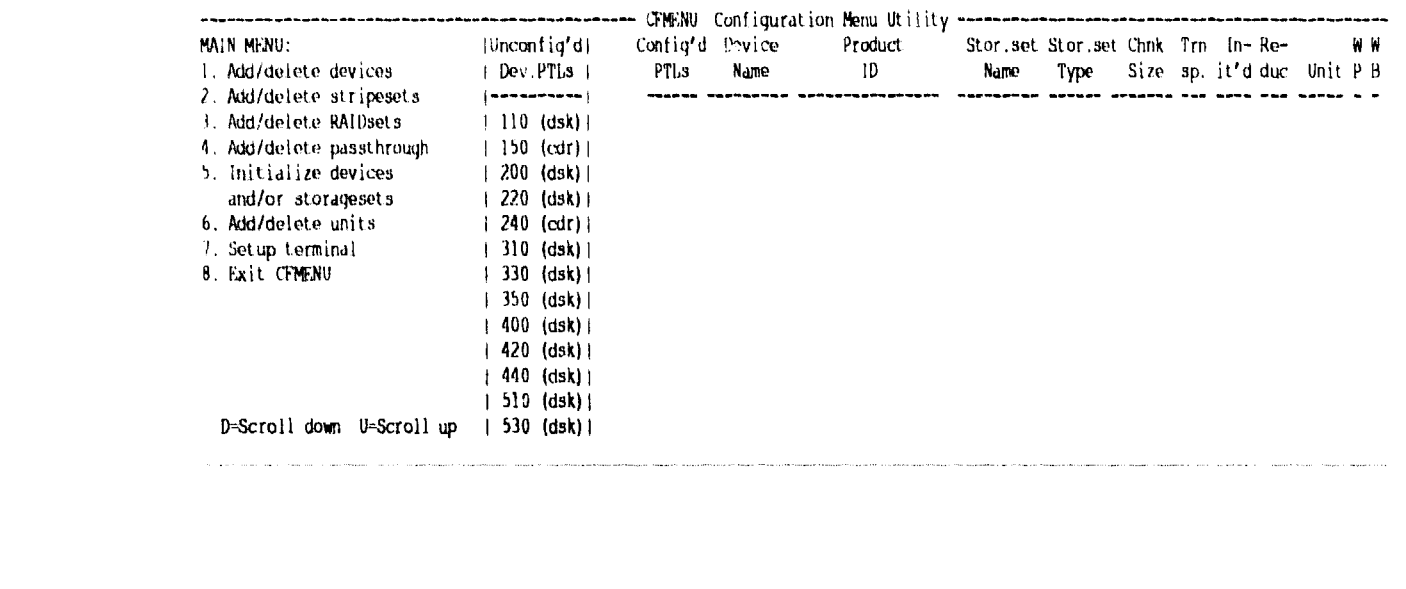
7.6.2 Main Menu

Run CFMENU during a terminal session by invoking it from the CLI prompt:

```
CLI> RUN CFMENU
```

The main menu appears as shown in Figure 7–2. Any attached devices that have not been added to your configuration appears, with their PTLs, in the column to the right of the menu options.

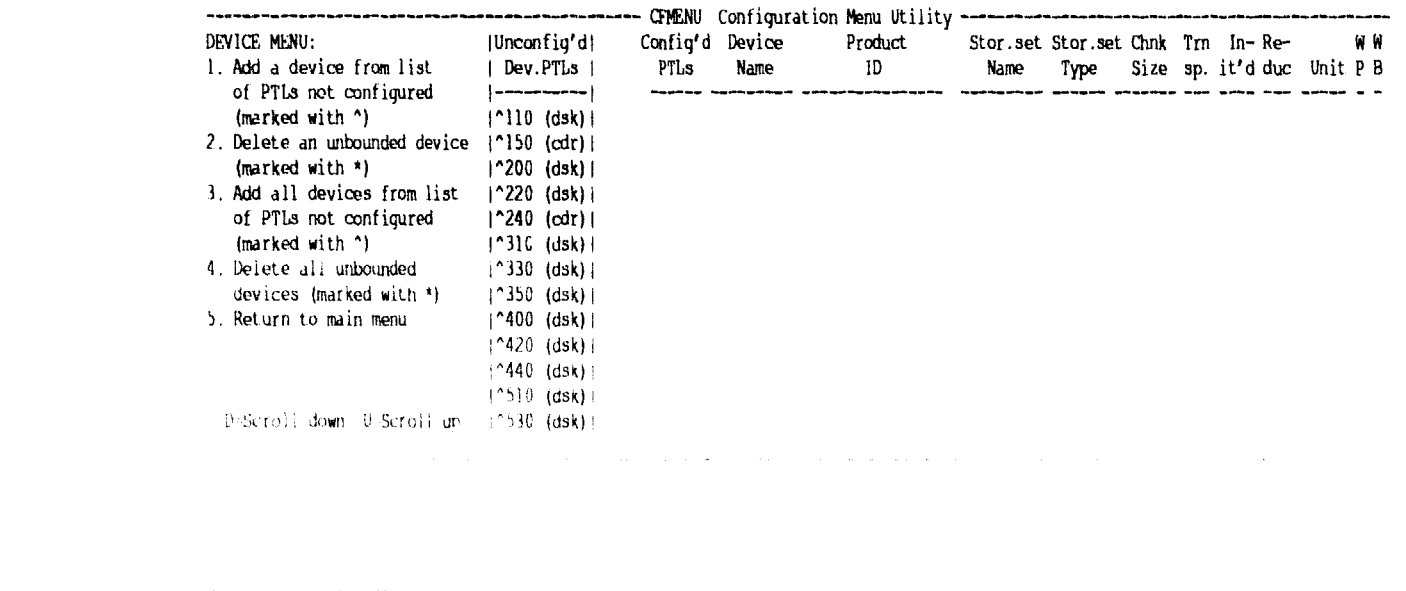
Figure 7–2 CFMENU Main Menu



7.6.3 Adding Devices

From the main menu, enter option 1 (press 1 followed by the Return key). The device menu appears as shown in Figure 7–3. The same list of unconfigured devices remains to the right of the options. When the list is too long to be shown on one screen, you can enter D or U to scroll the information down/up.

Figure 7–3 CFMENU Device Menu (Before Adding Devices)



In Figure 7–3, none of the devices have been added to the configuration yet. If you enter option 1, CFMENU asks you whether (y/n/q)[2] to add each device on the

[2] Yes/no/quit.

list. If you enter option 3, CFMENU adds *all* the unknown devices, in the same way as the CONFIG utility (refer to Section 7.5).

After entering option 1 and adding some devices, the screen resembles Figure 7–4.

Figure 7–4 CFMENU Device Menu

```
------------------------------------------- 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 ^)           |^310 (dsk)|disks:  110 * DISK110 RZ25     (C) DEC                          N    N
2. Delete an unbounded device|^350 (dsk)|        200 * DISK200 RZ35     (C) DEC                          N    N
   (marked with *)           |^400 (dsk)|        220 * DISK220 RZ25     (C) DEC                          N    N
3. Add all devices from list |^510 (dsk)|        330 * DISK330 RZ26     (C) DEC                          N    N
   of PTLs not configured    |^600 (dsk)|        420   DISK420 RZ25     (C) DEC                          N    Y
   (marked with ^)           |          |        440 * DISK440 RZ26     (C) DEC                          N    Y
4. Delete all unbounded      |          |        530 * DISK530 RZ26     (C) DEC                          N    N
   devices (marked with *)   |          |        550 * DISK550 RZ25     (C) DEC                          N    N
5. Return to main menu       |          |        620 * DISK620 RZ25     (C) DEC                          Y    N
                             |          |        640 * DISK640 RZ26     (C) DEC                          N    Y
                             |          |cdrms:  150 * CDROM150 RRD44   (C) DEC
                             |          |        240 * CDROM240 RRD44   (C) DEC
                             |          |loadr:  540 *   LDR540 TL820   (C) DEC

Enter menu choice (1,5) [5] ?
```

CFMENU has added information to the following fields:

- **Config'd PTLs**—Configured PTLs. The program lists the PTL of each device added to your configuration.
- **Device Name**—CFMENU automatically assigns a name to each device. The name contains both the device type and PTL. You cannot override the automatic CFMENU naming convention.
- **Product ID**—Information identifying the device model appears here.
- **Trnsp**—Transportable. This field identifies whether the device is transportable.
- **Init'd**—Initialized. This field identifies whether the device is initialized.

After adding devices, return to the main menu.

7.6.4 Adding Stripesets

Enter option 2 from the main menu to work with stripesets. From the stripeset menu (see Figure 7–5), enter option 1 to add a stripeset. CFMENU asks you how many and which devices from the configured PTLs list you wish to include in the stripeset (2 to 14 devices allowable). In Figure 7–5, a stripeset was created from disks at PTLs 110 and 220.

Figure 7–5 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 ^)                       | 310 (dsk)|disks:  200 ^  DISK200 RZ35     (C) DEC                                 N    N
2. Delete an unbounded         | 350 (dsk)|        330 ^  DISK330 RZ26     (C) DEC                                 N    N
   stripeset (marked by *)     | 400 (dsk)|        420 ^  DISK420 RZ25     (C) DEC                                 N    Y
3. Delete all unbounded        | 510 (dsk)|        440 ^  DISK440 RZ26     (C) DEC                                 N    Y
   stripesets (marked by *)    | 600 (dsk)|        530 ^  DISK530 RZ26     (C) DEC                                 N    N
4. Return to main menu         |          |        550 ^  DISK550 RZ25     (C) DEC                                 N    N
                               |          |        620 ^  DISK620 RZ25     (C) DEC                                 N    N
                               |          |        640 ^  DISK640 RZ26     (C) DEC                                 N    Y
                               |          |cdrms:  150    CDROM150 RRD44   (C) DEC
                               |          |        240    CDROM240 RRD44   (C) DEC
                               |          |loadr:  540      LDR540 TL820   (C) DEC
                               |          |strps:  110    DISK110 RZ25     (C) DEC *  S1     STRP     unk          N
                               |          |        220    DISK220 RZ25     (C) DEC    "      "        "            "

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

CFMENU updates the following fields after adding a stripeset:

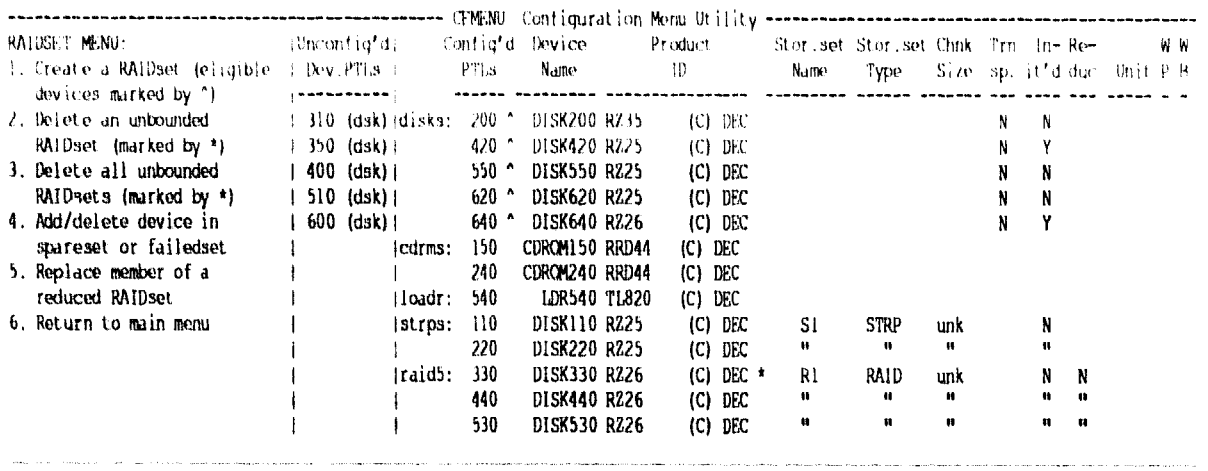
- **Stor.set Name**—Storageset name. CFMENU automatically assigns the name "Sx" to a stripeset. You cannot alter this automatic CFMENU naming convention.
- **Stor.set Type**—Storageset type. This field reads "STRP" for storagesets that are stripesets.
- **Chnk Size**—Chunk size. This field reads "unk" (unknown) until you initialize the stripeset.

After adding stripesets, return to the main menu.

7.6.5 Adding RAIDsets

Enter option 3 from the main menu to work with RAIDsets. From the RAIDset menu (see Figure 7–6), enter option 1 to add a RAIDset. CFMENU asks you how many and which devices from the configured PTLs list you wish to include in the RAIDset (3 to 14 devices allowable). In Figure 7–6, a RAIDset was created from disks at PTLs 330, 440, and 530.

Figure 7–6 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         | 310 (dsk)|disks:  200 ^ DISK200 RZ35    (C) DEC                              N    N
   RAIDset (marked by *)       | 350 (dsk)|        420 ^ DISK420 RZ25    (C) DEC                              N    Y
3. Delete all unbounded        | 400 (dsk)|        550 ^ DISK550 RZ25    (C) DEC                              N    N
   RAIDsets (marked by *)      | 510 (dsk)|        620 ^ DISK620 RZ25    (C) DEC                              N    N
4. Add/delete device in        | 600 (dsk)|        640 ^ DISK640 RZ26    (C) DEC                              N    Y
   spareset or failedset       |          |cdrms:  150   CDROM150 RRD44   (C) DEC
5. Replace member of a         |          |        240   CDROM240 RRD44   (C) DEC
   reduced RAIDset             |          |loadr:  540     LDR540 TL820   (C) DEC
6. Return to main menu         |          |strps:  110   DISK110 RZ25    (C) DEC     S1        STRP      unk          N
                               |          |        220   DISK220 RZ25    (C) DEC     "         "         "            "
                               |          |raid5:  330   DISK330 RZ26    (C) DEC *   R1        RAID      unk          N    N
                               |          |        440   DISK440 RZ26    (C) DEC     "         "         "            "    "
                               |          |        530   DISK530 RZ26    (C) DEC     "         "         "            "    "

Enter menu choice (1,6) [6] ?
```

CFMENU updates the following fields after adding a RAIDset:

- **Stor.set Name**—Storageset name. CFMENU automatically assigns the name "Rx" to a RAIDset. You cannot alter this automatic CFMENU naming convention.
- **Stor.set Type**—Storageset type. This field reads "RAID" for storagesets that are RAIDsets.
- **Chnk Size**—Chunk size. This field reads "unk" (unknown) until you initialize the RAIDset.

7.6.6 Adding to Sparesets

Enter option 4 from the RAIDset menu to work with the sparesets and failedsets associated with RAIDsets. From the spareset/failedset menu (see Figure 7–7), you can, for example, enter option 1 to add a device to the spareset. CFMENU asks you which devices from the configured PTLs list that you wish to include in the spareset. In Figure 7–7, one device, PTL 420, was added to the spareset.

Figure 7–7 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 ^)       | 310 (dsk)|disks:  200 ^  DISK200 RZ35     (C) DEC                          N    N
2. Remove a device from the   | 350 (dsk)|        550 ^  DISK550 RZ25     (C) DEC                          N    N
   spareset                   | 400 (dsk)|        620 ^  DISK620 RZ25     (C) DEC                          N    N
3. Move a device from a       | 510 (dsk)|        640 ^  DISK640 RZ26     (C) DEC                          N    Y
   RAIDset to the failedset   | 600 (dsk)|cdrms:  150    CDROM150 RRD44   (C) DEC
   (eligible devices marked   |          |        240    CDROM240 RRD44   (C) DEC
   by *)                      |          |loadr:  540      LDR540 TL820   (C) DEC
4. Remove a device from the   |          |strps:  110    DISK110 RZ25     (C) DEC   S1      STRP    unk        N
   failedset                  |          |        220    DISK220 RZ25     (C) DEC   "       "       "          "
5. Return to RAIDset menu     |          |raid5:  330    DISK330 RZ26     (C) DEC   R1      RAID    unk        N    N
                              |          |        440    DISK440 RZ26     (C) DEC   "       "       "          "    "
                              |          |        530    DISK530 RZ26     (C) DEC   "       "       "          "    "
                              |          |spare:  420    DISK420 RZ25     (C) DEC
-------------------------------------------------------------------------------------------------------------------------

Enter menu choice (1,5) [5] ?
```

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

7.6.7 Adding Passthrough Containers

Enter option 5 from the main menu to work with passthrough containers. From the passthrough menu (see Figure 7–8), enter option 1 to add a passthrough container. CFMENU asks you which device from the configured PTLs list you wish to include in the passthrough container.

Although you may add any generic SCSI device, you normally add devices requiring the use of MSCP read and write operations as a transport for SCSI commands (such as "jukebox" style loaders). In Figure 7–8, a passthrough container was created from the tape library (loader) at PTL 540.

Figure 7–8 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 ^)                   | 310 (dsk)|disks:  200 ^ DISK200 RZ35     (C) DEC                          N   N
2. Delete an unbounded     | 350 (dsk)|        550 ^ DISK550 RZ25     (C) DEC                          N   N
   passthrough (marked     | 400 (dsk)|        620 ^ DISK620 RZ25     (C) DEC                          N   N
   by *)                   | 510 (dsk)|        640 ^ DISK640 RZ26     (C) DEC                          N   Y
3. Delete all unbounded    | 600 (dsk)|cdrms:  100   CDROM150 RRD44   (C) DEC
   passthroughs (marked    |          |        240   CDROM240 RRD44   (C) DEC
   by *)                   |          |strps:  110   DISK110 RZ25     (C) DEC ^  S1     STRP   unk       N
4. Return to main menu     |          |        220   DISK220 RZ25     (C) DEC    "      "      "         "
                           |          |raid5:  330   DISK330 RZ26     (C) DEC ^  R1     RAID   unk       N   N
                           |          |        440   DISK440 RZ26     (C) DEC    "      "      "         "   "
                           |          |        530   DISK530 RZ26     (C) DEC    "      "      "         "   "
                           |          |spare:  420   DISK420 RZ25     (C) DEC
                           |          |pass:   540   LDR540 TL820     (C) DEC *  P1     PASS

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

CFMENU updates the following fields after adding a passthrough:

- **Stor.set Name**—Storageset name. CFMENU automatically assigns the name "P*x*" to a passthrough. You cannot alter this automatic CFMENU naming convention.
- **Stor.set Type**—Storageset type. This field reads "PASS" for passthroughs.

After adding passthroughs, return to the main menu.

7.6.8 Initializing Containers

Enter option 6 from the main menu to initialize containers (devices or storagesets). From the initialization menu (see Figure 7–9), enter option 1. CFMENU asks you (y/n/q) if you want to initialize each eligible container.

In addition, CFMENU asks you to decide on other operating qualifiers, depending on whether the container is a device, stripeset, or RAIDset. See the descriptions of the ADD or SET commands in Appendix E if you need help understanding the qualifiers.

In Figure 7–9, the RAIDset created earlier (R1) is initialized. Notice that R1 now has a chunksize of 32. The disk at PTL 640 has been initialized throughout this example, because CFMENU detected its condition when it was added to the configuration.

Figure 7–9 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 ^)    | 310 (dsk)|disks:  200 ^ DISK200 RZ35     (C) DEC                          N    N
2. Return to main menu        | 350 (dsk)|        550 ^ DISK550 RZ25     (C) DEC                          N    N
                              | 400 (dsk)|        620 ^ DISK620 RZ25     (C) DEC                          N    N
                              | 510 (dsk)|        640 ^ DISK640 RZ26     (C) DEC                          N    Y
                              | 600 (dsk)|cdrms:  150   CDROM150 RRD44   (C) DEC
                              |          |        240   CDROM240 RRD44   (C) DEC
                              |          |strps:  110   DISK110 RZ25     (C) DEC ^  S1       STRP    unk          N
   WARNING! Initialization of |          |        220   DISK220 RZ25     (C) DEC    "        "       "            "
   any device or storageset   |          |raid5:  330   DISK330 RZ26     (C) DEC ^  R1       RAID    32           Y    N
   will destroy all of its    |          |        440   DISK440 RZ26     (C) DEC    "        "       "            "    "
   current data.              |          |        530   DISK530 RZ26     (C) DEC    "        "       "            "    "
                              |          |spare:  420   DISK420 RZ25     (C) DEC
                              |          |pass:   540   LDR540 TL820     (C) DEC *  P1       PASS

Enter menu choice (1,2) [2] ?
```

After initializing containers, return to the main menu.

7.6.9 Adding Units

Enter option 7 from the main menu to work with units. From the unit menu (see Figure 7–10), enter option 1 to add a unit. CFMENU asks you which initialized containers you wish to create units from.

CFMENU also asks 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 7–10.) In addition, CFMENU asks you to decide on other unit qualifiers. See the description of the ADD *unit* or SET *unit* commands in Appendix E if you need help understanding the qualifiers.

In Figure 7–10, two units were created from RAIDset R1 and the disk at PTL 640.

Figure 7–10 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  | 310 (dsk)|disks:  200   DISK200 RZ35     (C) DEC                        N    N
   units marked by *)       | 350 (dsk)|        550   DISK550 RZ25     (C) DEC                        N    N
3. Return to main menu      | 400 (dsk)|        620   DISK620 RZ25     (C) DEC                        N    N
                            | 510 (dsk)|        640   DISK640 RZ26     (C) DEC                        N    Y      * D464 N Y
                            | 600 (dsk)|cdrms:  150 ^ CDROM150 RRD44   (C) DEC
                            |          |        240 ^ CDROM240 RRD44   (C) DEC
                            |          |strps:  110   DISK110 RZ25     (C) DEC  S1       STRP   unk      N
                            |          |        220   DISK220 RZ25     (C) DEC  "        "      "        "
                            |          |raid5:  330   DISK330 RZ26     (C) DEC  R1       RAID   32       Y    N * D465 N Y
                            |          |        440   DISK440 RZ26     (C) DEC  "        "      "        "    "   "    " "
                            |          |        530   DISK530 RZ26     (C) DEC  "        "      "        "    "   "    " "
                            |          |spare:  420   DISK420 RZ25     (C) DEC
                            |          |pass:   540   LDR540 TL820     (C) DEC *  P1     PASS

Enter menu choice (1,3) [3] ?
```

CFMENU updates the following fields after adding a unit:

- **Reduc**—Reduced. This field indicates whether a RAIDset is running reduced (missing one member).
- **Unit**—The unit number you assign appears here, preceded by the letter "D" or "T."
- **WP**—Write protect. This value indicates whether the unit is write protected.
- **WB**—Write-back. This value indicates whether the unit is set for write-back caching.

7.6.10 Terminal Setup

You can enter option 8 from the main menu to set the number of rows CFMENU displays. This feature is available primarily for terminals with the capability of displaying more than 24 rows.

7.6.11 Messages

This section lists the messages, other than the standard CLI messages, that CFMENU displays. However, most messages you see are those sent by the CLI, and are described in Appendix E.

CFMENU cannot complete request without exceeding array boundary.

Explanation: CFMENU detected an unexpected condition which would exceed an array boundary and possibly require controller reinitialization, so it aborted your request.

...CFMENU is updating its configuration tables. Please be patient...

Explanation: CFMENU is polling the SCSI ports to see what physical devices are in place, as well as checking the configuration information.

Not enough eligible devices to complete the storageset.

Explanation: You chose to create a stripeset or a RAIDset and specified how many members to use, but there are not enough eligible devices to make up a storageset of this size. Eligible devices are disks that have the NOTRANSPORTABLE switch set, and that are not already used in any higher-level configuration such as a unit, storageset, spareset, or failedset.

Not enough members specified for a non-reduced RAIDset

Explanation: You chose to create a RAIDset and specified that it is not a previously reduced RAIDset. However, when choosing how many members to add to the RAIDset, you specified a number that is legal only for a reduced RAIDset and is too low for a nonreduced RAIDset.

...Polling for unconfigured devices...

Explanation: CFMENU is polling the SCSI ports to see what physical devices are in place.

Port *port#* is blocked. No devices are configured on port *port#*

Explanation: In order to check each device bus to discover what devices are present, firmware must also see if any port is currently blocked. A port can be blocked for various reasons, such as when its bus is quiesced. When a port is blocked, CFMENU does not access devices on that bus.

RAID5 license is not enabled; cannot create RAID5 set.

Explanation: You may not configure any RAIDsets unless the RAID5 license has been enabled by running the FLS utility.

Received user request to terminate CFMENU...

Explanation: You pressed Ctrl/C or Ctrl/Y to abort CFMENU.

There are no devices available to use as a replacement.

Explanation: You chose to replace a member of a reduced RAIDset, but there are no disks eligible to use for the replacement member. Eligible disks must have the NOTRANSPORTABLE switch set and may not be part of any higher-level configuration such as units, storagesets, or the spareset or failedset.

There are no devices available to use in a passthrough

Explanation: You chose to create a passthrough container, but there are no devices eligible to be added to a passthrough. Eligible devices are any SCSI device that is not already used in any higher-level configuration such as a unit, storageset, spareset, or failedset. Disk devices must be set NOTRANSPORTABLE to be eligible.

There are no *devices/stripesets/RAIDsets/passthroughs* eligible for deletion.

Explanation: You chose to delete a device, stripeset, RAIDset, or passthrough; but there are none that are eligible for deletion. A stripeset, RAIDset, or passthrough may not be deleted if it is configured as a unit. A device may not be deleted if it is configured as a unit or if it is used in a storageset, spareset, or failedset.

There are no devices eligible to be added to the configuration.

Explanation: You chose to add a device, but there are no devices available to add. The only devices that are eligible to be added are devices that CFMENU has detected as being physically present on a SCSI port and that are not already configured as devices on the controller.

There are no devices eligible to be added to the spareset.

Explanation: You chose to add a device to the spareset, but no devices are eligible. The only devices that can be added to the spareset are disks that have the NOTRANSPORTABLE switch set, and that are not already used in any higher-level configuration such as a unit, storageset, spareset, or failedset.

There are no devices in the failedset.

Explanation: You chose to delete devices from the failedset but the failedset currently is empty.

There are no devices in the spareset.

Explanation: You chose to delete devices from the spareset but the spareset currently is empty.

There are no entities eligible for initialization.

Explanation: You chose to initialize a device or storageset, but there are currently no devices or storagesets that are eligible to be initialized. Tapes, CDROMs, loaders, and passthrough devices may not be initialized. Optical memory devices, stripesets, and RAIDsets can be initialized, but only if they are not already configured as a unit. Disks can be initialized only if they are not already configured as a unit or as part of a storageset, spareset, or failedset.

There are no entities eligible to be added as units.

Explanation: You chose to add a unit, but there are no devices or storagesets that are eligible to become units. Tape and CDROM devices can be added as units unless they are already configured as units. Passthrough containers also can be added as units unless they are already configured as units. Disks, optical memory, stripesets, and RAIDsets must first be initialized before they can be added as units. Disks may not be made into units if they are currently in the spareset or the failedset.

There are no reduced RAIDset units with NOPOLICY set.

Explanation: You chose to replace a member of a reduced RAIDset. CFMENU is unable to find any RAIDsets that are eligible for a manual replacement. In order to be eligible, the RAIDset must be configured as a unit, it must be in a reduced state, and it must have the NOPOLICY switch set.

There are no units to delete.

Explanation: You chose to delete a unit, but there are no units configured on the controller.

Unable to allocate memory, CFMENU terminating.

Explanation: There is not enough memory available for CFMENU to run.

User has not picked enough eligible devices to complete the storageset.

Explanation: You chose to create a stripeset or a RAIDset and specified how many members to use, but when CFMENU prompted for devices, you did not select enough to complete the storageset.

Waiting for completion of CLI command...

Explanation: Some CLI commands take a long time to complete, such as initializing a large RAIDset or adding a tape drive that needs to have the tape rewound. CFMENU prints out this message periodically to inform you that it is still waiting for the last CLI command to finish.

7.6.12 Exiting CFMENU

Enter the last option from the main menu to stop CFMENU and return to the CLI. (You also may enter Ctrl/C or Ctrl/Y to abort CFMENU.)

7.7 Code Patch Utility

The Code Patch utility allows you to alter programs stored in the controller's NVMEM while the controller is active.

Note

The Code Load portion of the Code Load/Code Patch (CLCP) utility is not supported on the HS1CP.

7.7.1 Invoking the Code Patch Utility

Invoke the Code Patch utility using the RUN command from the CLI interface as follows:

```
CLI> RUN CLCP
 Select an option from the following list:

 Code Load & Code Patch local program Main Menu

 0: Exit
 1: Enter Code LOAD local program
 2: Enter Code PATCH local program

   Enter option number (0..2) [0] ?
```

Select option 2. The Code Patch utility executes, prompting for further inputs.

7.7.2 Code Patching

The Code Patch function of the CLCP utility allows you to enter small changes (provided by Digital Equipment Corp.) to the executing controller code. The changes that you enter are placed directly into the controller's NVMEM, and become active after the next controller initialization.

The Code Patch function also allows you to list patches already installed and delete any of those patches.

You can run the Code Patch function from either a maintenance terminal or a virtual host terminal.

Note

The Code Patch function of the CLCP utility is supported on all HS array controller models.

7.7.2.1 Code Patch Considerations

Be aware of the following when using the Code Patch function of the CLCP utility:

- The controller reserves enough nonvolatile memory for approximately 10 patches. However, this number varies according to the size of the patches you enter.
- Each patch is associated with only one firmware version, and the Code Patch program verifies the patch against the currently installed firmware version.
- Patches are hierarchical. In other words, patch number 1 must be entered before you enter patch number 2, and so on. Furthermore, there are no "0" patches. Patches are always numbered sequentially beginning with "1."

- Because of the hierarchical patch structure, removing any patch also removes all higher numbered patches. For example, deleting patch number 2 also removes patches 3, 4, and so on.
- Controllers in dual-redundant configurations must have the same patches applied, and patches must be entered twice (once in each controller).

7.7.2.2 Entering a Patch

The following is an example of entering a patch to the current firmware version loaded on your controller:

```
CLI> RUN CLCP

 Select an option from the following list:

Code Load & Code Patch local program Main Menu

 0: Exit
 1: Enter Code LOAD local program
 2: Enter Code PATCH local program

Enter option number (0..3) [0] ? 2

 You have selected the Code Patch local program.  This program is used to
 manage firmware code patches.  Select an option from the following list:

 Type ^Y or ^C (then RETURN) at any time to abort the Code Patch.

Code Patch Main Menu

 0: Exit
 1: Enter a Patch
 2: Delete Patches
 3: List Patches

Enter option number (0..3) [0] ? 1

 This is the Enter Patch option.  The program prompts you
 for the patch information, one line at time.  Be careful
 to enter the information exactly as it appears on the patch release.
 Patches may be installed for any version of firmware; however,
 patches entered for firmware versions other than XLZ6 are not
 applied until the matching version of firmware is installed.

 To enter any patch, you must first install all patches with lower patch
 numbers, beginning with patch number 1, for the specific firmware version.
 If you incorrectly enter the patch information, you are given the option to
 review the patch one line at a time.

 Type ^Y or ^C (then RETURN) at any time to abort Code Patch.

 Do you wish to continue (y/n) [y] ?

 Version:       ? XLZ6
 Length:        ? 10
 Patch Type:    ? 0
 Patch Number:  ? 1

 Count:         ? 1
 Address:       ? 10
 Value[  0]     ? 0

 Count:         ? 0

 Verification:  ? B2103761

 The patch you have just entered is not applied until
 the controller is restarted.
```

7.7.2.3 Listing Patches

The List Patches option allows you to display a listing of controller firmware versions, and the currently installed patches that apply to them.

The following is an example of the List Patches option:

```
CLI> RUN CLCP

 Select an option from the following list:

 Code Load & Code Patch local program Main Menu

 0: Exit
 1: Enter Code LOAD local program
 2: Enter Code PATCH local program

  Enter option number (0..2) [0] ? 2

 You have selected the Code Patch local program.  This program is used to
 manage firmware code patches.  Select an option from the following list:

 Type ^Y or ^C (then RETURN) at any time to abort Code Patch.

Code Patch Main Menu

 0: Exit
 1: Enter a Patch
 2: Delete Patches
 3: List Patches

  Enter option number (0..3) [0] ? 3

The following patches are currently stored in the patch area:

   Firmware Version - Patch number(s)

        123456  -      1
        XLZ6    -      2,1

 Currently, 91% of the patch area is free.
```

Note

The SHOW THIS_CONTROLLER command also provides patch information in the form of a "dash number" following the firmware version. In the following example, firmware Version 2.5 has had patches applied to patch number 3:

```
CLI> SHOW THIS_CONTROLLER

Controller:
        HSZ40 ZG33400026 Firmware V25J-3, Hardware 0000
        .
        .
        .
```

7.7.2.4 Deleting a Patch

The Delete Patches option allows you to remove installed patches from controller non-volatile memory. The program displays the currently installed patches so that you can identify the patches to be deleted.

The following is an example of how to delete a patch and a firmware version that are no longer needed:

```
CLI> RUN CLCP

 Select an option from the following list:

 Code Load & Code Patch local program Main Menu

 0: Exit
 1: Enter Code LOAD local program
 2: Enter Code PATCH local program

  Enter option number (0..2) [0] ? 2

 You have selected the Code Patch local program.  This program is used to
 manage firmware code patches.  Select an option from the following list:

 Type ^Y or ^C (then RETURN) at any time to abort Code Patch.

Code Patch Main Menu

 0: Exit
 1: Enter a Patch
 2: Delete Patches
 3: List Patches

  Enter option number (0..3) [0] ? 2

 This is the Delete Patch option.  The program prompts you
 for the firmware version and patch number you wish to delete.
 If you select a patch for deletion that is required for another patch,
 all dependent patches are also selected for deletion.  The
 program lists your deletion selections and asks if you wish to continue.

 Type ^Y or ^C (then RETURN) at any time to abort Code Patch.

 The following patches are currently stored in the patch area:

   Firmware Version - Patch number(s)

       123456  -     1
       XLZ6    -     2,1

 Currently, 91% of the patch area is free.

 Firmware Version of patch to delete ? 123456
 Patch Number to delete ? 1

 The following patches have been selected for deletion:

  Firmware Version - Patch #

       123456  -     1

 Do you wish to continue (y/n) [y] ?

 The patch you have just deleted is  currently applied, but it will not be
 applied when the controller is  restarted.

 The following patches  are currently stored in the patch area:

  Firmware Version - Patch #

       XLZ6    -     2,1

 Currently, 94% of the patch area is free.
```

7.7.2.5 Messages

The Code Patch utility does not allow you to incorrectly enter or delete patch information. The program provides messages to assist you with understanding any problems that you may encounter and suggests corrective actions.

Firmware Version *x* does not have any patches to delete.

Explanation: You cannot delete a patch because the firmware version entered does not have any patches entered.

Firmware Version *x* does not have patch number *x* to delete.

Explanation: You cannot delete this patch because the firmware version entered does not have the specified patch entered.

The patch you entered is already installed on this controller.

Explanation: The specified patch is already present in the patch area of controller memory. If you wish to reenter this patch, first use the Delete Patch option.

The patch you are entering requires other patches to be entered.

Explanation: You have attempted to enter a patch without first entering the lower numbered patches in the hierarchy. Enter all patches for this firmware version that have lower numbers than the current patch. Then enter the current patch.

WARNING The patch you are entering is not for the current firmware version *x*.

Explanation: The patch you are entering applies to a firmware version other than the one currently installed in the controller. Code Patch allows you to enter the patch; however, the patch is not applied until its correct firmware version is installed.

You incorrectly entered the patch information.

Explanation: The patch information was not entered exactly. The program prompts you for each line of the patch entry, with the default from your previous response. Verify that each entry is exactly the same as the patch release. If you choose not to continue, or if you abort during this review procedure, the patch information you entered is lost and you must enter the entire patch again. You may enter Ctrl/Z followed by Return at any prompt to choose the default for the remaining entries.

The patch you have just entered is not applied until the controller firmware is changed to Version *x*.

Explanation: The patch entered applies to a firmware version other than the one currently installed in the controller. Code Patch does not apply the patch until its correct firmware version is installed.

You have requested deletion of a patch number that another patch requires.

Explanation: You are attempting to delete a patch in the hierarchy that has higher numbered patches entered. Code Patch allows you to proceed; however, the program deletes all the higher numbered patches in the hierarchy (for this firmware version) along with the specified patch.

7.7.2.6 Exiting Code Patch

Exit Code Patch by choosing option 0 from the main menu. (Pressing Ctrl/C or Ctrl/Y at any time during Code Patch also aborts and exits the program.)

7.8 Firmware Licensing System

The firmware licensing system (FLS) enables or disables the licensed value-added software features of the HS1CP controller. You can use the FLS utility to perform the following tasks:

- Enable or disable optional functions for your controller
- Try an optional feature before purchasing the license to use it
- Change your license key for an option

Start FLS from the CLI prompt. After starting, the FLS display shows the current status of the value-added options for your controller and contains menu choices for each function of the utility.

7.8.1 Enabling Options

You can turn on any option at any time with FLS, but if you enable an option for which you are not licensed, an error message appears on your CLI terminal and an error is logged in the host error log. These error indications are repeated at least once each hour while the unlicensed option remains enabled.

7.8.2 Disabling Options

You cannot disable an option if that option is currently in use. The following table lists the conditions under which you can disable an FLS option.

Option	Conditions Required to Disable
RAID	No RAIDset configured Write-back caching not in use
WBCA	Write-back caching not in use

7.8.3 License Key

When you first run FLS, the license key is cleared. If you purchase a license for a firmware option, you receive a customer license key. This key contains two parts: a customer identification string from 6 to 32 characters long, and an 8-character cyclic redundancy check (CRC) string. You must enter the customer identification string with the CRC string appended to it when you use FLS.

7.8.4 Using the Menu

You can perform the following operations from the FLS menu:

Select	Action	Submenu Choices	Result
1	Enable an option	List each option and its status	Selection enabled
2	Disable an option	List each option and its status	Selection disabled
3	Enter a license key	Prompt for new license key	Entered key checked for validity
4	Clear a license key	Prompt for license key to clear	Entered key becomes invalid

The following example shows the FLS main menu:

```
CLI> RUN FLS

------------------------------------------------------------------------------
     Firmware Licensing System (FLS) on node MASS
Option❶   State❷       License❸       Key❹
------------------------------------------------------------------------------
 RAID       ENABLED ******INVALID!****** *none*
 WBCA       ENABLED        VALID         ACME_WIDGET_CORP........

              RAID = Raid5 + Writeback Cache ❺
              WBCA = Writeback Cache ONLY
------------------------------------------------------------------------------
    1. Enable a firmware option
    2. Disable a firmware option
    3. Enter a license key for a firmware option
    4. Clear a license key for a firmware option
    0. Exit FLS
Enter selection (0:4) [0] ?
```

❶ **Option**–The RAID and write-back cache (WBCA) options are available.

❷ **State**–Both RAID and WBCA are enabled. You can use any option that is enabled, regardless of whether you have a valid license key.

❸ **License**–RAID is running without a valid license. This status shows when you are running an option on a trial basis. The license becomes valid when you enter a license key that FLS verifies as valid. You receive this key when you purchase a software option.

❹ **Key+CRC**–The license key is ACME_WIDGET_CORP; the 8-character CRC portion of the key is shown as hidden text (.....).

❺ **Description of Option**–A short description of each option is given.

7.8.5 Example

To perform an operation, enter the choice number and any information requested by the Submenu or prompts. The following example demonstrates how to enter a license key and enable write-back caching.

```
CLI> RUN FLS

------------------------------------------------------------------------------
     Firmware Licensing System (FLS) on node MASS
Option      State        License        Key
------------------------------------------------------------------------------
 RAID       DISABLED     INVALID        *none*
 WBCA       DISABLED     INVALID        *none*

              RAID = Raid5 + Writeback Cache
              WBCA = Writeback Cache ONLY
------------------------------------------------------------------------------
    1. Enable a firmware option
    2. Disable a firmware option
    3. Enter a license key for a firmware option
    4. Clear a license key for a firmware option
    0. Exit FLS
Enter selection (0:4) [0] ? 3

    1. Enter new license key+CRC for RAID   (current key is invalid)
    2. Enter new license key+CRC for WBCA   (current key is invalid)
    0. Return to main menu
Enter selection (0:2) [0] ? 2❶

Enter new WBCA key, including 8-character CRC, or enter 0
       to return to main menu: ACME_WIDGET_CORPVB8UWQ9C❷
```

```
*** License key verified ***

-------------------------------------------------------------------------------
     Firmware Licensing System (FLS) on node MASS
Option      State        License        Key
-------------------------------------------------------------------------------
 RAID      DISABLED      INVALID        *none*
 WBCA      DISABLED       VALID         ACME_WIDGET_CORP........

              RAID = Raid5 + Writeback Cache
              WBCA = Writeback Cache ONLY
-------------------------------------------------------------------------------
    1. Enable a firmware option
    2. Disable a firmware option
    3. Enter a license key for a firmware option
    4. Clear a license key for a firmware option
    0. Exit FLS
Enter selection (0:4) [0] ? 1

    1. Enable RAID
    2. Enable WBCA
    0. Return to main menu
Enter selection (0:2) [0] ? 2●
*** WBCA enabled ***

-------------------------------------------------------------------------------
     Firmware Licensing System (FLS) on node MASS
Option      State        License        Key
-------------------------------------------------------------------------------
 RAID      DISABLED      INVALID        *none*
 WBCA      ENABLED        VALID●        ACME_WIDGET_CORP........

              RAID = Raid5 + Writeback Cache
              WBCA = Writeback Cache ONLY
-------------------------------------------------------------------------------
    1. Enable a firmware option
    2. Disable a firmware option
    3. Enter a license key for a firmware option
    4. Clear a license key for a firmware option
    0. Exit FLS
Enter selection (0:4) [0] ? Return
FLS - Normal Termination

CLI>
```

● **The user chooses to enter a new license key for WBCA.**

● **The user enters the new license key, along with the customer license key, which is displayed as it is entered.**

● **The user enables write-back cache.**

● **This entry in the FLS display shows that write-back cache is enabled under a valid license.**

7.8.6 Messages

This section lists the message that you can receive from FLS.

option has been turned on without a valid license

Explanation: You have activated the option named by *option* without entering a valid license key. You can evaluate this option for a time to determine its value, and you receive a valid license key when you purchase the license for the option.

Error *nnnn*: *option* support is not enabled on this controller

Explanation: The *option* you are attempting to use is not enabled in FLS. For example, if you try to turn on write-back caching, you receive this error if either RAID or write-back caching is not enabled by FLS (the RAID option includes write-back caching).

WARNING: This is an invalid license Key+CRC

Explanation: The license key you entered is not between 6 and 32 characters, or the customer license key is not valid. Verify that the key is correct and reenter.

***Error: Disabling *option* is not possible at this time, option is in use ***

Explanation: You have attempted to disable the option named by *option* while it is in use. Refer to Section 7.8.2 for more information on disabling options.

7.9 Fault Management Utility

The Fault Management Utility (FMU) allows you to do the following:

- Control the spontaneous event logging and last failure logging displays
- Display controller last failure and memory system failure information

FMU also can provide a convenient way to review some error log information during your terminal session. Run FMU interactively during a terminal session by invoking it from the CLI prompt:

```
CLI> RUN FMU
```

FMU interprets only errors that occur after you install and run controller firmware containing FMU. In other words, FMU cannot search for "older" errors, and you cannot install FMU in an attempt to troubleshoot a preexisting error.

After invoking FMU, you can perform the functions described in the following sections. Defaults are specified by "D."

7.9.1 The SET Command

The SET command controls the output from the FMU utility. The SET command changes options for the following two spontaneous displays:

- Event logging (EVL)
- Last failure logging (LFL)

Also, the SET command controls the options for the interactive displays available under the SHOW command.

SET EVENT_LOGGING
SET NOEVENT_LOGGING (D)

This command enables/disables the event log display on the maintenance terminal. With the event log display enabled, the controller spontaneously displays EIP information during your terminal session. The first line of an event log display begins with "%EVL."

Event log displays are inhibited during the execution of both CLI commands and utilities invoked from a maintenance terminal. Events that are reported while a maintenance terminal is in use do not appear when the terminal again becomes available. (The %EVL display is lost.)

Note

Execution of a CLI command or utility does not begin until you press Return. If FMU reports an event during command line input before Return is pressed, the %EVL display interrupts the input.

Following the %EVL display, the CLI prompt and command input entered prior to the interruption redisplays. You then can complete the current command line (unless FMU reports another event).

SET LAST_FAILURE_LOGGING
SET NOLAST_FAILURE_LOGGING (D)

This command enables/disables the last failure log display on the maintenance terminal. With the last failure log display enabled, the controller spontaneously displays information relevant to the sudden termination of executing firmware. The first line of a last failure log display begins with "%LFL."

In cases where an automatic hardware reset occurs (such as power failure, pressing the Reset button, and so on) the last failure log display is inhibited because automatic reset does not allow sufficient time to complete the display.

SET EVENT_LOGGING REPAIR_ACTION_DISPLAY
SET EVENT_LOGGING NOREPAIR_ACTION_DISPLAY (D)
SET LAST_FAILURE_LOGGING REPAIR_ACTION_DISPLAY
SET LAST_FAILURE_LOGGING NOREPAIR_ACTION_DISPLAY (D)

This command and qualifier enables/disables recommended repair action display for event logging and last failure logging displays. With recommended repair action display enabled, the controller displays all of the recommended repair actions associated with the Instance Code and/or Last Failure Code used to describe an event.

SET EVENT_LOGGING VERBOSE
SET EVENT_LOGGING NOVERBOSE (D)
SET LAST_FAILURE_LOGGING VERBOSE
SET LAST_FAILURE_LOGGING NOVERBOSE (D)

This command and qualifier enables/disables descriptive text for event logging and last failure logging displays.

The display always identifies the various fields and their numeric content that comprise an event/last failure log. With verbosity enabled, the controller also displays a description of the numeric value in each log field.

SET PROMPT_DISPLAY
SET NOPROMPT_DISPLAY (D)

This command enables/disables the CLI prompt string display within the first line of event logging and last failure logging displays, as shown in the following example (using %EVL):

```
%EVL-- Instance Code: 01010302
%EVL--CLI> -- Instance Code: 01010302
```

SET TIMESTAMP_DISPLAY
SET NOTIMESTAMP_DISPLAY (D)

This command enables/disables current timestamp string display within the first line of event logging and last failure logging displays, as shown in the following example (using %EVL):

```
%EVL-- Instance Code: 01010302
%EVL--07-JUL-1994 07:44:48-- Instance Code: 01010302
```

You can use combinations of the SET [NO]PROMPT_DISPLAY and SET [NO]TIMESTAMP_DISPLAY commands to provide the following types of event logging and last failure logging first line displays (examples using %EVL):

```
%EVL--CLI> --07-JUL-1994 07:44:48-- Instance Code: 01010302

%EVL--07-JUL-1994 07:44:48-- Instance Code: 01010302

%EVL--CLI> Instance Code: 01010302

%EVL-- Instance Code: 01010302
```

The last example shown is recommended when VCS is in use, since the controller identification (prompt string) and timestamp information is already supplied by VCS.

SET FMU REPAIR_ACTION_DISPLAY
SET FMU NOREPAIR_ACTION_DISPLAY (D)

This command enables/disables the recommended repair action display for FMU SHOW LAST_FAILURE and SHOW MEMORY_SYSTEM_FAILURE command output. With recommended repair action display enabled, the command output displays all of the recommended repair actions associated with the Instance Code and/or Last Failure Code used to describe an event.

SET FMU VERBOSE
SET FMU NOVERBOSE (D)

This command enables/disables descriptive text for FMU SHOW LAST_FAILURE and SHOW MEMORY_SYSTEM_FAILURE command output. The output always identifies the various fields and their numeric content that comprise an event/last failure log. With verbosity enabled, the controller also displays a description of the numeric value in each log field.

SET EVENT_LOGGING [qualifier ... qualifier] PERMANENT
SET NOEVENT_LOGGING PERMANENT
SET LAST_FAILURE_LOGGING [qualifier ... qualifier] PERMANENT
SET NOLAST_FAILURE_LOGGING PERMANENT
SET FMU [qualifier ... qualifier] PERMANENT

The PERMANENT qualifier stores the parameter setting specified by the primary keyword and optional qualifiers in nonvolatile memory so that the setting is preserved across controller resets. In addition, when PERMANENT is specified, the given setting takes effect immediately.

If the PERMANENT qualifier is not specified, the given setting takes effect immediately. However, it remains in effect only as long as the current FMU session remains active or until the setting is changed by a subsequent SET command.

When running FMU from a maintenance terminal, changing EVENT_LOGGING parameters without specifying the PERMANENT qualifier has no effect. However, the same action while running FMU from a virtual terminal is effective. This permits the EVENT_LOGGING operation to be changed on a temporary basis only from a virtual terminal.

You can specify multiple additional qualifiers on the same command line for the SET EVENT_LOGGING, SET LAST_FAILURE_LOGGING, and SET FMU commands. For example, the following are all valid commands:

```
FMU> SET EVENT_LOGGING PERMANENT
FMU> SET LAST_FAILURE_LOGGING NOREPAIR_ACTION_DISPLAY PERMANENT
FMU> SET EVENT_LOGGING REPAIR_ACTION_DISPLAY NOVERBOSE PERMANENT
```

7.9.2 The SHOW Command

The SHOW command controls the interactive reviewing of last failure and memory system failure information.

SHOW LAST_FAILURE qualifier [additional qualifier]

This command interactively displays the last failure information stored in nonvolatile memory. Information related to the most recent and three previous last failure events is stored in nonvolatile memory.

Qualifier	Description
MOST_RECENT	Displays the most recent last failure information.
ALL	Displays all the errors available in the buffer (up to four). Errors are displayed in descending order, starting with the most recent.
ENTRY *n*	Displays one of the entries from the buffer. You must supply an entry number (range 1–4).

Additional Qualifier	Description
FULL	When included on the same command line with the MOST_RECENT, ALL, or ENTRY qualifier, displays extended information valuable to Digital Multivendor Customer Services.

SHOW MEMORY_SYSTEM_FAILURE qualifier

This command interactively displays memory system failure information from any of the last failure entries stored in nonvolatile memory.

Qualifier	Description
MOST_RECENT	Displays the most recent memory system failure information contained in any of last failure information entries. (Note that the most recent memory system failure may not be the most recent last failure.)
ALL	Displays all memory system failure information contained in any or all of the four last failure information entries, in most recent to least recent order.
ENTRY *n*	Displays memory system failure information contained in one last failure information entry (range 1–4). Note that the FULL additional qualifier is not available with the SHOW MEMORY_SYSTEM_FAILURE command. To obtain the extended information associated with the selected memory system failure, perform a SHOW LAST_FAILURE ENTRY *n* FULL, where *n* is the last failure entry number identified in the memory system failure display.

Note

The following message appears when FMU cannot access error information for the SHOW command:

```
(***Last Failure Entry x EDC bad; translation terminated***)
```

SHOW PARAMETERS

This command displays the current/permanent setting of parameters affected by the SET command.

SHOW TIME

This command displays the current controller time (that is, wall clock time) and the current controller power on time in the following format:

```
Time: 07-AUG-1995 10:25:35
Power On Time: 0.Years, 0. Days, 8. Hours, 27. Minutes, 8. Seconds
```

If the controller time has not been set yet when the show times command is issued, the format looks like this:

```
Time: 07-AUG-1995 10:25:35 (time not set)
Power On Time: 0.Years, 0. Days, 8. Hours, 27. Minutes, 8. Seconds
```

7.9.3 The DESCRIBE Command

The DESCRIBE command displays descriptive text for a numeric value contained in a particular event log field.

DESCRIBE ASC_ASCQ_CODE[3]
DESCRIBE COMPONENT_CODE
DESCRIBE CONTROLLER_UNIQUE_ASC_ASCQ_CODE[3]
DESCRIBE DEVICE_TYPE_CODE
DESCRIBE EVENT_THRESHOLD_CODE

[3] Requires entry of multiple numeric values.

DESCRIBE INSTANCE_CODE
DESCRIBE LAST_FAILURE_CODE
DESCRIBE REPAIR_ACTION_CODE
DESCRIBE RESTART_TYPE
DESCRIBE SCSI_COMMAND_OPERATION_CODE[3]
DESCRIBE SENSE_DATA_QUALIFIERS[3]
DESCRIBE SENSE_KEY_CODE
DESCRIBE TEMPLATE_CODE

DESCRIBE CI_OPERATION_CODE
DESCRIBE CONNECTION_STATE_CODE
DESCRIBE HIS_STATUS_CODE
DESCRIBE PPD_OPERATION_CODE
DESCRIBE SCS_OPERATION_CODE
DESCRIBE VC_STATE_CODE

All DESCRIBE qualifiers require at least one numeric value parameter. DESCRIBE qualifiers requiring multiple numeric value parameters are footnoted as such.

Type a question mark (?) in place of a numeric value parameter to identify the value and range required, as shown in the following example. Note that when sequential values are required, you must supply values for the earlier parameters before entering a question mark for the later parameter in the sequence.

```
FMU> DESCRIBE ASC_ASCQ_CODE ?
Your options are:
   ASC value (range: 0 through FF hexadecimal)
FMU> DESCRIBE ASC_ASCQ_CODE 0 ?
Your options are:
   ASCQ value (range: 0 through FF hexadecimal)
FMU> DESCRIBE ASC_ASCQ_CODE 0 0 ?
Your options are:
   SCSI Device Type value (range: 0 through FF hexadecimal)
FMU>
```

7.9.4 The EXIT Command

The EXIT command terminates FMU and returns you to the CLI prompt. (You also can enter Ctrl/C or Ctrl/Y to abort FMU.)

7.9.5 Examples

This section presents examples that show some of the output information available when using FMU. Values enclosed in parentheses are hexadecimal translations of decimal numbers.

```
FMU> SHOW LAST_FAILURE ENTRY 4

Last Failure Entry: 4❶ Flags: 0007FA80❷
```

```
%FMU-01-Last Failure Event, Instance Code: 01010302❶
 Power On Time: 0 Years, 41 Days, 4 Hours, 49 Minutes, 8 Seconds❷
 ❸Controller Model: HS1CP Serial Number: ZG30355555 Hardware Version: 0000(00)
 Controller Identifier:
  Unique Device Number: 000130355555 Model: 40(28) Class: 1(01)
 HSOF Version: V20(20)❹
 Node Name: "HS1CP1" CI Node Number: 7(07)❺
 Informational Report
 ❻Instance Code 01010302 Description:
  An unrecoverable hardware detected fault occurred.
 ❼Last Failure Code: 018800A0 (No Last Failure Parameters)
 Last Failure Code 018800A0 Description:
  A processor interrupt was generated with an indication that the program
  card was removed.
```

```
FMU> SHOW MEMORY_SYSTEM_FAILURE ENTRY 2

Last Failure Entry: 2❶ Flags: 0007FA8C❷
 Memory System Failure indicated

%FMU-14-Memory System Failure Event, Instance Code: 016E2D02❸
 Power On Time: 0 Years, 41 Days, 21 Hours, 5 Minutes, 39 Seconds❹
 ❺Controller Model: HSICP Serial Number: ZG30355555 Hardware Version: 0000(00)
 Controller Identifier:
  Unique Device Number: 000130355555 Model: 40(28) Class: 1(01)
 HSOF Version: V20(20)❻
 Node Name: "HSICP1" CI Node Number: 7(07)❼
 Reported via non-maskable interrupt
 ❽Memory Address: 40000000
 Byte Count: 0(00000000)
 DRAB Registers:
  DSR:  2D17403F  CSR:  8000A220 DCSR:  00003403  DER:  00001C00  EAR:  04000000
  EDR.  F4000003  ERR:  00000000  RSR:  09805432  CHC:  E7FFFFFC  CMC:  90A5FEF0
 Diagnostic Registers:
  RDR0: E7FFFFFC  RDR1: 90A5FEF0  WDR0: 7F021000  WDR1: FF06020D
 ❸Instance Code 016E2D02 Description:
  The CACHEA0 DRAB detected a Nonexistent Memory Error condition during an
  I960 attempt to read CACHEA0 memory.
```

❶ **Entry**—A number representing the error position in the buffer.

❷ **Flags**—This value should be recorded and reported to Digital Multivendor Customer Services.

❸ **Instance Code**—The instance code (and description) associated with this failure.

❹ **Power On Time**—The time of failure.

❺ **Controller**—Information identifying your controller.

❻ **HSOF Version**—Firmware version.

❼ **Node Name**—Node identifier.

❽ **Last Failure Code**—Last failure code and description.

❾ **Memory Address**—The memory address, byte count, and register contents should be recorded and reported to Digital Multivendor Customer Services.

```
FMU> SHOW LAST_FAILURE ENTRY 4 FULL

Last Failure Entry: 4 Flags: 0007FA80

%FMU-01-Last Failure Event, Instance Code: 01010302
 Power On Time: 0 Years, 41 Days, 4 Hours, 49 Minutes, 8 Seconds
 Controller Model: HS1CP Serial Number: ZG30355555 Hardware Version: 0000(00)
 Controller Identifier:
  Unique Device Number: 000130355555 Model: 40(28) Class: 1(01)
 HSOF Version: V20(20)
 Node Name: "HS1CP1" CI Node Number: 7(07)
 Informational Report
 Instance Code 01010302 Description:
  An unrecoverable hardware detected fault occurred.
 Last Failure Code: 018800A0 (No Last Failure Parameters)
 Last Failure Code 018800A0 Description:
  A processor interrupt was generated with an indication that the program
  card was removed.
 Current Thread: NULL Current I960 Priority: 001F0000
 Interrupt Stack Guard is intact
 Thread Stack Guard State Flags (ID# Bit; 0=intact,1=not intact): 00000000
 I960 Stack:
  Levels: 2
  Level 0:
   Return type: Interrupt
   PFP: 201C15F7  SP:  201FABD0  RIP: 200CF898  R3:  201C15F7
   R4:  00000000  R5:  00000000  R6:  00000000  R7:  00000000
   R8:  00000000  R9:  00000000  R10: 00000000  R11: 00000000
   R12: 00000000  R13: 00000000  R14: 00000000  R15: D87FA8FE
  Level 1:
   Return type: Local
   PFP: 00000000  SP:  201C1680  RIP: 200D0AC0  R3:  00000000
   R4:  00000000  R5:  00000000  R6:  00000000  R7:  00000000
   R8:  00000000  R9:  00000000  R10: 00000000  R11: 00000000
   R12: 00000000  R13: 00000000  R14: 00000000  R15: 00000000
   G0:  00000000  G1:  00000000  G2:  2011DFF4  G3:  00000000
   G4:  00000000  G5:  00400000  G6:  201148B0  G7:  00000003
   G8:  000000CC  G9:  00000001  G10: 00000004  G11: 2011F108
   G12: 200E969C  G13: 201D3D8C  G14: 00000000  FP:  201FAB50
 Diagnostic Registers:
  RDR0: E7FFFFFC  RDR1: E7FFFEF0  WDR0: 7F021000  WDR1: FF06020D
 Master DRAB Registers:
  DSR:  2D170CBC  CSR:  00000000 DCSR:  0009FFFF  DER:  00001C00  EAR:  00200400
  EDR:  00000000  ERR:  00000023  RSR:  00801432  CHC:  005BB41A  CMC:  002E853E
 CACHEA0 DRAB Registers:
  DSR:  2D17003F  CSR:  00000000 DCSR:  00003403  DER:  00001C76  EAR:  00000000
  EDR:  FFFFFFFF  ERR:  00000000  RSR:  09805432  CHC:  00000000  CMC:  00000000
 CACHEA1 DRAB Registers:
  Not Available.
 CACHEB0 DRAB Registers:
  DSR:  2C17003D  CSR:  00000000 DCSR:  00003402  DER:  00001C76  EAR:  00000000
  EDR:  FFFFFFFF  ERR:  00000000  RSR:  09805432  CHC:  00000000  CMC:  00000000
 CACHEB1 DRAB Registers:
  Not Available.
 FX Registers:
  PCX[0]:   00000000  PCX[1]:   00000001  PCX[2]:   00000201  PCX[3]:   00000401
  PCX[4]:   00000601  PCX[5]:   00000801  PCX[6]:   00000A01  PCX[7]:   00000000
  CSR:      0087E002  GEN_PCX:  00000000  UNUSED0:  00000000  UNUSED1:  00000000
  DILP:     00000000  DADDR:    00000000  DCMD:     00000000
 Host Port Registers (YACI):
  SET:    072E004A  CIA:    F8070700  PCS:    0505400F  RTS:    00000000
  RADILP: 201BEE08  RBDILP: 201BEE2C  TADILP: 201BEE50  TADILP: 201BEE74
  RADFPA: 2063E760  RBDFPA: 20619560  TADFPA: 20637FA0  TBDFPA: 2062D5A0
  RADNPA: 2063E760  RBDNPA: 20619560  TADNPA: 20637FA0  TBDNPA: 2062D5A0
  REV:    00000003  DIAG:   00000000
 Device Port 0 Registers (NCR710):
```

```
SCNTL0:  CA  SCNTL1:  20  SDID:    00  SIEN:    AF  SCID:    80  SXFER:   00
SODL:    00  SOCL:    00  SFBR:    00  SIDL:    00  SBDL:    00  SBCL:    00
DSTAT:   90  SSTAT0:  00  SSTAT1:  00  SSTAT2:  00  DSA:     00000000
CTEST0:  70  CTEST1:  F0  CTEST2:  21  CTEST3:  08  CTEST4:  00  CTEST5:  00
CTEST6:  FC  CTEST7:  80  TEMP:    00000000      DFIFO:   00  ISTAT:   00
CTEST8:  21  LCRC:    00  DBC:     000000        DCMD:    54
DNAD:    20578F50         DSP:     20578F50      DSPS:    0000015C
SCRATCH: 00000000         DMODE:   80  DIEN:    27  DWT:     4E  DCNTL:  21
ADDER:   205790AC
Device Port 1 Registers (NCR710):
SCNTL0:  DA  SCNTL1:  20  SDID:    08  SIEN:    AF  SCID:    80  SXFER:   18
SODL:    41  SOCL:    00  SFBR:    00  SIDL:    00  SBDL:    00  SBCL:    00
DSTAT:   90  SSTAT0:  00  SSTAT1:  00  SSTAT2:  0F  DSA:     3E000000
CTEST0:  70  CTEST1:  F0  CTEST2:  25  CTEST3:  41  CTEST4:  00  CTEST5:  00
CTEST6:  B0  CTEST7:  80  TEMP:    205788B0      DFIFO:   00  ISTAT:   00
CTEST8:  21  LCRC:    08  DBC:     000000        DCMD:    54
DNAD:    205788B8         DSP:     205788B8      DSPS:    0000015C
SCRATCH: 205788B0         DMODE:   80  DIEN:    27  DWT:     4E  DCNTL:  21
ADDER:   20578A14
Device Port 2 Registers (NCR710):
SCNTL0:  CA  SCNTL1:  20  SDID:    00  SIEN:    AF  SCID:    80  SXFER:   00
SODL:    00  SOCL:    00  SFBR:    00  SIDL:    00  SBDL:    00  SBCL:    00
DSTAT:   90  SSTAT0:  00  SSTAT1:  00  SSTAT2:  00  DSA:     00000000
CTEST0:  70  CTEST1:  F0  CTEST2:  21  CTEST3:  08  CTEST4:  00  CTEST5:  00
CTEST6:  FC  CTEST7:  80  TEMP:    00000000      DFIFO:   00  ISTAT:   00
CTEST8:  21  LCRC:    00  DBC:     000000        DCMD:    54
DNAD:    20578220         DSP:     20578220      DSPS:    0000015C
SCRATCH: 00000000         DMODE:   80  DIEN:    27  DWT:     4E  DCNTL:  21
ADDER:   2057837C
Device Port 3 Registers (NCR710):
SCNTL0:  DA  SCNTL1:  20  SDID:    01  SIEN:    AF  SCID:    80  SXFER:   18
SODL:    02  SOCL:    00  SFBR:    00  SIDL:    00  SBDL:    00  SBCL:    00
DSTAT:   90  SSTAT0:  00  SSTAT1:  00  SSTAT2:  0F  DSA:     3E000000
CTEST0:  70  CTEST1:  F0  CTEST2:  25  CTEST3:  02  CTEST4:  00  CTEST5:  00
CTEST6:  80  CTEST7:  80  TEMP:    20577B80      DFIFO:   00  ISTAT:   00
CTEST8:  21  LCRC:    01  DBC:     000000        DCMD:    54
DNAD:    20577B88         DSP:     20577B88      DSPS:    0000015C
SCRATCH: 20577B80         DMODE:   80  DIEN:    27  DWT:     4E  DCNTL:  21
ADDER:   20577CE4
Device Port 4 Registers (NCR710):
SCNTL0:  DA  SCNTL1:  20  SDID:    01  SIEN:    AF  SCID:    80  SXFER:   18
SODL:    42  SOCL:    00  SFBR:    00  SIDL:    00  SBDL:    00  SBCL:    00
DSTAT:   90  SSTAT0:  00  SSTAT1:  00  SSTAT2:  0F  DSA:     3E000000
CTEST0:  70  CTEST1:  F0  CTEST2:  25  CTEST3:  42  CTEST4:  00  CTEST5:  00
CTEST6:  E8  CTEST7:  80  TEMP:    205774E8      DFIFO:   00  ISTAT:   00
CTEST8:  21  LCRC:    01  DBC:     000000        DCMD:    54
DNAD:    205774F0         DSP:     205774F0      DSPS:    0000015C
SCRATCH: 205774E8         DMODE:   80  DIEN:    27  DWT:     4E  DCNTL:  21
ADDER:   2057764C
Device Port 5 Registers (NCR710):
SCNTL0:  DA  SCNTL1:  20  SDID:    02  SIEN:    AF  SCID:    80  SXFER:   68
SODL:    F6  SOCL:    00  SFBR:    00  SIDL:    00  SBDL:    00  SBCL:    00
DSTAT:   90  SSTAT0:  00  SSTAT1:  00  SSTAT2:  0F  DSA:     3E000000
CTEST0:  70  CTEST1:  F0  CTEST2:  25  CTEST3:  80  CTEST4:  00  CTEST5:  00
CTEST6:  50  CTEST7:  80  TEMP:    20576E50      DFIFO:   00  ISTAT:   00
CTEST8:  21  LCRC:    20  DBC:     000000        DCMD:    54
DNAD:    20576E58         DSP:     20576E58      DSPS:    0000015C
SCRATCH: 20576E50         DMODE:   80  DIEN:    27  DWT:     4E  DCNTL:  21
ADDER:   20576FB4
```

- **Information presented after this callout is extended information available with the FULL option. The output shows register contents for various hardware as well as other data. You should print a copy of this display so that Digital Multivendor Customer Services has more information with which to troubleshoot the controller.**

7.10 The CLONE Utility

CLONE is a utility that partially automates the process of mirroring units to create a snapshot copy of host unit data.

Using the mirroring facility to create a snapshot copy of host unit data is a multistep process described in the *StorageWorks Array Controllers HS Family of Array Controllers User's Guide*. If your data snapshot needs are straightforward, you can use CLONE to relieve you of much of the chore of performing this common operation. CLONE does nothing that cannot be achieved by issuing the appropriate CLI commands yourself; in fact, CLONE operates by issuing CLI commands for you. These commands are printed on the terminal as part of the program output so that you can see what it has done. This is useful both as a learning tool, and in the event that you need to recover an incomplete CLONE operation manually.

Any unit created by CLONE has a mirrorset level in the configuration hierarchy, even if the original unit did not have mirroring. The CLONE utility uses mirrorsets during copying, and must maintain that structure when it adds the cloned disk as a newly-created unit. The following steps illustrate this concept:

1. CLONE is run on a single-disk unit.
2. CLONE creates a mirrorset from the single disk and adds the target disk to the mirrorset. The target disk is initialized *as a member of a mirrorset* when it is added.
3. The data on the first member of the mirrorset is copied to the new (target) member.
4. When the copy is complete, CLONE removes the target drive from the mirrorset, and then sets the original disk so that it is no longer a mirrorset.
5. Because the target disk was initialized as a mirrorset member, CLONE makes it into a single-member mirrorset to preserve the metadata before adding it as a unit.

The new mirrorset level only affects clones of single-disk units and stripesets. Mirrorsets and striped mirrorsets already use mirroring, and any clones from such units have the same structure as the original units.

Although the operations involved in most CLONE processes are commonplace, the reaction in response to failures during the operation is heavily influenced by each installation's operational needs. For this reason, CLONE does not attempt to automate the error recovery process. If you encounter an error during a CLONE operation, you will need to resolve the situation by manually issuing the appropriate CLI commands.

The following circumstance causes CLONE to cease operation:

1. If the controller is reset or there is a power loss.
2. If the host moves the unit being copied from one controller to the other.
3. If you change the configuration of the unit while CLONE is running.
4. If a disk device being used by CLONE fails.
5. If the CLONE utility is aborted via Ctrl/Y or Ctrl/C.
6. If the controller fails.

In most cases, CLONE can automate the operation entirely and completes successfully without intervention. In those cases when it cannot complete, it stops at the point the error occurred. The CLI commands performed up to that point remain on the screen.

Recovering partially completed CLONE operations requires you to be familiar with the data snapshot process as described in the *StorageWorks Array Controllers HS Family of Array Controllers User's Guide*. The recovery strategy is to decide what follow-on operations will best meet your operational needs. You can continue the data snapshot operation to completion by manually issuing the remaining commands in the sequence, or you can undo the actions CLONE has taken so far by issuing appropriate CLI commands to reverse them.

Note

CLONE does not operate on stripesets that have different types of members. All stripeset members must be single disks, or all must be mirrorsets.

Each of the four CLONE examples that follow contain a detailed discussion of the typical sequence of operations that CLONE performs for each of the legal target configurations. If a CLONE operation does not complete, use these template sequences and your knowledge of data snapshot operations in general to guide a manual recovery session.

7.10.1 Cloning a Single Device Unit

```
CLI> RUN CLONE

Clone Local Program Invoked
Units available for cloning:
 110
 799
Enter unit to clone ? 110

Clone will create a new unit which is a copy of unit 110.
Enter the unit number which you want assigned to the new unit ? 797
The new unit may be added using one of the following methods:
  1. Clone will pause after all members have been copied. The user must
     then press RETURN to cause the new unit to be added.
  2. After all members have been copied, the unit will be added automatically.
Under which above method should the new unit be added [] ? 1
Devices available for clone targets:
  DISK300 (size=2050353)
  DISK330 (size=2050353)
Use available device DISK300(size=2050353) for
member DISK110(size=2050353) (y,n) [y] ? y
       mirror DISK110 C_M
       set C_M nopolicy
       set C_M members=2
       set C_M replace=DISK300
Copy in progress for each new member. Please be patient...
  copy from DISK110 to DISK300 is 6% complete
  copy from DISK110 to DISK300 is 12% complete
  copy from DISK110 to DISK300 is 19% complete
    .
    .
    .
  copy from DISK110 to DISK300 is 98% complete
  copy from DISK110 to DISK300 is 100% complete

Press RETURN when you want the new unit to be created
       reduce  DISK300
       unmirror DISK110
       add mirrorset C_M        DISK300
       init C_M       nodestroy
       add unit D797 C_M

D797 has been created. It is a clone of D110.

Clone - Normal Termination

HS1CP1>
```

The following is the list of commands (in sequential order) you would use if you need to manually clone the single device unit from the example above:

1. SHOW DEVICES
2. SHOW STORAGESETS
3. MIRROR *disk-device-name1 container-name*
4. SET *mirrorset-container-name* POLICY= *policy-type*
5. SET *mirrorset-container-name* MEMBERSHIP= *number-of- members*
6. SET *mirrorset-container-name* REPLACE= *disk-device-name*

   Wait for member NORMALIZATION
7. REDUCE *disk-device-name*
8. UNMIRROR *disk-device-name*

9. ADD MIRRORSET *mirrorset-container-name disk-device-name*
10. INITIALIZE *container-name* NODESTROY
11. ADD UNIT *unit-number container-name*

The following is an example of the specific manual commands used to clone the single device unit from the example above:

```
HS1CP1> SHOW DEVICES ❶
HS1CP1> SHOW STORAGESETS ❷
HS1CP1> MIRROR DISK110 C_M ❸
HS1CP1> SET C_M NOPOLICY ❹
HS1CP1> SET C_M MEMBERSHIP=2 ❺
HS1CP1> SET C_M REPLACE=DISK300 ❻
HS1CP1> REDUCE DISK300 ❼
HS1CP1> UNMIRROR DISK110 ❽
HS1CP1> ADD MIRRORSET C_M DISK300 ❾
HS1CP1> INITIALIZE C_M NODESTROY ❿
HS1CP1> ADD UNIT D797 C_M ⓫
```

❶ Displays the units available for cloning.

❷ Displays the storagesets that are currently configured.

❸ Converts the physical device DISK110 into a one-member mirrorset and names the mirrorset C_M.

❹ Sets the replacement policy for mirrorset C_M to NOPOLICY so that spareset members are not automatically added to the mirrorset

❺ Sets the nominal membership of mirrorset C_M to two members.

❻ Places disk DISK300 into mirrorset C_M. The controller automatically starts copying data from the original member to the new member.

❼ Removes DISK300 from mirrorset C_M, after a copy from the existing mirrorset member to DISK300 has completed.

❽ Converts the one-member mirrorset DISK110 back to a physical device.

❾ Creates a one-member mirrorset from DISK300 and names the mirrorset C_M.

❿ Initializes mirrorset C_M, but does not destroy any of the forced error metadata on the disk.

⓫ Creates a logical unit to the controller from the initialized mirrorset C_M and names the unit D797.

7.10.2 Cloning a Stripeset

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

ST1           stripeset                     DISK130          D799
                                            DISK200
HS1CP1> run clone

Clone Local Program Invoked
Units available for cloning:
 110
 799
Enter unit to clone ? 799
Clone will create a new unit which is a copy of unit 799.
Enter the unit number which you want assigned to the new unit ? 798
The new unit may be added using one of the following methods:
  1. Clone will pause after all members have been copied. The user must
     then press RETURN to cause the new unit to be added.
  2. After all members have been copied, the unit will be added automatically.
Under which above method should the new unit be added [] ? 1

Devices available for clone targets:
  DISK220 (size=832317)
  DISK240 (size=832317)
  DISK310 (size=832317)
Use available device DISK220(size=832317) for
member DISK130(size=832317) (y,n) [y] ? y
        mirror DISK130 C_MA
        set C_MA nopolicy
        set C_MA members=2
        set C_MA replace=DISK220
Devices available for clone targets:
  DISK240 (size=832317)
  DISK310 (size=832317)
Use available device DISK240(size=832317) for
member DISK200(size=832317) (y,n) [y] ? y
        mirror DISK200 C_MB
        set C_MB nopolicy
        set C_MB members=2
        set C_MB replace=DISK240
Copy in progress for each new member. Please be patient...
  copy from DISK130 to DISK220 is 15% complete
  copy from DISK200 to DISK240 is 11% complete
  copy from DISK130 to DISK220 is 27% complete
  copy from DISK200 to DISK240 is 23% complete
    .
    .
    .
  copy from DISK130 to DISK220 is 100% complete
  copy from DISK200 to DISK240 is 100% complete

Press RETURN when you want the new unit to be created
        reduce  DISK220 DISK240
        unmirror DISK130
        unmirror DISK200
        add mirrorset C_MA       DISK220
        add mirrorset C_MB       DISK240
        add stripeset C_ST1  C_MA C_MB
        init C_ST1     nodestroy chunk=128
        add unit D798 C_ST1

D798 has been created. It is a clone of D799.

Clone - Normal Termination
```

```
HS1CP1> SHOW DEVICES
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------
DISK130       disk             1    3    0                    ST1
DISK200       disk             2    0    0                    ST1
DISK220       disk             2    2    0                    C_MA
DISK240       disk             2    4    0                    C_MB
DISK310       disk             3    1    0

HS1CP1> SHOW STRIPESETS
Name          Storageset                    Uses               Used by
-------------------------------------------------------------------------
C_ST1         stripeset                     C_MA               D798
                                            C_MB
ST1           stripeset                     DISK130            D799
                                            DISK200

HS1CP1>
```

To manually clone a stripeset, use the following commands in sequential order:

1. SHOW DEVICES
2. SHOW UNITS
3. MIRROR *disk-device-name1 container-name1*
4. SET *mirrorset-container-name1* NOPOLICY
5. SET *mirrorset-container-name1* MEMBERSHIP=
6. SET *mirrorset-container-name1* REPLACE=
7. MIRROR *disk-device-name2 mirrorset-container-name2*
8. SET *mirrorset-container-name2* NOPOLICY
9. SET *mirrorset-container-name2* MEMBERSHIP=
10. SET *mirrorset-container-name2* REPLACE=
    Wait for NORMALIZATION of *all* new mirrorset members.
11. REDUCE *disk-device-name disk-device-name[N]*
12. UNMIRROR *disk-device-name*
13. UNMIRROR *disk-device-name*
14. ADD MIRRORSET *mirrorset-container-name1 disk-device-name*
15. ADD MIRRORSET *mirrorset-container-name2 disk-device-name*
16. ADD STRIPESET *stripeset-container-name mirrorset-container-name1 mirrorset-container-name2*
17. INITIALIZE *stripeset-container-name* NODESTROY CHUNK=
18. ADD UNIT *unit-name stripeset-container-name*

The following is an example of the manual commands for cloning a stripeset:

```
HS1CP1> SHOW DEVICES ❶
HS1CP1> SHOW STORAGESETS ❷
HS1CP1> MIRROR DISK130 C_MA ❸
HS1CP1> SET C_MA NOPOLICY ❹
HS1CP1> SET C_MA MEMBERSHIP=2 ❺
HS1CP1> SET C_MA REPLACE=DISK220 ❻
HS1CP1> MIRROR DISK200 C_MB ❼
HS1CP1> SET C_MB NOPOLICY ❽
HS1CP1> SET C_MB MEMBERSHIP=2 ❾
HS1CP1> SET C_MB REPLACE=DISK240 ❿
...wait for NORMALIZATION OF DISK220 and DISK240...
HS1CP1> REDUCE DISK220 DISK240 ⓫
HS1CP1> UNMIRROR DISK130 ⓬
HS1CP1> UNMIRROR DISK200 ⓭
HS1CP1> ADD MIRRORSET C_MA DISK220 ⓮
HS1CP1> ADD MIRRORSET C_MB DISK240 ⓯
HS1CP1> ADD STRIPESET C_ST1 C_MA C_MB ⓰
HS1CP1> INITIALIZE C_ST1 NODESTROY CHUNK=128 ⓱
HS1CP1> ADD UNIT D798 C_ST1 ⓲
HS1CP1> SHOW DEVICES ⓳
HS1CP1> SHOW STRIPESETS ⓴
```

❶ Displays the units available for cloning.

❷ Displays the storagesets that are currently configured.

❸ Converts the physical device DISK130 into a one-member mirrorset and names the mirrorset C_MA.

❹ Sets the replacement policy for mirrorset C_MA to NOPOLICY so that spareset members are not automatically added to the mirrorset

❺ Sets the nominal membership of mirrorset C_MA to two members.

❻ Places disk DISK220 into mirrorset C_MA. The controller automatically starts copying data from the original member to the new member.

❼ Converts the physical device DISK200 into a one-member mirrorset and names the mirrorset C_MB.

❽ Sets the replacement policy for mirrorset C_MB to NOPOLICY so that spareset members are not automatically added to the mirrorset

❾ Sets the nominal membership of mirrorset C_MB to two members.

❿ Places disk DISK240 into mirrorset C_MB. The controller automatically starts copying data from the original member to the new member.

⓫ Removes DISK220 and DISK240 from their respective mirrorsets, after the copy operations from the existing mirrorset members has completed.

⓬ Converts the one-member mirrorset DISK130 back to a physical device.

⓭ Converts the one-member mirrorset DISK200 back to a physical device.

⓮ Creates a one-member mirrorset from DISK220 and names the mirrorset C_MA.

⓯ Creates a one-member mirrorset from DISK240 and names the mirrorset C_MB.

- Creates a stripeset called C_ST1 from the two newly-created mirrorsets C_MA and C_MB.
- Initializes stripeset C_ST1 with a chunksize of 128 blocks, but does not write new metadata onto the members.
- Creates a logical unit for the host from the initialized stripeset C_ST1 and names the unit D798.
- Shows the device configuration after the clone is complete.
- Shows the stripesets after the clone is complete.

7.10.3 Cloning a Mirrorset

```
HS1CP1> SHOW DEVICES
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------

DISK110       disk             1    1    0                    M1
DISK150       disk             1    5    0
DISK220       disk             2    2    0                    M1
DISK310       disk             3    1    0
DISK350       disk             3    5    0
DISK420       disk             4    2    0
DISK510       disk             5    1    0
DISK550       disk             5    5    0
DISK620       disk             6    2    0

HS1CP1>

HS1CP1> SHOW MIRRORSETS
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------

M1            mirrorset                     DISK110          D411
                                            DISK220

HS1CP1>

HS1CP1> RUN CLONE

Clone Local Program Invoked
Units available for cloning:
 411
Enter unit to clone ?411
Enter the unit number which you want assigned to the new unit?499
The new unit may be added using one of the following methods:
  1. Clone will pause after all members have been copied. The user must
     then press RETURN to cause the new unit to be added.
  2. After all members have been copied, the unit will be added automatically.
Under which above method should the new unit be added []?2
Devices available for clone targets:
  DISK150 (size=832317)
  DISK310 (size=832317)
  DISK350 (size=832317)
  DISK420 (size=832317)
  DISK510 (size=832317)
  DISK550 (size=832317)
  DISK620 (size=832317)
Use available device DISK150(size=832317) for
member DISK110(size=832317) (y,n) [y] ?N
Use available device DISK310(size=832317) for
member DISK110(size=832317) (y,n) [y] ?Y
        set M1 nopolicy
        set M1 members=3
        set M1 replace=DISK310
```

```
Copy in progress for each new member. Please be patient...
  copy from DISK110 to DISK310 is 5% complete
  copy from DISK110 to DISK310 is 10% complete
  .
  .
  .
  copy from DISK110 to DISK310 is 97% complete
  copy from DISK110 to DISK310 is 100% complete
        reduce DISK310
        add mirrorset C_M1       DISK310
        init C_M1      nodestroy
        add unit D499 C_M1

D499 has been created. It is a clone of D411.

Clone - Normal Termination

HS1CP1>

HS1CP1> SHOW DEVICES
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------

DISK110       disk             1    1    0                    M1
DISK150       disk             1    5    0
DISK220       disk             2    2    0                    M1
DISK310       disk             3    1    0                    C_M1
DISK350       disk             3    5    0
DISK420       disk             4    2    0
DISK510       disk             5    1    0
DISK550       disk             5    5    0
DISK620       disk             6    2    0

HS1CP1>

HS1CP1> SHOW MIRRORSETS
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------

C_M1          mirrorset                     DISK310          D499

M1            mirrorset                     DISK110          D411
                                            DISK220

HS1CP1>
```

To manually clone a mirrorset, use the following commands in sequential order:

1. SHOW DEVICES
2. SHOW MIRRORSETS
3. SET *mirrorset-container-name* NOPOLICY
4. SET *mirrorset-container-name* MEMBERSHIP=3
5. SET *mirrorset-container-name* REPLACE= *disk-device-name*
   Wait for NORMALIZATION.
6. REDUCE *disk-device-name*
7. ADD MIRRORSET *mirrorset-container-name disk-device-name*
8. INITIALIZE *mirrorset-container-name* NODESTROY CHUNK=
9. ADD UNIT *unit-name mirrorset-container-name*
10. SHOW DEVICES

11. SHOW MIRRORSETS

The following is an example of the commands for cloning a mirrorset:

```
HS1CP1> SHOW DEVICES ❶
HS1CP1> SHOW MIRRORSETS ❷
HS1CP1> SET M1 NOPOLICY ❸
HS1CP1> SET M1 MEMBERSHIP=3 ❹
HS1CP1> SET M1 REPLACE=DISK310 ❺

...wait for NORMALIZATION of M1...

HS1CP1> REDUCE DISK310 ❻
HS1CP1> ADD MIRRORSET C_M1 DISK310 ❼
HS1CP1> INITIALIZE C_M1  NODESTROY CHUNK=128 ❽
HS1CP1> ADD UNIT D499 C_M1 ❾
HS1CP1> SHOW DEVICES ❿
HS1CP1> SHOW MIRRORSETS ⓫
```

❶ Displays the units available for cloning.

❷ Displays the mirrorsets that are currently configured.

❸ Sets the replacement policy for mirrorset M1 to NOPOLICY so that spareset members are not automatically added to the mirrorset

❹ Sets the nominal membership of mirrorset M1 to three members.

❺ Places disk DISK310 into mirrorset M1. The controller automatically starts copying data to the new member to make it identical to the existing NORMAL members.

❻ Removes DISK310 from mirrorset M1 after the copy operation has completed.

❼ Creates a mirrorset called C_M1 from DISK310.

❽ Initializes mirrorset C_M1 with a chunksize of 128 blocks, but does not write new metadata onto the members.

❾ Creates a logical unit for the host from mirrorset C_M1 and names the unit D499.

❿ Shows the device configuration after the clone is complete.

⓫ Shows the mirrorsets after the clone is complete.

7.10.4 Cloning a Striped Mirrorset

```
HS1CP1> SHOW DEVICES
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------
DISK110       disk             1    1    0                    M1
DISK150       disk             1    5    0
DISK220       disk             2    2    0                    M1
DISK310       disk             3    1    0                    M2
DISK350       disk             3    5    0
DISK420       disk             4    2    0                    M2
DISK510       disk             5    1    0                    M3
DISK550       disk             5    5    0
DISK620       disk             6    2    0                    M3

HS1CP1> SHOW MIRRORSETS
Name          Storageset                    Uses               Used by
-------------------------------------------------------------------------
M1            mirrorset                     DISK110            ST1
                                            DISK220

M2            mirrorset                     DISK310            ST1
                                            DISK420

M3            mirrorset                     DISK510            ST1
                                            DISK620

HS1CP1>

HS1CP1> SHOW STORAGESETS
Name          Storageset                    Uses               Used by
-------------------------------------------------------------------------
ST1           stripeset                     M1                 D411
                                            M2
                                            M3

HS1CP1>

HS1CP1> RUN CLONE
```

```
Clone Local Program Invoked
Units available for cloning:
 411
Enter unit to clone ?411
Clone will create a new unit which is a copy of unit 411.
Enter the unit number which you want assigned to the new unit?499
The new unit may be added using one of the following methods:
  1. Clone will pause after all members have been copied. The user must
     then press RETURN to cause the new unit to be added.
  2. After all members have been copied, the unit will be added automatically.
Under which above method should the new unit be added []?1
Devices available for clone targets:
  DISK150 (size=832317)
  DISK350 (size=832317)
  DISK550 (size=832317)
Use available device DISK150(size=832317) for
member DISK110(size=832317) (y,n) [y] ?N
Use available device DISK350(size=832317) for
member DISK110(size=832317) (y,n) [y] ?Y
        set M1 nopolicy
        set M1 members=3
        set M1 replace=DISK350
Devices available for clone targets:
  DISK150 (size=832317)
  DISK550 (size=832317)
Use available device DISK150(size=832317) for
member DISK310(size=832317) (y,n) [y] ?Y
        set M2 nopolicy
        set M2 members=3
        set M2 replace=DISK150
Devices available for clone targets:
  DISK550 (size=832317)
Use available device DISK550(size=832317) for
member DISK510(size=832317) (y,n) [y] ?Y
        set M3 nopolicy
        set M3 members=3
        set M3 replace=DISK550
Copy in progress for each new member. Please be patient...
  copy from DISK110 to DISK350 is 7% complete
  copy from DISK310 to DISK150 is 6% complete
  copy from DISK510 to DISK550 is 5% complete
  .
  .
  .
  copy from DISK110 to DISK350 is 100% complete
  copy from DISK310 to DISK150 is 100% complete
  copy from DISK510 to DISK550 is 100% complete

Press RETURN when you want the new unit to be created
        reduce DISK350 DISK150 DISK550
        add mirrorset C_M1      DISK350
        add mirrorset C_M2      DISK150
        add mirrorset C_M3      DISK550
        add stripeset C_ST1     C_M1 C_M2 C_M3
        init C_ST1    nodestroy chunk=128
        add unit D499 C_ST1

D499 has been created. It is a clone of D411.

Clone - Normal Termination

HS1CP1>

HS1CP1> SHOW DEVICES

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

```
DISK110       disk            1    1    0                    M1
DISK150       disk            1    5    0                    C_M2
DISK220       disk            2    2    0                    M1
DISK310       disk            3    1    0                    M2
DISK350       disk            3    5    0                    C_M1
DISK420       disk            4    2    0                    M2
DISK510       disk            5    1    0                    M3
DISK550       disk            5    5    0                    C_M3
DISK620       disk            6    2    0                    M3

HS1CP1>

HS1CP1> SHOW MIRRORSETS

Name          Storageset                    Uses              Used by
------------------------------------------------------------------------

C_M           mirrorset                     DISK350           C_ST1

C_MA          mirrorset                     DISK150           C_ST1

C_MB          mirrorset                     DISK550           C_ST1

M1            mirrorset                     DISK110           ST1
                                            DISK220

M2            mirrorset                     DISK310           ST1
                                            DISK420

M3            mirrorset                     DISK510           ST1
                                            DISK620

HS1CP1>

HS1CP1> SHOW STORAGESETS

Name          Storageset                    Uses              Used by
------------------------------------------------------------------------

C_ST1         stripeset                     C_M1              D499
                                            C_M2
                                            C_M3

ST1           stripeset                     M1                D411
                                            M2
                                            M3

HS1CP1>
```

To manually clone a striped mirrorset, use the following commands in sequential order:

1. SHOW DEVICES
2. SHOW MIRRORSETS
3. SHOW STORAGESETS
4. SET *mirrorset-container-name1* NOPOLICY
5. SET *mirrorset-container-name1* MEMBERSHIP=
6. SET *mirrorset-container-name1* REPLACE=
7. SET *mirrorset-container-name2* NOPOLICY
8. SET *mirrorset-container-name2* MEMBERSHIP=
9. SET *mirrorset-container-name2* REPLACE=
10. SET *mirrorset-container-name3* NOPOLICY
11. SET *mirrorset-container-name3* MEMBERSHIP=
12. SET *mirrorset-container-name3* REPLACE=

    Wait for NORMALIZATION of *all* new mirrorset members.
13. REDUCE *disk-device-name disk-device-name[N] disk-device-name [N]*
14. ADD MIRRORSET *mirrorset-container-name disk-device-name*
15. ADD MIRRORSET *mirrorset-container-name disk-device-name*
16. ADD MIRRORSET *mirrorset-container-name disk-device-name*
17. ADD STRIPESET *stripeset-name mirrorset-container-name1 mirrorset-container- name2 mirrorset-container-name3*
18. INITIALIZE *stripeset-name* NODESTROY CHUNK=
19. ADD UNIT *unit-name stripeset-name*
20. SHOW DEVICES
21. SHOW MIRRORSETS
22. SHOW STORAGESETS

The following is an example of the commands for cloning striped mirrorsets:

```
HS1CP1> SHOW DEVICES ❶
HS1CP1> SHOW MIRRORSETS ❷
HS1CP1> SHOW STORAGESETS ❸
HS1CP1> SET M1 NOPOLICY ❹
HS1CP1> SET M1 MEMBERSHIP=3 ❺
HS1CP1> SET M1 REPLACE=DISK350 ❻
HS1CP1> SET M2 NOPOLICY ❼
HS1CP1> SET M2 MEMBERSHIP=3 ❽
HS1CP1> SET M2 REPLACE=DISK150 ❾
HS1CP1> SET M3 NOPOLICY ❿
HS1CP1> SET M3 MEMBERSHIP=3 ⓫
HS1CP1> SET M3 REPLACE=DISK550 ⓬
...wait for NORMALIZATION of M1, M2, and M3...
```

```
HS1CP1> REDUCE DISK350 DISK150 DISK550 ❶
HS1CP1> ADD MIRRORSET C_M1  DISK350 ❷
HS1CP1> ADD MIRRORSET C_M2  DISK150 ❸
HS1CP1> ADD MIRRORSET C_M3  DISK550 ❹
HS1CP1> ADD STRIPESET C_ST1 C_M1 C_M2 C_M3 ❺
HS1CP1> INITIALIZE C_ST1 NODESTROY ❻
HS1CP1> ADD UNIT D499 C_ST1 ❼
HS1CP1> SHOW DEVICES ❽
HS1CP1> SHOW MIRRORSETS ❾
HS1CP1> SHOW STORAGESETS ❿
```

❶ Displays the devices available for cloning.

❷ Displays the configured mirrorsets.

❸ Displays the configured storagesets.

❹ Sets the replacement policy for mirrorset M1 to NOPOLICY so that spareset members are not automatically added to the mirrorset

❺ Sets the nominal membership of mirrorset M1 to three members.

❻ Places disk DISK350 into mirrorset M1. The controller automatically starts copying data from existing NORMAL members to the new member.

❼ Sets the replacement policy for mirrorset M2 to NOPOLICY so that spareset members are not automatically added to the mirrorset

❽ Sets the nominal membership of mirrorset M2 to three members.

❾ Places disk DISK150 into mirrorset M2. The controller automatically starts copying data from existing NORMAL members to the new member.

❿ Sets the replacement policy for mirrorset M3 to NOPOLICY so that spareset members are not automatically added to the mirrorset

⓫ Sets the nominal membership of mirrorset M3 to three members.

⓬ Places disk DISK550 into mirrorset M3. The controller automatically starts copying data from existing NORMAL members to the new member.

⓭ Removes DISK350, DISK150, and DISK550 from their respective mirrorsets, after the copy operations from the existing NORMAL mirrorset members has completed.

⓮ Creates new one-member mirrorset C_M1 that contains DISK350.

⓯ Creates new one-member mirrorset C_M2 that contains DISK150.

⓰ Creates new one-member mirrorset C_M3 that contains DISK550.

⓱ Creates a stripeset called C_ST1 from the three new mirrorsets C_M1, C_M2, and C_M3.

⓲ Initializes stripeset C_ST1 with a chunksize of 128 blocks, but does not write new metadata onto the members.

⓳ Creates a logical unit for the host from stripeset C_ST1 and names the unit D499.

⓴ Shows the device configuration after the clone is complete.

- Shows the mirrorsets after the clone is complete.
- Shows the storagesets after the clone is complete.

8

Removing and Replacing Field Replaceable Units

This chapter describes how to remove and replace/install field replaceable units (FRUs) in both dual-redundant and nonredundant configurations. This chapter provides information on the following:

- Electrostatic discharge protection
- Accessing FDDI Server components
- Connecting a terminal to the server processor
- Replacing a server processor in a dual server processor FDDI Server
- Preparing to replace server processor components
- Replacing internal server processor components accessible from the rear
- Replacing internal server processor components accessible from the front
- Reassembling the server processor
- Reinstalling the server processor
- Replacing the cable distribution unit
- Replacing FDDI Server shelf components
- Removing and replacing the HS1CP and HS1CP components
- Removing and replacing fans
- Replacing cabinet cables and cords

This chapter also discusses how to warm swap some of the FDDI Server components.

Note

Warm swapping internal server processor components is not recommended at any time.

Appendix B provides a list of FRUs, expendable parts, and required tools.

8.1 Electrostatic Discharge Protection

Electrostatic discharge (ESD) can damage subsystem components. This section describes the necessary procedures for protecting the subsystem components against ESD.

Use the following strategies to minimize electrostatic discharge problems:

- Maintain more than 40 percent humidity in the room where your subsystem resides.
- Place the subsystem cabinet away from heavy traffic paths.
- If possible, do not work in a carpeted area. If carpeting is required, choose an antistatic carpet. If any type of carpeting is installed, place antistatic mats around the system to decrease electrostatic discharge.

ESD Grounding Procedure

Prior to removing or replacing any module, do the following:

1. Obtain and attach an ESD wrist strap to your wrist. Ensure that the strap fits snugly around your wrist.
2. Plug or clip the other end of the ESD wrist strap to a ESD bolt or ground stud usually located on the cabinet's vertical rail that is common for both the HS1CP shelves and the storage shelves.
3. Obtain and use an approved antistatic bag and/or a grounded antistatic mat.

CAUTION

Attempting to replace or repair components within FRUs can cause equipment damage. Use the controller fault indications and error logs to isolate FRU-level failures.

8.2 Accessing FDDI Server Components

It is not always necessary to remove the top and sides of the SW800 cabinet to perform maintenance activities on FDDI Server components. Use the appropriate steps of the following procedure to remove power and gain access to the cabinet interior for equipment removal and replacement.

Digital recommends full and unrestricted access to the cabinet interior if you are replacing and routing cables during maintenance activities.

WARNING

While working in the cabinet interior, you must remove ac power from cabinet components. Failure to do so can result in personal injury as a result of electric shock.

If you need to power down the full cabinet, first spin down all disk drives and halt all tape operation in the cabinet. Switch the circuit breakers on the server processor, HS1CP, and the CDUs to the OFF (○) position.

1. Remove power from the failed server processor, HS1CP, or CDU. Use the shutdown procedure described for each component in the appropriate section of this manual.
2. If required, remove power cords from the wall receptacles.

3. Move the cabinet away from any adjacent cabinets to allow 360 degree access.

Note

You must remove the sides and top of the cabinet to access the cables or cords.

4. Using Figure 8–1 for reference, remove the side panels as follows:

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.

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

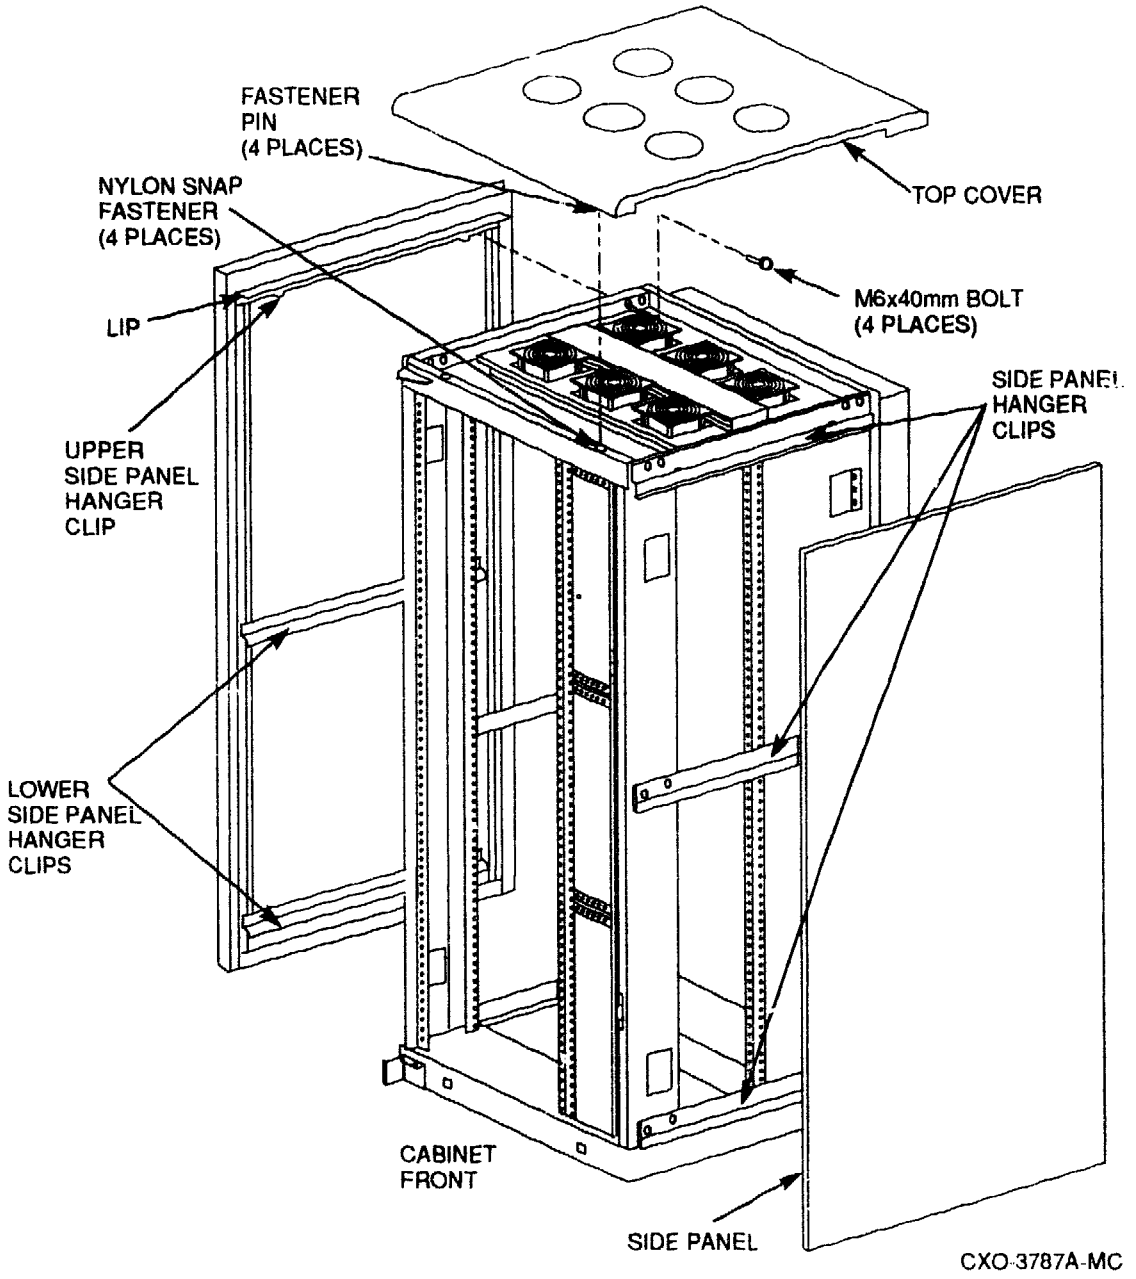
WARNING

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

   b. Using two people, lift the top cover from the cabinet and set it aside.
   c. Remove the bolts attaching the side panels to the top side rails of the cabinet.
   d. Grasp a panel along its front and rear edges and lift up until the hanger clips disengage. Lift the panel away from the cabinet.
   e. Repeat the previous step to remove the opposite panel.
5. Release the door locks mounted on the smooth vertical panel of each door by turning the locks counterclockwise with a 5/32-inch Allen wrench.
6. Open the cabinet doors.

Replace the cabinet doors, side panels, and top cover by reversing the previous procedures.

Figure 8–1 Exterior Cabinet Panel Removal and Installation



8.3 Connecting a Terminal to the Server Processor

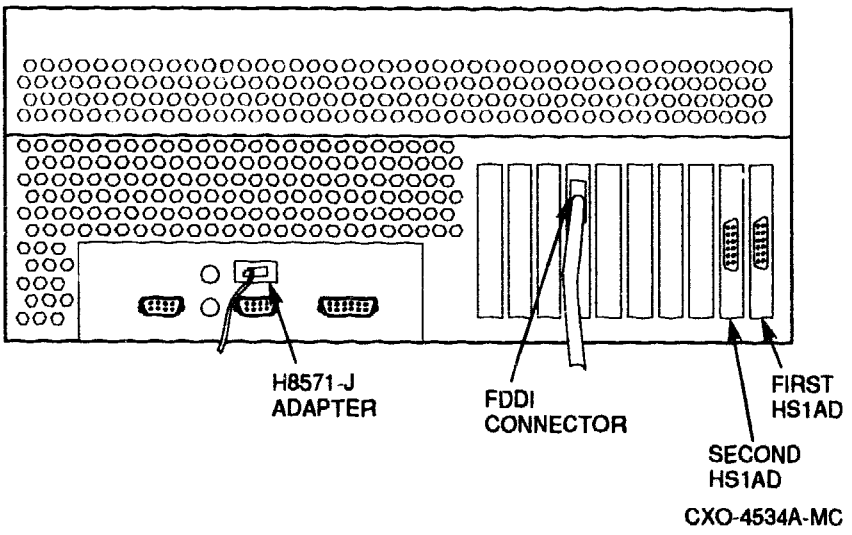
Note

You must have a console terminal connected to the server processor to complete the procedures described in this chapter. If you already have a terminal connected, proceed to the first step of the replacement process. If you do not have a terminal connected, use the following procedure. Any terminal that supports ANSI control sequences can be used, including

graphics displays that provide emulation of an ANSI-compatible video terminal.

1. Make sure the power switch on the back of the terminal is OFF (○).
2. Connect one end of the terminal cable to the back of the terminal.
3. Using a 5/32-inch hex wrench, open the rear door of the FDDI Server cabinet.
4. Connect the other end of the terminal cable to the EIA terminal port on the rear of the server processor using the female MMJ adapter (part number H8571–J) as shown in Figure 8–2.
5. Turn the terminal power switch to ON (|).
6. Set the terminal's communication setup to 9600 baud, with 8 data bits, 1 stop bit, and no parity. Refer to your terminal documentation for terminal setup instructions.

Figure 8–2 Back View of the Server Processor



8.4 Replacing a Server Processor in a Dual Server Processor FDDI Server

Digital recommends that you should power everything off prior to servicing equipment such as a server processor CPU. However, if the FDDI Server must remain running while servicing one of the server processors, then the following procedure provides a means to hot swap a server processor.

There are slight variations to this procedure depending on whether terminated or non-terminated internal bus cables are installed in the FDDI Server. These variations are highlighted during the procedure.

Note

The Digital Part # for the non-terminated internal bus cable is 17–03855–*xx*.

The Digital Part # for the terminated internal bus cable is 17–04288–*xx*. In addition, the terminated cable has a ferrite bead approximately six inches from the straight connector end of the cable.

8.4.1 Removing the Internal Bus Cables

Use the following procedure to remove the internal bus cable connecting the HS1CP device channel processors to the failed server processor and install terminators, if necessary, on the HS1CPs while the server processor is out-of-service.

1. Minimize Server traffic as much as possible by stopping all unnecessary jobs.
2. Run the VMS shutdown procedure on the failing server processor by entering the following command:

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

   at the DCL **$** prompt. When shutdown completes and the **>>>** prompt is displayed, then power off the failed server processor. Disconnect all ac power cords from the failed server processor.
3. From the Server console terminal at the DCL prompt enter the following 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). HS1CP1 and HS1CP2 will be used throughout this procedures to reference the two device channel processors in a dual-redundant pair.
4. You will now be at the Command Line Interpreter of HS1CP1 and the prompt **HS1CP1>** will appear on your screen.
5. Failover HS1CP2 by entering the following command:

   ```
   HS1CP1> SHUTDOWN OTHER OVERRIDE_ONLINE
   ```

   Wait for the following message **BEFORE** proceeding to the next step:

   ```
   %PAA0, Software is Closing Virtual Circuit - REMOTE NODE HS1CP2
   ```
6. It may take some time based on work load for HS1CP2 to failover to HS1CP1. You can monitor this failover process by entering

   ```
   HS1CP1> SHOW THIS
   ```

   When the display shows the following message, failover is complete:

   ```
   Configured for dual-redundancy with ZG43100034
   All devices failed over to this controller
   ```
7. Disconnect the internal bus cable (see appropriate steps below):

   Nonterminated Internal Bus Cable (Digital Part Number 17–03855–*xx*)

   WARNING

   The following steps are very CRITICAL when working with a powered-on HS1CP. With the nonterminated internal bus cable (Digital Part Number 17–03855–*xx*), there is a RISK of damaging the device bus adapter in the server processor if you are not careful.

During the removal/installation of internal bus cables and terminators while power is asserted, you must make sure that the two mating pieces between the cable and trilink connector are perfectly aligned horizontally and vertically. If these connecting pieces are skewed in any manner upon attachment, there is a very strong possibility of shorting two adjacent pins together and blowing one or more fuses. Removing the trilink connector assembly (with internal bus cables still attached) from the HS1CP will help minimize this risk.

Digital does not recommend working on the DSSI bus with power on.

- At the trilink connector of HS1CP2, disconnect the internal bus cable that attaches this failed over HS1CP to the powered off server processor. Detach the other end of the internal bus cable from the server processor.
- On the trilink of HS1CP2, install the DSSI terminator in the connector from which you just removed the DSSI cable.

Terminated Internal Bus Cable (Digital Part Number 17-04288-*xx*)

- Detach the internal bus cable from the server processor **ONLY**. Do **NOT** detach the internal bus cable from the HS1CP.

Note

The HS1CP is terminated since the terminated internal bus cable has a terminator incorporated in the straight-angle connector that attaches to the server processor.

---

8. At the **HS1CP1>** prompt on the console terminal enter the following command:

```
HS1CP1> RESTART OTHER
```

9. Now depress the green Restart button on HS1CP2.

Important

Wait for the LED on the Restart button to start blinking BEFORE proceeding to the next step. The blinking LED indicates a successful hardware restart has occurred.

---

10. At the **HS1CP1>** prompt on the console terminal enter the following command:

```
HS1CP1> SHOW OTHER
```

If the message:

```
Configured for dual-redundancy with ZG43100039
In dual-redundant configuration
    .
    .
    .
Other controller restarted
```

is **NOT** displayed, then HS1CP2 has not completed its restart.

Continue to enter SHOW OTHER until this message appears.

11. At the **HS1CP1>** prompt on the console terminal enter the following command:

```
HS1CP1> EXIT
```

12. Perform Steps 3-11 for the second HS1CP in this dual-redundant pair.
13. Unless you are working on a model HS241 FDDI Server, you are now finished with this procedure until you are ready to reinstall a replacement server processor or you are have replaced the failed device bus adapter. At that time you should proceed to Section 8.4.3. If you are working on a model HS241 FDDI Server, you must repeat the steps in this section to disconnect the internal bus cables from the second HS1CP dual-redundant pair to the failed server processor.

Note

You can proceed with normal workload processing through your Server. You might notice degraded performance until the second server processor is replaced.

8.4.2 Exchanging the Server Processor

Use the following procedure to remove the failed server processor from the cabinet:

CAUTION

The server processor can weigh nearly 80 pounds when fully configured. Unless proper safety precautions are used when removing or installing the server processor, injury can result. The server processor requires two people to assure safe handling during removal or replacement.

CAUTION

Do not replace the server processor or its internal components while the power is applied or the system remains on-line.

1. Disconnect the FDDI connections from the rear of the server processor. Make sure that the FDDI cable is secured to prevent damage during the maintenance activity.
2. Disconnect the terminal data connection from the EIA terminal port.
3. Remove the red shipping brackets that hold the server processor in place.
4. With all the cables disconnected, slide the server processor toward the rear until it is stopped by the slide insert locks.
5. To remove the server processor from the cabinet, press in on each of the slide insert locks to move the server processor to the end of its slide, and remove it from the cabinet.

After performing maintenance on the server processor or receiving a replacement server processor, reverse the above steps to reinstall the server processor in the cabinet. For more detailed server processor installation procedures, see Sections 8.8 and 8.9.

8.4.3 Reinstalling the Internal Bus Cables

After reinstalling the server processor in the cabinet perform the following procedure to reinstall the internal bus cables between the HS1CPs and the serviced server processor:

1. Minimize Server traffic as much as possible by stopping all unnecessary jobs.
2. From the Server console terminal at the DCL prompt enter the following:

   ```
   $ 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). HS1CP1 and HS1CP2 will be used throughout this procedures to reference the two device channel processors in a dual-redundant pair.
3. You will now be at the Command Line Interpreter of HS1CP1 and the prompt **HS1CP1>** will appear on your screen.
4. Failover HS1CP2 by entering:

   ```
   HS1CP1> SHUTDOWN OTHER OVERRIDE_ONLINE
   ```

   Wait for the following message **BEFORE** proceeding to the next step:

   ```
   %PAA0, Software is Closing Virtual Circuit - REMOTE NODE HS1CP2
   ```
5. It may take some time based on work load for HS1CP2 to failover to HS1CP1. You can monitor this failover process by entering

   ```
   HS1CP1> SHOW THIS
   ```

   When the display shows the following message, failover is complete:

   ```
   Configured for dual-redundancy with ZG43100034
   All devices failed over to this controller
   ```
6. Reconnect the internal bus cable (see appropriate steps below):

   Nonterminated Internal Bus Cable (Digital Part Number 17–03855–*xx*)

   WARNING

   The following steps are very CRITICAL when working with a powered-on HS1CP. With the nonterminated internal bus cable (Digital Part Number 17–03855–*xx*), there is a RISK of damaging the device bus adapter in the server processor if you are not careful.

   During the removal/installation of internal bus cables and terminators while power is asserted, you must make sure that the two mating pieces between the cable and trilink connector are perfectly aligned horizontally and vertically. If these connecting pieces are skewed in any manner upon attachment, there is a very strong possibility of shorting two adjacent pins together and blowing one or more fuses. Removing the trilink connector assembly (with internal bus cables still attached) from the HS1CP will help minimize this risk.

   Digital does not recommend working on the DSSI bus with power on.

- To remove the trilink connector from HS1CP2, insert a thin flathead screwdriver between the internal bus cable and the terminator and loosen the top and bottom screws on the trilink connector.
- Remove the trilink assembly from the HS1CP.
- Disconnect the DSSI terminator from the trilink assembly.
- Connect the right-angle end of the internal bus cable to the trilink connector previously removed from the HS1CP and connect the other end of the cable to the serviced server processor.
- Install the trilink connector assembly back onto the HS1CP.

Terminated Internal Bus Cable (Digital Part Number 17–04288–*xx*)

- Reattach the internal bus cable to the server processor.

Note

The right-angle end of the internal bus cable should still be attached to the trilink connector on the HS1CP.

7. At the **HS1CP1>** prompt on the console terminal enter:

```
HS1CP1> RESTART OTHER
```

8. Depress the green Restart button on HS1CP2.

Important

Wait for the LED on the Restart button to start blinking **BEFORE** proceeding to the next step. The blinking LED indicates a successful hardware restart has occurred.

9. At the **HS1CP1>** prompt on the console terminal enter:

```
HS1CP1> SHOW OTHER
```

If the message:

```
Configured for dual-redundancy with ZG43100039
In dual-redundant configuration
   .
   .
   .
Other controller restarted
```

is **NOT** displayed, then HS1CP2 has not completed its restart.

Continue to enter SHOW OTHER until this message appears.

10. At the **HS1CP1>** prompt on the console terminal enter:

```
HS1CP1> EXIT
```

11. Perform Steps 2-10 to reinstall the internal bus cables on the second HS1CP in this dual-redundant pair.
12. For reinstallation of the internal bus cables for a model HS241 FDDI Server, repeat the steps in this section for the second HS1CP dual-redundant pair.

13. Power on the serviced server processor. The console terminal should be connected to this server processor so you can monitor activity.
14. On the console terminal at the >>> prompt enter:

```
>>> BOOT
```

15. After the second server processor has booted, depress the reset button on both HS1CPs so that load balancing can be achieved between the two HS1CP controllers.
16. Use the following procedure to verify that the HS221 FDDI Server is correctly integrated into the VMScluster environment:
    1. Enter the SHOW CLUSTER command on the terminal that is connected to one of the server processors.
    2. Enter the SHOW CLUSTER command from one or more client nodes to ensure that the client nodes can see both nodes of this dual server processor FDDI Server.
    3. Enter the SHOW DEVICE command to ensure the FDDI Server is serving storage to the clients. All storage units, stripesets, and RAIDsets should appear in the resultant display.
17. Maintenance is complete and you have full use of the FDDI Server once again.

8.5 Preparing to Replace Server Processor Components

Before removing and replacing components, you need to assemble the required equipment, familiarize yourself with antistatic precautions (refer to Section 8.1), and access the individual cables or FDDI Server components (see Section 8.2).

8.5.1 Required Equipment

You need the following equipment to perform the installation and removal procedures described in the following sections:

- Flathead screwdriver
- Phillips screwdriver, large—#2
- Phillips screwdriver, small—#0
- Antistatic wrist strap
- Replacement option board or component

8.5.2 Shutting Down the System

Use the following procedure to shut down the system:

1. Stop all work on the server processor in which a component is to be replaced.
2. Perform an orderly shutdown of the operating system in the server processor by entering the following command at the DCL $ prompt:

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

3. When shutdown is complete, open the front door on the FDDI Server cabinet and depress the power button on the server processor OCP to turn off power.
4. Set the On/Off switches on all external connections to the system to the Off position.

5. Set the On/Off switch on the system unit to the Off position.
6. If required, remove the FDDI and EISA data cables from the rear of the server processor.

8.5.3 Accessing the Server Processor

You may need to remove the server processor from the cabinet to access components. The server processor can be removed either from the front or the back of the cabinet, but the power cord must be accessed through the back door. If you need to remove the server processor, use the following procedure:

CAUTION

The server processor can weigh nearly 80 pounds when fully configured. Unless proper safety precautions are used when removing or installing the server processor, injury can result. The server processor requires two people to assure safe handling during removal or replacement.

CAUTION

Do not replace the server processor or its internal components while the power is applied or the system remains on-line.

1. Stop all work on the FDDI Server in which a component is to be replaced.
2. Perform an orderly shutdown of the operating system in the server processor using the procedure in Section 8.5.2.
3. Disconnect the power cable from the front of the server processor.
4. Disconnect the FDDI connections from the rear of the server processor (see Figure 8–2). Make sure that the FDDI cable is secured to prevent damage during the maintenance activity.
5. Disconnect the bus cable connectors from the device bus adapter connections on rear of the server processor and stow the cables out of the way.
6. Disconnect the terminal data connection from the EIA terminal port.
7. Remove the red shipping brackets that hold the server processor in place.
8. With all the cables disconnected, slide the server processor toward the rear until it is stopped by the slide insert locks.
9. To remove the server processor from the cabinet (if necessary), press in on each of the slide insert locks to move the server processor to the end of its slide, and remove it from the cabinet.

   At this point, you can either access the internal components of the server processor or exchange it for a replacement unit. See Section 8.5.4 for procedures on removing the covers from the server processor to access components. Refer to 8–1 for a list of components accessible from the front or back compartment of the server processor.

To close the server processor and replace it in the cabinet, reverse the previous procedure.

Table 8–1 Server Processor Components

Front	Rear
Power supply	HS1AD
Media tray	CPU daughterboard
Floppy drive	FDDI adapter
CD–ROM drive	Memory modules
OCP tray	Motherboard
Fan tray assembly	SCSI hard drive
(Fan) Baffle assembly	

8.5.4 Server Processor Cover Panels

The following sections discuss how to remove the cover panels of the server processor.

Removing the Rear Top Cover

To remove the rear top cover panel, refer to Figure 8–3 and follow these steps:

1. Loosen the six Phillips screws with a quarter turn counter clockwise.
2. Lift the cover slightly and pull forward.

CAUTION

Damage to the equipment can result if any of the metallic edges of the finger stock (EMF shield) around the inside of the top cover break off and fall into the internal server processor circuits.

3. Inspect the finger stock for missing edges.
4. Inspect the interior of the server processor for finger stock edges. Remove broken edges before applying power.

Reverse the steps in this procedure to replace the top cover.

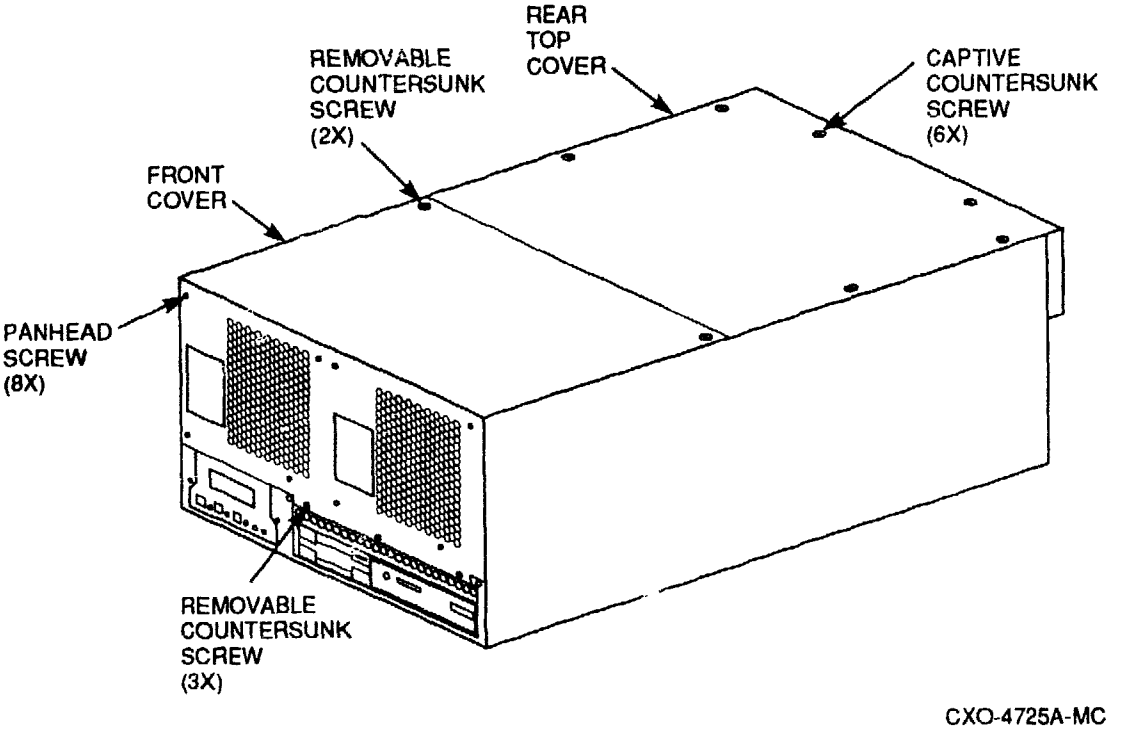
Removing the Front Cover

To remove the front cover panel, refer to Figure 8–3 and follow these steps:

1. If you have not done so already, remove the power cable from the server processor.
2. Remove the screws that attach the power supply to the front cover.
3. Remove the screws attaching the front cover to the chassis using a Phillips screwdriver.
4. Remove the cover by pulling forward slightly and lifting off.

Reverse the steps in this procedure to replace the front cover.

Figure 8–3 Removing the Cover Panels



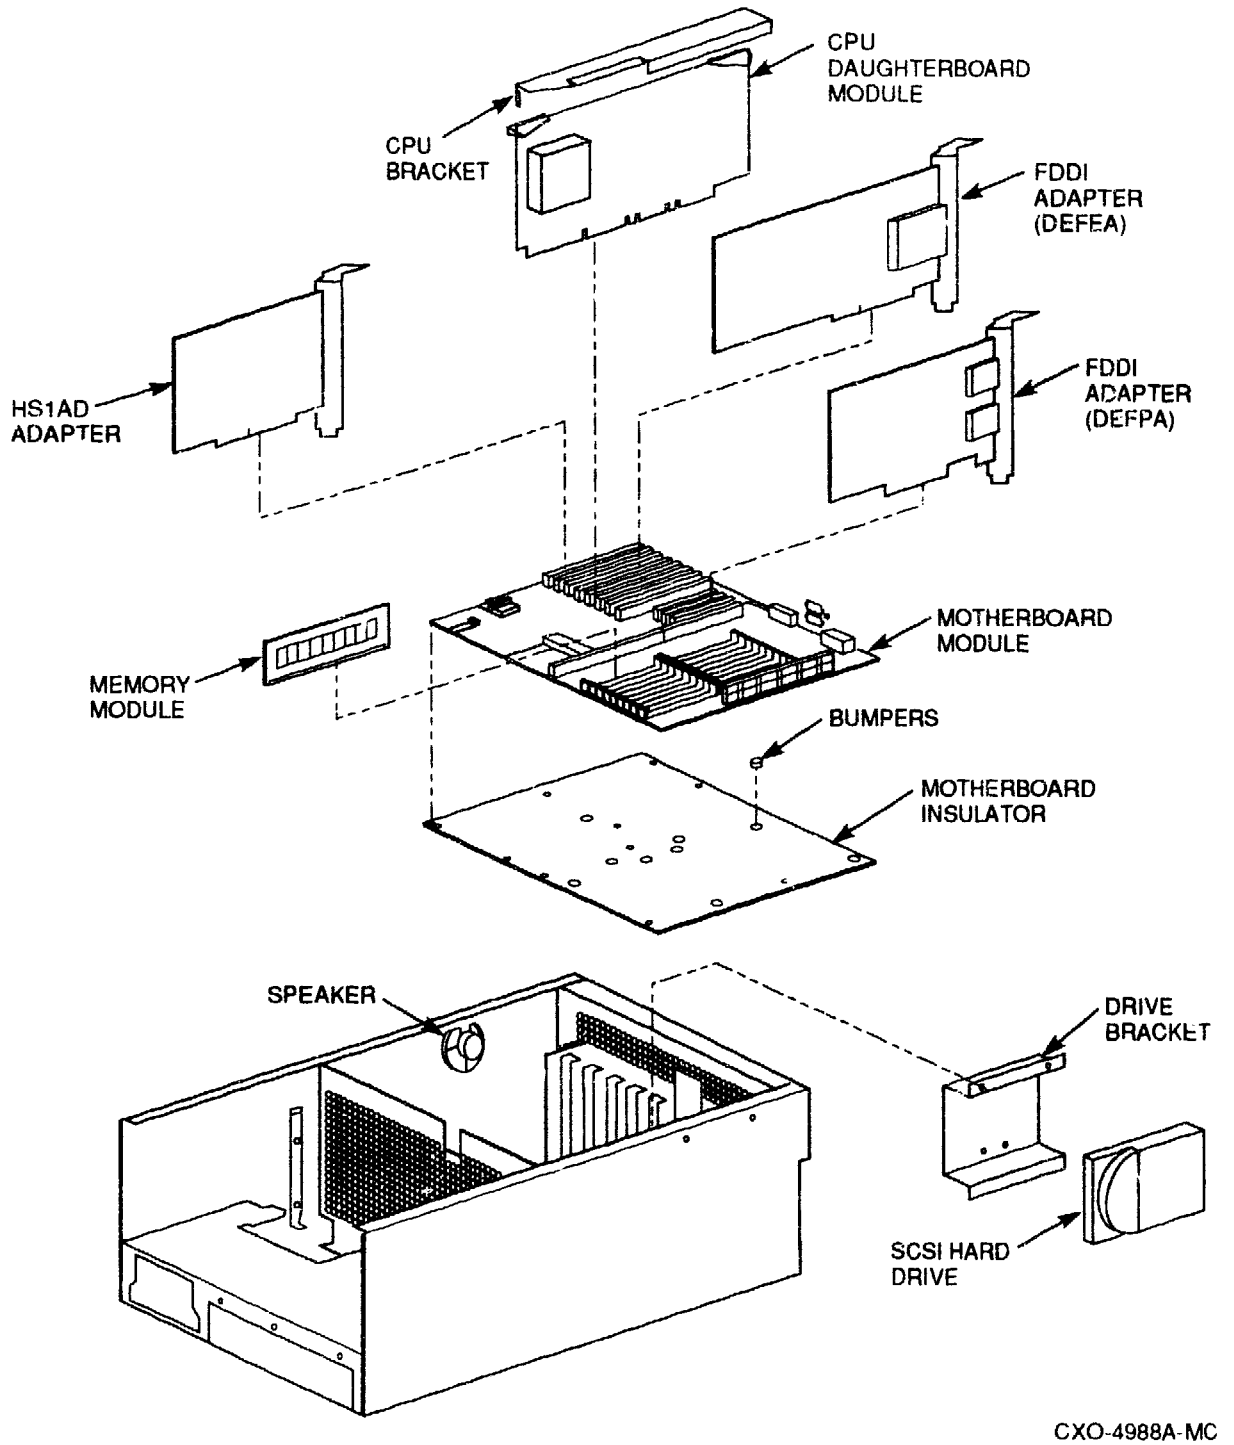
8.6 Replacing Internal Server Processor Components Accessible from the Rear

The following sections describe the procedures for removing and replacing internal server processor FRUs. Before removing and replacing any FRUs, ensure that the following have occurred:

- Data input to the server processor has been stopped.
- Power has been removed. (Refer to Section 8.5.2 for the system shutdown procedures.)

See Figure 8–4 for the layout of the server processor components accessible from the rear of the server processor.

Figure 8–4 Server Processor Components Accessible from the Rear



Note

Note that the motherboard insulator and bumpers shown in Figure 8–4 are bonded to the surface of the server processor chassis and usually are not replaceable items.

Note

Figure 8–4 shows both a DEFEA FDDI adapter (for HS1*xx* FDDI Servers) and a DEFPA FDDI adapter (for HS2*xx* FDDI Servers. Your server processor has one or the other, but not both.

8.6.1 Removing and Installing SIMMs

Removing and replacing Single In-line Memory Modules (SIMMs) includes the following tasks:

- Removing SIMMs
- Replacing SIMMs
- Verifying memory configuration
- Booting the system

The following sections describe these tasks.

8.6.1.1 Removing SIMMs

Note

If SIMMs are installed in Bank 3, you may need to disconnect and remove the hard drive to have enough space to remove or replace a full complement of SIMMs. Refer to 8.6.3.

WARNING

Memory and CPU modules have parts that operate at high temperatures. Wait 2 minutes after power is removed before handling these modules.

To prevent damage to SIMMs adjacent to the ones you want to remove, it is recommended that SIMMs be removed and installed only in successive order. For example, if you have SIMMs already installed in Banks 1, 2, and 3, and you want to remove a SIMM from Bank 0, follow this process:

- Remove SIMMs in Bank 3.
- Remove SIMMs in Bank 2.
- Remove SIMMs in Bank 1.
- Remove the SIMM you need to remove in Bank 0.
- Replace the SIMMs using the reverse order.

Refer to Figure 8–5 for clarification of the SIMM layout on the motherboard. Figure 8–6 illustrates how a SIMM is held and removed from the motherboard.

Use the following procedure to remove SIMMs:

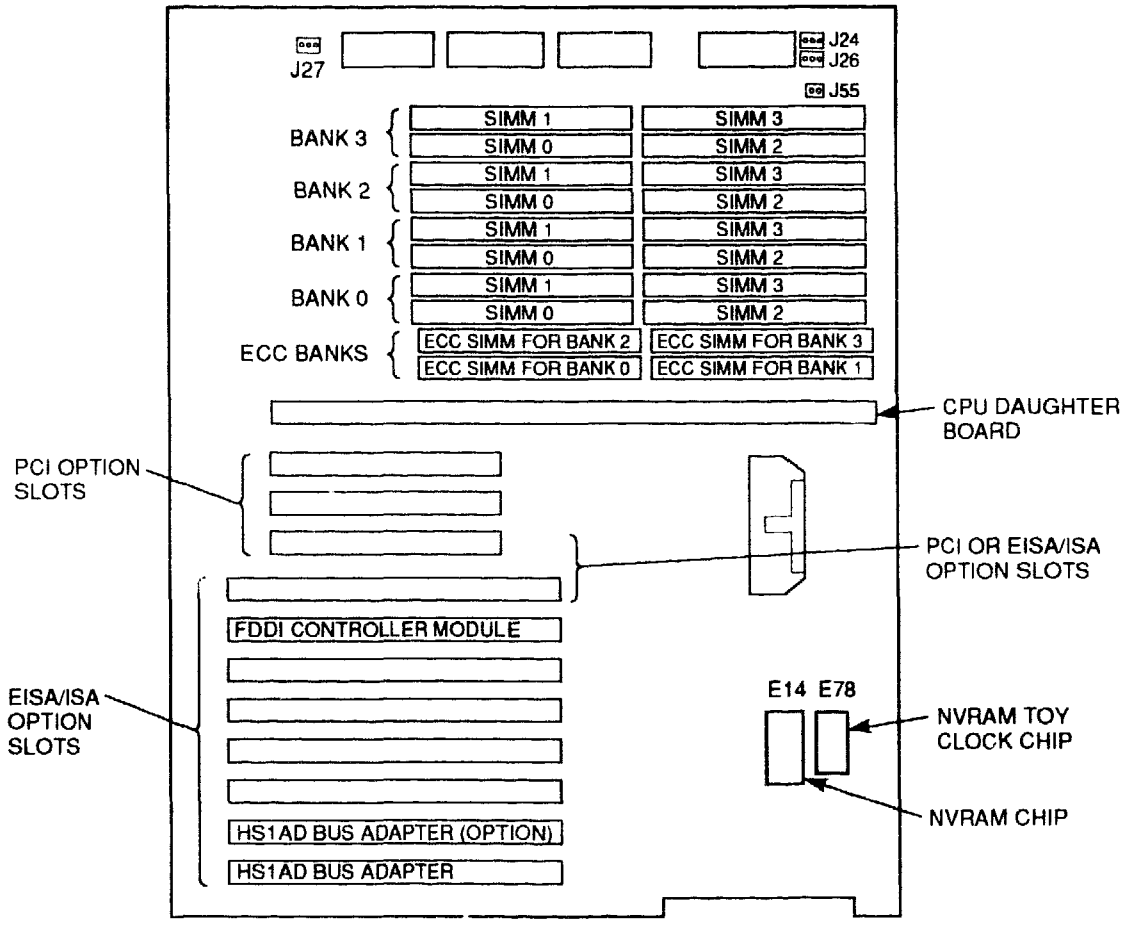
1. Remove a SIMM by pressing the metal clips on both sides of the SIMM connector to the side. (See callout 1 in Figure 8–6.)

Note

If you use a screwdriver to press back on the metal clips, take care not to damage the motherboard.

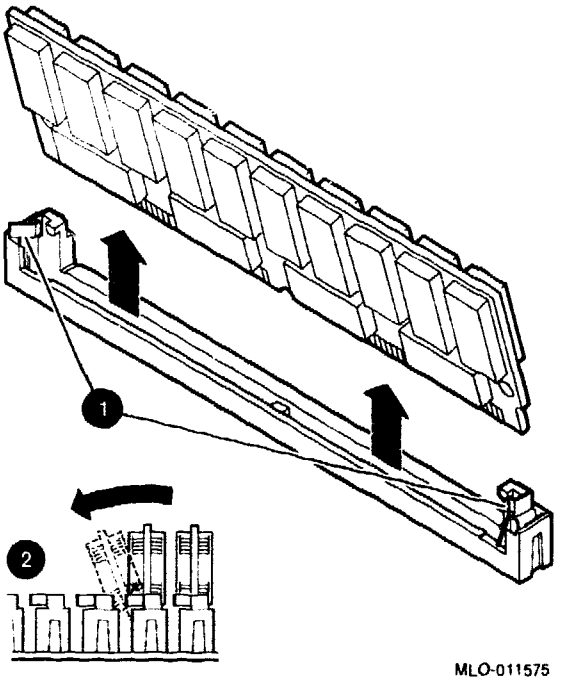
2. Note the position of each of the SIMMs you remove.
3. Tilt the memory module and lift it out of its connector. (See callout 2 in Figure 8–6.)

Figure 8–5 Server Processor Motherboard (for HS1*xx*)



CXO-4626A-MC

Figure 8–6 Removing SIMMs from the Motherboard



8.6.1.2 Preparing to Install SIMMs

Observe the following rules when replacing SIMMs:

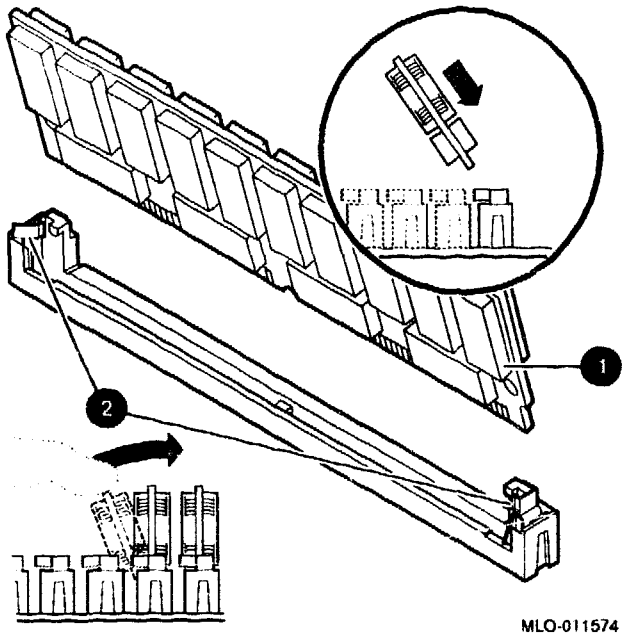
- Bank 0 must contain a memory option. A memory option consists of 5 SIMMs: 0, 1, 2, and 3, and 1 error correction code (ECC) SIMM.
- When you install memory in a memory bank, you must have a SIMM installed in all of the connectors in that bank.
- All SIMMs within a bank must be of the same memory capacity.
- Observe the rules and recommendations regarding the procedure for removing and installing SIMMs (see Section 8.6.1.1).

8.6.1.3 Replacing SIMMs

Use the following procedure for replacing SIMMs:

1. Tilt the connecting end of the module and press gently on the module so it slips over the two posts located at each end of the slot (see Figure 8–7).
2. Snap the SIMM into place in the correct upright position, as shown in Figure 8–7. Make sure that the SIMM is fully seated. The two latches on the ends of each SIMM connector should lock around the edges of the SIMM.
3. Install all SIMMs using this same procedure, including any you removed for SIMM access.
4. Visually inspect all SIMMs to insure that the spring clips are locked over the edge of each SIMM connector.
5. If the hard drive was removed for SIMM access, reinstall and reconnect the drive. (Refer to Section 8.6.3.)

Figure 8–7 Installing a SIMM



8.6.1.4 Verifying Memory Configuration

Before booting your system, you should verify the memory configuration using the SHOW MEMORY command, as follows:

```
>>> SHO MEM
```

The SHOW MEMORY command displays information for each memory module in the system. The command and a sample display are shown in Example 8–1.

Example 8–1 SHOW MEMORY Command

```
>>> SHO MEM
     128 Meg of System Memory
     Bank 0 = 64 Mbytes(16 MB Per SIMM) Starting at 0x00000000
     Bank 1 = 64 Mbytes(16 MB Per SIMM) Starting at 0x04000000
     Bank 2 = No Memory Detected
     Bank 3 = No Memory Detected

>>> CONT
```

8.6.1.5 Booting Your System

After replacing the SIMMs and verifying your memory, boot your system by entering the BOOT command at the >>> prompt.

8.6.2 Removing and Replacing the EISA or PCI Option Cards

Removing and replacing EISA and PCI option cards includes the following tasks:

- Removing an option card
- Replacing an option card
- Testing the option card installation
- Configuring the option card

8.6.2.1 Preparing to Remove an EISA Device Bus Adapter Card

Digital recommends that you should power everything off prior to servicing equipment such as a server processor device bus adapter card. However, if the FDDI Server must remain running while servicing one of the server processors in a dual server processor FDDI Server, then the following procedure provides a means to hot swap components in one of the server processors.

1. Detach the internal bus cable from the device bus adapter to be replaced as described in Steps 1-7 of Section 8.4.1.
2. Remove and replace the device bus adapter as described in Sections 8.6.2.2 and 8.6.2.3.
3. Reattach the internal bus cable to the replaced device bus adapter as described in Steps 6-10 and 13-17 of Section 8.4.3.
4. Test the new device bus adapter as described in Section 8.6.2.4.

8.6.2.2 Removing an EISA or PCI Option Card

CAUTION

Electrostatic discharge damages components. Always use proper ESD grounding procedures when handling components. Refer to Section 8.1 for proper grounding procedures.

To remove an EISA or PCI option card, refer to Figure 8–5 and perform the following steps:

1. Disconnect any cables connected to the external or internal ports on the option card you want to remove.
2. Remove the slot cover screw securing the option card to the chassis.

   If the option card to be replaced is in HS1AD slot 0 (between the side of the server processor and the option card in slot 1), you may need to remove the card in slot 1 first.
3. Carefully disconnect the option card from the slot connectors on the system board and remove it from the system.

Note

If you leave an option slot vacant, install a slot cover and secure it to the chassis using the screw that you removed from the option card.

8.6.2.3 Replacing an EISA or PCI Option Card

To replace an option card, refer to Figure 8–5 and perform the following steps:

1. Remove the new option card from its ESD Antistatic bag and carefully install it into the appropriate slot connectors on the system board. Press it firmly into place.
2. Secure the option card to the chassis using the screw you removed from the original card.

Note

If the top EISA option slot is used, the bottom PCI slot cannot be used. If the bottom PCI slot is used, the top EISA slot cannot be used.

8.6.2.4 Testing an EISA or PCI Option Card Installation

To test an option card, you must have a terminal connected to the server processor, all data cables must be connected, and power must be restored to the system. Refer to Section 8.3 for information on connecting a terminal to the system.

Use the following procedure to test the option card installation:

1. Enter the SHOW CONFIG command to display the system configuration. The system responds with a display similar to that shown in Example 8–2.
2. Examine the EISA or PCI bus information in the display to make sure that the new option is listed. If the option is not listed, see Chapter 4 for help.
3. Determine the device name of the new option. You must know the device name to run diagnostic tests on the option. Enter the SHOW DEVICE command to see the device name.
4. Run a self-test on the option by entering the TEST command followed by its device name. If the self-test passes, the system displays the word OK before displaying the console prompt >>>.

   If the system displays a failure message, see Chapter 4 for help.

The following example shows the information displayed when you enter the SHOW CONFIG command:

Example 8–2 The SHOW CONFIG Command

```
    Digital Equipment Corporation Alpha AXP (tm) Server

SRM Console x3.7-239 VMS  PALcode x5.48-49

MEMORY 128 MB of System Memory
   Bank 0 = 64 MB (16 MB per SIMM) starting at 0x00000000
   Bank 1 = 64 MB (16 MB per SIMM) starting at 0x20000000
   Bank 2 = No Memory Detected
   Bank 3 = No Memory Detected

AlphaServer 1000 Serial ROM version:    x0.6

EISA BUS
    Slot 3    DEC4220     era0.0.0.3.1  08-00-2B-35-F5- 24
    Slot 6     CPQ30110
```

8.6.2.5 Configuring an EISA or PCI Option Card

Depending on the type of option card you install, you may not need to configure it. When installing an EISA or PCI option card, refer to the system startup display screen to see if configuration of the card is required. After you turn on the system, the system startup sequence examines the EISA or PCI option slots and reports whether you need to run the EISA or PCI Configuration Utility.

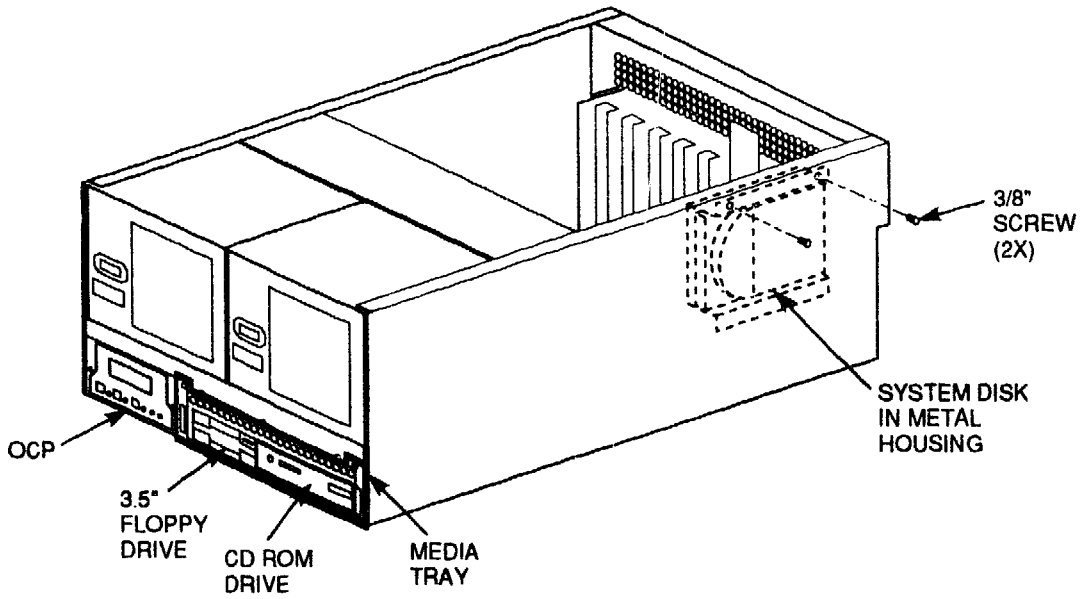
8.6.3 Removing and Replacing the SCSI Hard Drive

The system disk drive is the SCSI hard drive mounted on the sidewall of the server processor. When you replace a SCSI hard drive, you must assign it a unique SCSI ID number. Otherwise, the system does not recognize it. To remove the hard drive, refer to Figure 8–4 and Figure 8–8 and follow these steps:

CAUTION

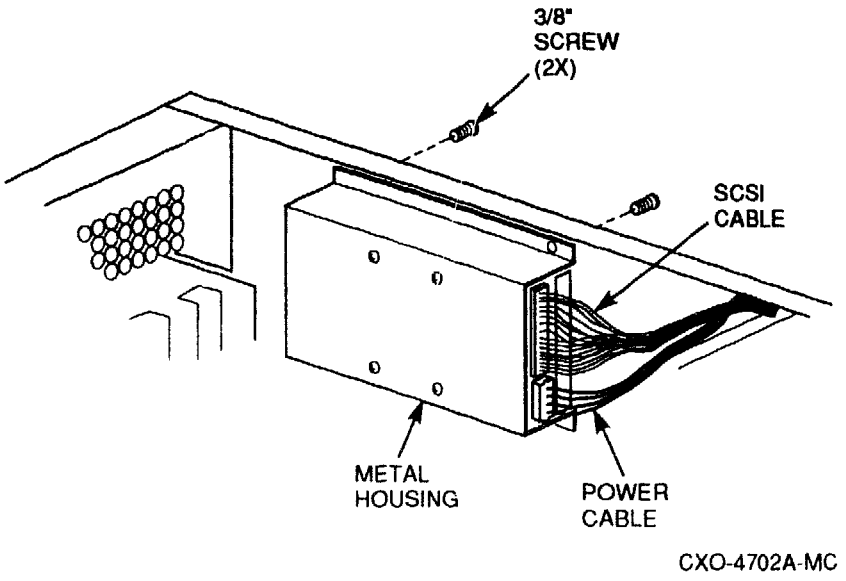
Electrostatic discharge damages components. Always use proper ESD grounding procedures when handling components. Refer to Section 8.1 for proper grounding procedures.

Figure 8–8 System Disk Location



1. Disconnect the SCSI cable and the power cable that are plugged into the back of the system disk (see Figure 8–9). Note the position of any jumpers or switches on the original device and set up the replacement devices using an identical configuration.

Figure 8–9 System Disk Cabling

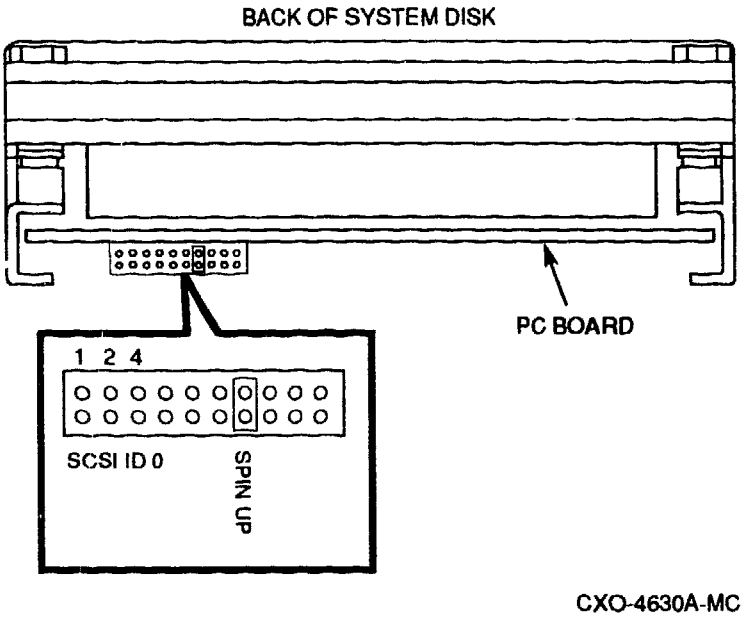


2. Using a Phillips screwdriver, remove the two screws on the side panel of the server processor (see Figure 8–8).
3. Carefully remove the system disk and its drive bracket from the inside of the server processor.

The drive bracket is the housing that holds the system disk.

4. Locate the option connector on the system disk (see Figure 8–10) and observe the jumper configuration. The option connector is a small, black 20 pin connector with one or two white jumpers installed.

Figure 8–10 System Disk Option Connector and SCSI ID 0



5. Set the jumper connections on the new disk drive to match the jumpers on the drive being removed.
6. Reinstall the system disk and drive bracket into the server processor by mounting it onto the side panel with the two Phillips screws that were previously removed.
7. Reattach both the SCSI cable and power cables to the system disk (see Figure 8–9).
8. If necessary, replace the nylon cable tie holding the power cord and data cables to the side of the server processor housing.

8.6.4 Removing and Replacing the Speaker

Use the following procedure to remove and replace the speaker:

1. Remove the rear top cover .
2. Locate the speaker on the side wall of the rear compartment.
3. Cut the tie wrap on the speaker wires.
4. Cut the tie wrap holding the speaker to the speaker bracket.
5. Remove the speaker.

Reverse the procedure to install the speaker.

8.6.5 Removing and Replacing the CPU Daughterboard

Use the following procedure to remove and replace the CPU daughterboard:

WARNING

CPU and memory modules have parts that operate at high temperatures. Wait 2 minutes after power is removed before handling these modules.

CAUTION

Electrostatic discharge damages components. Always use proper ESD grounding procedures when handling components. Refer to Section 8.1 for proper grounding procedures.

1. Remove the screw holding the CPU shipping bracket from the outside rear face.
2. Lift the CPU shipping bracket off by pushing forward towards the fans to clear the lip, lifting and turning to the side.
3. Remove the CPU daughterboard using the pry handles on each end of the board.
4. Install the new CPU daughterboard in the vacant slot, pressing firmly to make sure the board is properly seated.
5. Install the CPU shipping bracket using the screw that was previously removed.

Refer to Figure 8–4 for the relative location of motherboard components.

8.6.6 Removing and Replacing the FDDI Adapter

Refer to Figure 8–4 for a view of the components mounted on the motherboard.

Use the following procedure to remove and replace the FDDI adapter:

1. Remove the slot cover screw securing the FDDI adapter to the chassis.
2. Remove the FDDI adapter from the motherboard.
3. Install the new FDDI adapter, pressing firmly to make sure the board is properly seated.
4. Install slot cover screw that was previously removed.

8.6.7 Removing and Replacing the Motherboard

Use the following procedure to remove and replace the motherboard:

1. Remove data cables from EISA option and FDDI adapter.
2. Record the position of the EISA option.
3. Remove the EISA option (see Section 8.6.2).
4. Remove the CPU daughterboard (see Section 8.6.5).
5. Remove the FDDI adapter (see Section 8.6.6).
6. Remove the SIMMs (see Section 8.6.1).

7. Disconnect the speaker wire.
8. Disconnect each of the cable bundles from the motherboard.

Note

You may need to remove the cable ties holding various cables before removing the motherboard.

9. Remove the hard disk drive (see Section 8.6.3).
10. Remove the 14 Phillips screws holding the motherboard to the server processor.
11. Tilt and lift the motherboard out of the enclosure.

    Prior to reinstalling the motherboard, visually inspect the server processor case to insure that all foreign objects have been removed from the interior of the case.
12. Check that all jumpers on the new motherboard are configured the same as the old motherboard. See Table 8–2 for a description of the jumpers on the motherboard. See Figure 8–11 for the location of the jumpers on the motherboard.

Figure 8–11 Motherboard Jumper Locations

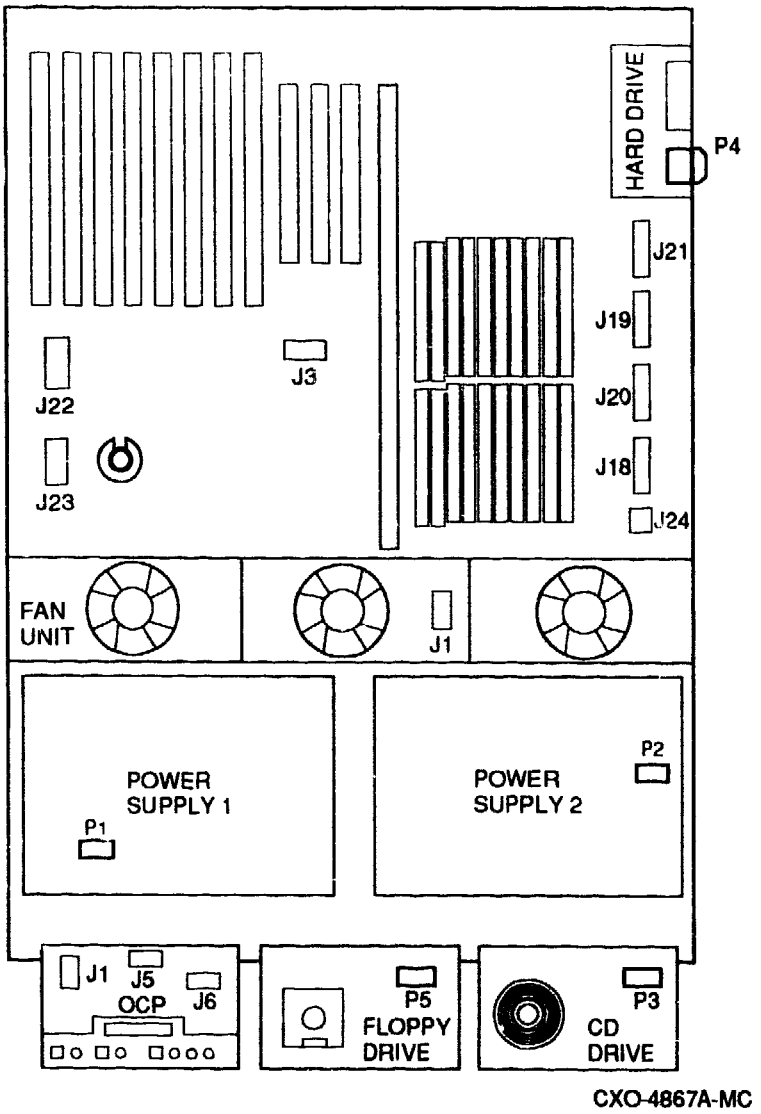


Table 8–2 Jumpers on the Motherboard

Jumper	Name	Description	Default Setting
J27	VGA Enable	When enabled (as shown in Figure 8–11), the on-board VGA logic is activated.	Enabled for on-board VGA; Disabled if an EISA- or PCI-based VGA option is installed.
J49	SCSI Termination	Allows the internal SCSI terminator to be disabled.	Enabled (as shown in Figure 8–11).
J50	Flash ROM VPP Enable	Permits the 12V voltage needed to update the Flash ROMs.	Jumper installed.
J52	Temperature Shutdown	Allows the temperature sensor to shut down the system in an over temperature condition.	Currently ships enabled (as shown in Figure 8–11).
J53	Fan Shutdown	Allows the software to shut down the system if a fan fails.	Enabled (as shown in Figure 8–11).
J55	Small Fan	Allows the small fan to be disabled to accommodate the rackmount enclosure.	This jumper is installed on server processor systems.
J56	Fan Fault	When enabled, the hardware forces the system to shut down if a fan fails. When disabled, the firmware generates a machine check; the system crashes and shuts down if a fan fails.	Enabled (as shown in Figure 8–11).

To install the new motherboard, reverse the step in the previous procedure.

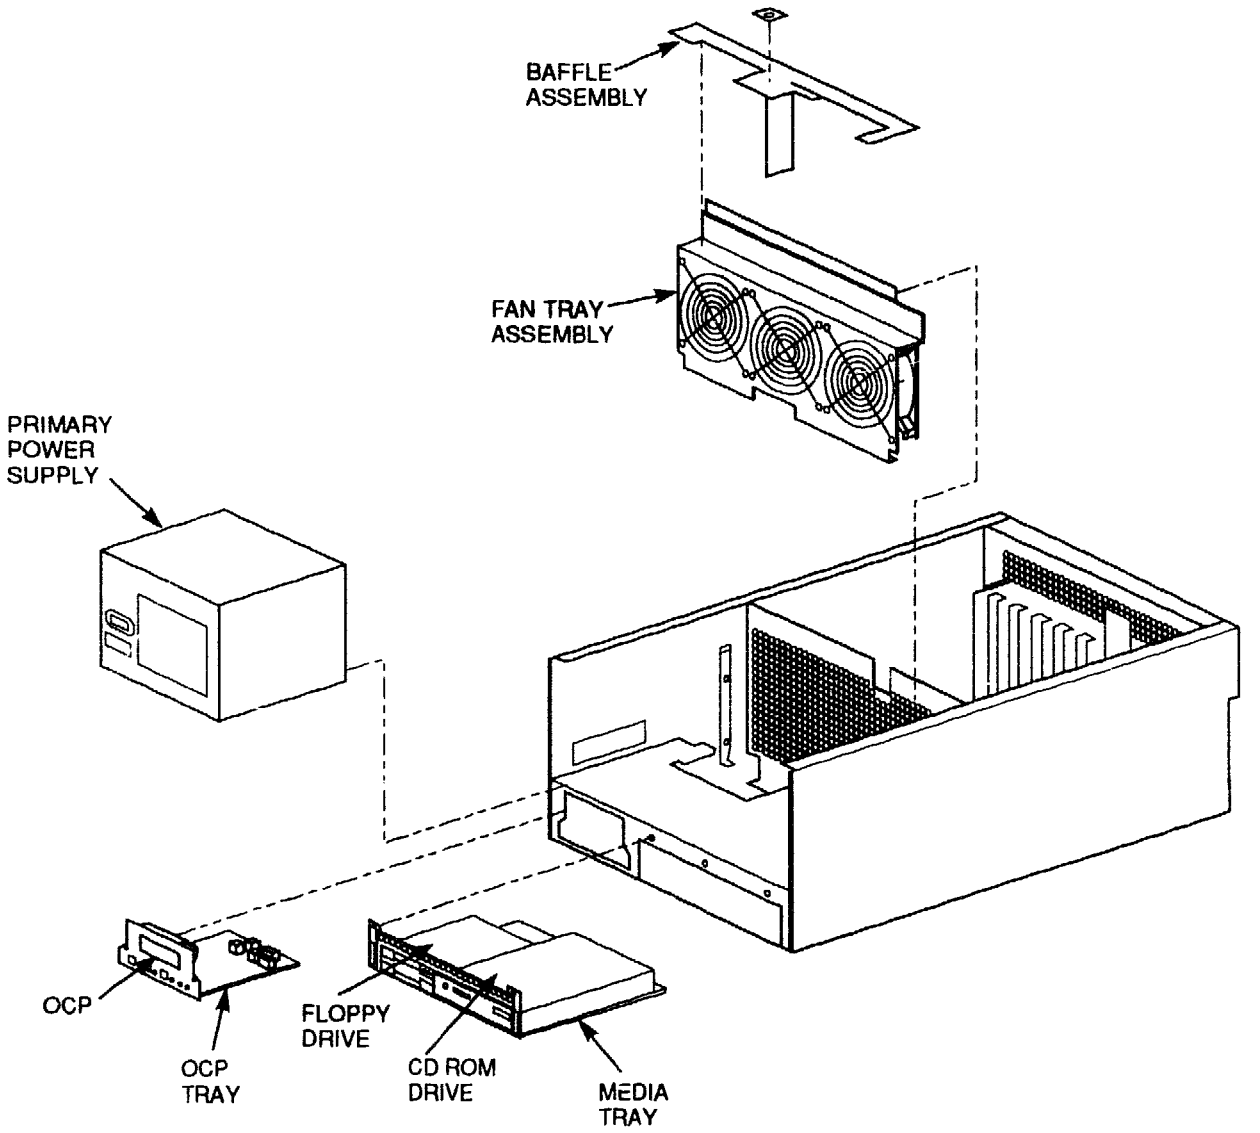
8.7 Replacing Internal Server Processor Components Accessible from the Front

The following sections describe the procedures for removing and replacing internal server processor FRUs. Before removing and replacing any FRUs, ensure that the following have occurred:

- Data input to the server processor has been stopped.
- Power has been removed. (Refer to Section 8.5.2 for the system shutdown procedures.)

See Figure 8–12 for the layout of the server processor components accessible from the front of the server processor.

Figure 8–12 Server Processor Components Accessible from the Front



8.7.1 Removing and Replacing the Fan

Use the following procedure to remove and replace the fans:

1. Remove the top front and rear top covers (see Section 8.5.4).
2. Disconnect all cables from the power supplies.
3. Remove the power supplies for clearance.
4. Move all cables out of the way by clipping the cable clamps on the fan assembly air baffle.
5. Disconnect the fan harness from the fan assembly.
6. Remove the two screws holding the air baffle to the fan assembly and remove the baffle.
7. Remove the three screws holding the fan assembly.
8. Pull the fan assembly straight up and out to clear the unit.

Reinstall the fan assembly by reversing the proceeding procedures.

8.7.2 Accessing the Media Tray

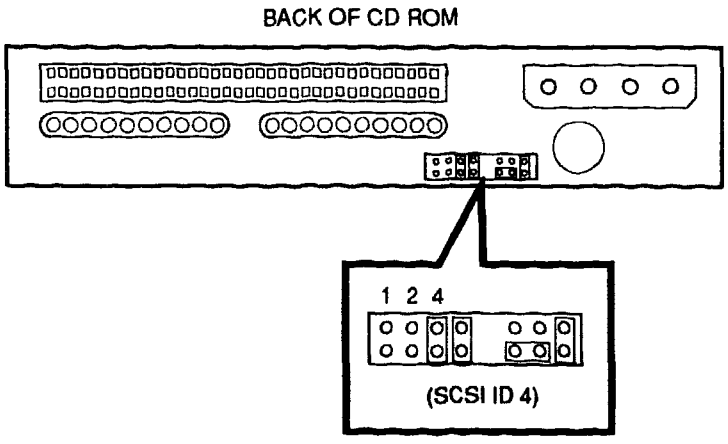
The media tray houses the CD–ROM drive and the floppy drive. You must remove the media tray before you can remove either the CD–ROM drive or the floppy drive. To remove the media tray, remove the two screws on the front of the media tray using a Phillips screwdriver and slide the media tray forward to access the floppy and CD–ROM drives.

8.7.2.1 Replacing a CD–ROM Drive

To replace a CD–ROM drive, follow these steps:

1. Disconnect the SCSI connector and power connector from the back of the CD–ROM drive (see Figure 8–13).
2. Remove the four screws holding the drive.
3. Replace the old drive with the new drive.
4. Make sure that the RRD43 terminators are installed and the jumpers on the new drive are set the same as on the drive being replaced.

Figure 8–13 Replacing a CD–ROM Drive



Note

The default SCSI ID for the CD–ROM is ID 5. For the HS121 configuration, a second FDDI Server CD–ROM needs to be set for SCSI ID 4. (Remove jumper ID 1.)

5. Slide the drive into the front of the media tray until its screw holes are aligned with the drive bracket holes and install the four mounting screws.
6. Connect the power cable and the data cable to the back of the drive.
7. Reinstall the media tray.

8.7.2.2 Removing and Replacing a Floppy Drive

Follow these steps to replace the 3.5-inch floppy drive:

1. With the media tray removed (see Section 8.7.2), disconnect the IDE and power cables from the back of the floppy drive.
2. Remove the screws attaching the floppy drive to the media tray.
3. Remove the two screws holding the floppy drive to its mounting bracket.
4. Slide the floppy drive into the front of the media tray until the screw holes line up with the holes on the bracket.
5. Attach the mounting bracket to the replacement floppy drive using the two screws that you removed from the original floppy drive.

Note

Make sure the key on the IDE cable connector faces up toward the top of the disk drive.

6. Connect the power cable and the data cable to the floppy drive.
7. Replace and secure the media tray into the server processor.

8.7.3 Removing and Replacing the OCP Module

Use the following procedure to remove and replace the OCP:

1. Remove the two screws on the front.
2. Remove the cables.

Note

Perform the next step only if you want to remove the LCD panel.

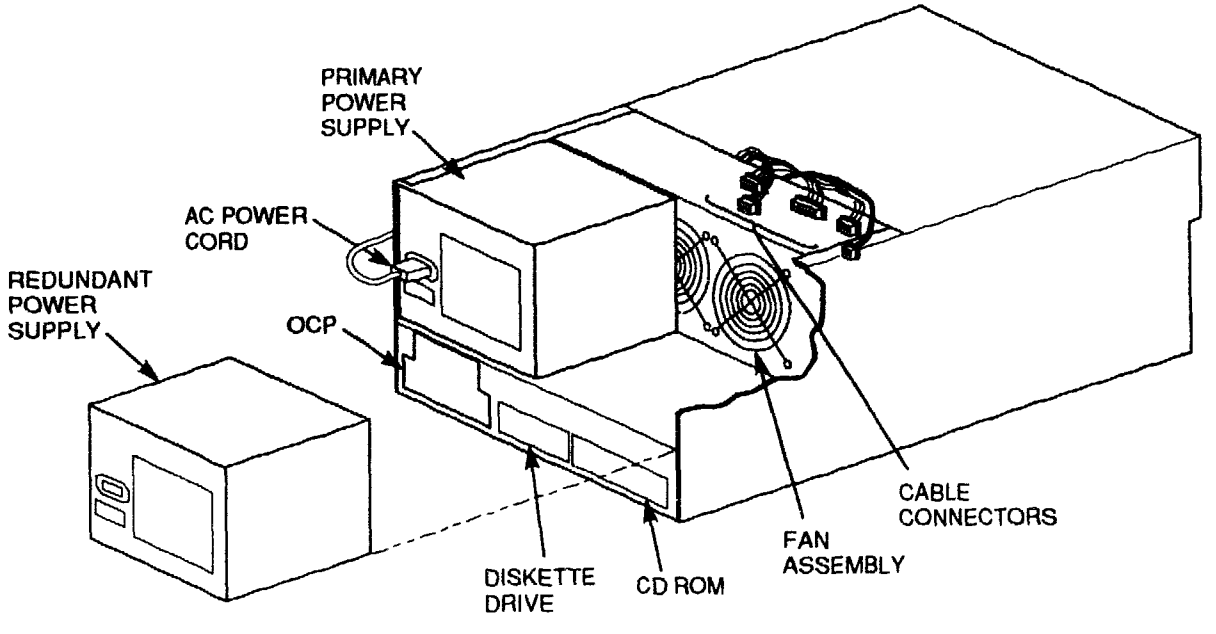
3. Remove the four screws, nuts, and spacers for the LCD panel.
4. Remove the OCP module.

Reverse the procedure to install the new OCP.

8.7.4 Removing and Replacing the Power Supply

The following procedures describe how to remove and replace either the primary or redundant power supplies. The primary power supply resides in the left-most position; the redundant or secondary power supply resides in the right-most position (see Figure 8–14). The procedures for installing and removing both power supplies are similar, unless otherwise noted.

Figure 8–14 Power Supply Locations



CXO-4627A-MC

Removing the Power Supply

To remove a power supply, follow these steps:

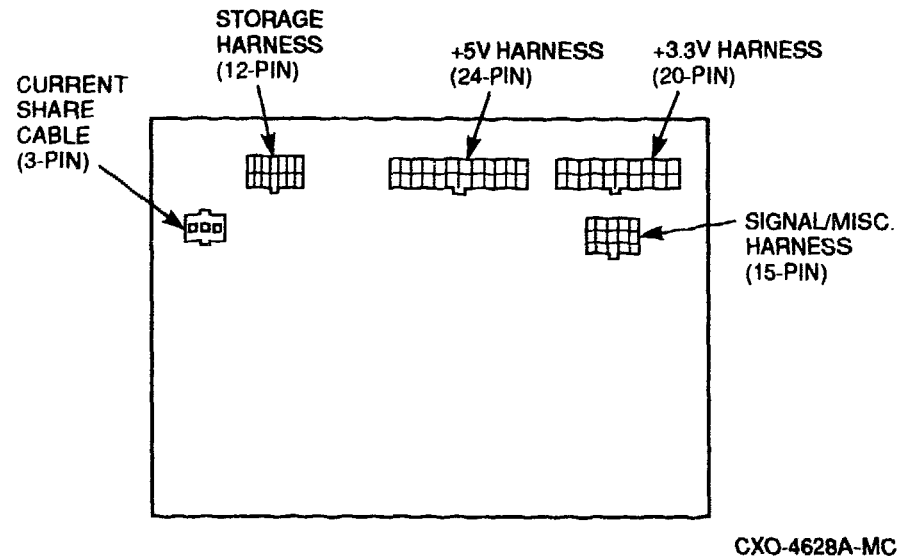
1. Remove the front cover from the server processor using the procedure described in Section 8.5.4.

Note

You may need to remove the rear top cover to access the power supply cables.

2. Disconnect all external power supply cables from one or both power supplies.
3. Slide the (right-most) power supply toward the front and lift out.
4. Repeat steps 1 through 3 on the primary power supply if a redundant power supply was removed to provide access to the primary power supply.

Figure 8–15 Power Supply Cable Connections



Replacing the Power Supply

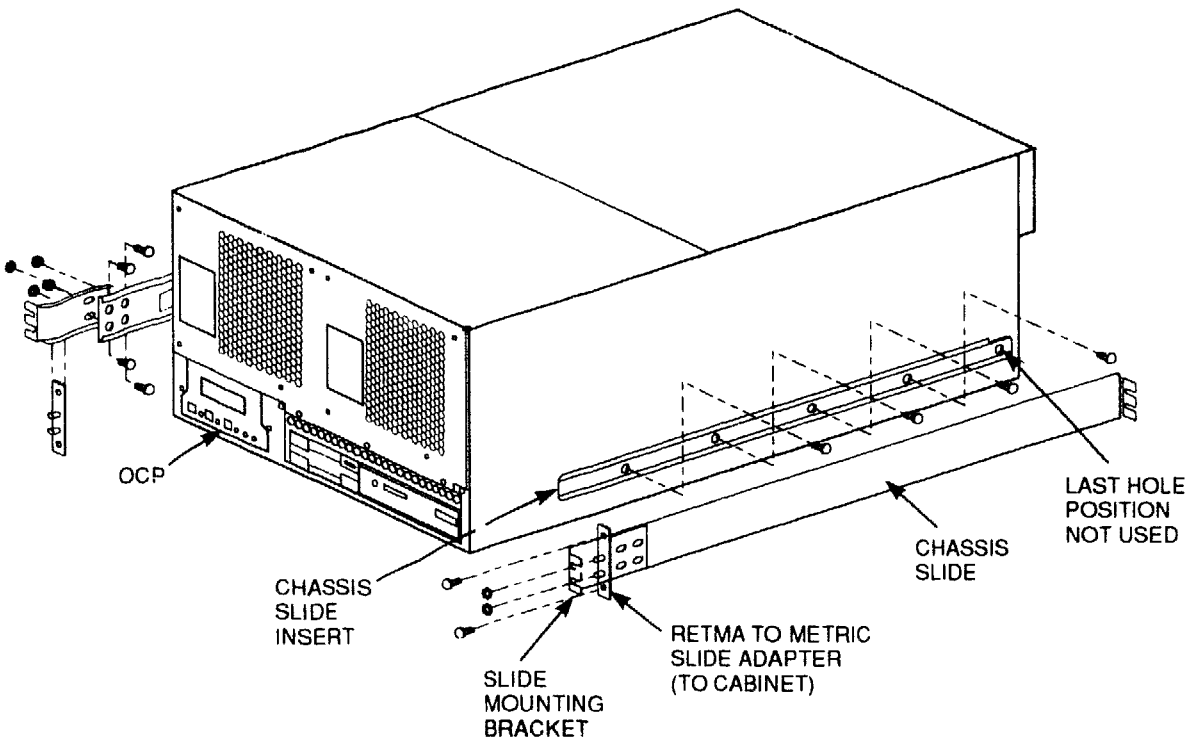
To replace or install a power supply, follow these steps:

1. Slide the replacement power supply into the chassis of the server processor (see Figure 8–14).
2. Connect all external cables to the connectors on the back of the power supply. Connectors are shown in Figure 8–15.
3. Replace the front cover.
4. Replace the screws that secure the power supply to the server processor.
5. Reattach the front cover to the chassis and to the installed power supplies using the screws previously removed.

8.8 Reassembling the Server Processor

If the server processor needs to be replaced, the chassis slide inserts need to be removed from the old server processor and be installed on the new server processor. Use the following procedure to change over the slide inserts:

Figure 8–16 Slide Insert Removal/Installation



CXO-4618A-MC

1. Remove the slide inserts from the old server processor by loosening and removing the four screws holding each insert.
2. Install one slide insert on each side of the new server processor as shown in Figure 8–16.
3. Position the slide insert so that the metal locking tab is outside and the *notch* on each slide insert is toward the rear of the server processor (see Figure 8–16). Align the screw holes on the server processor with the holes to the slide insert.

Note

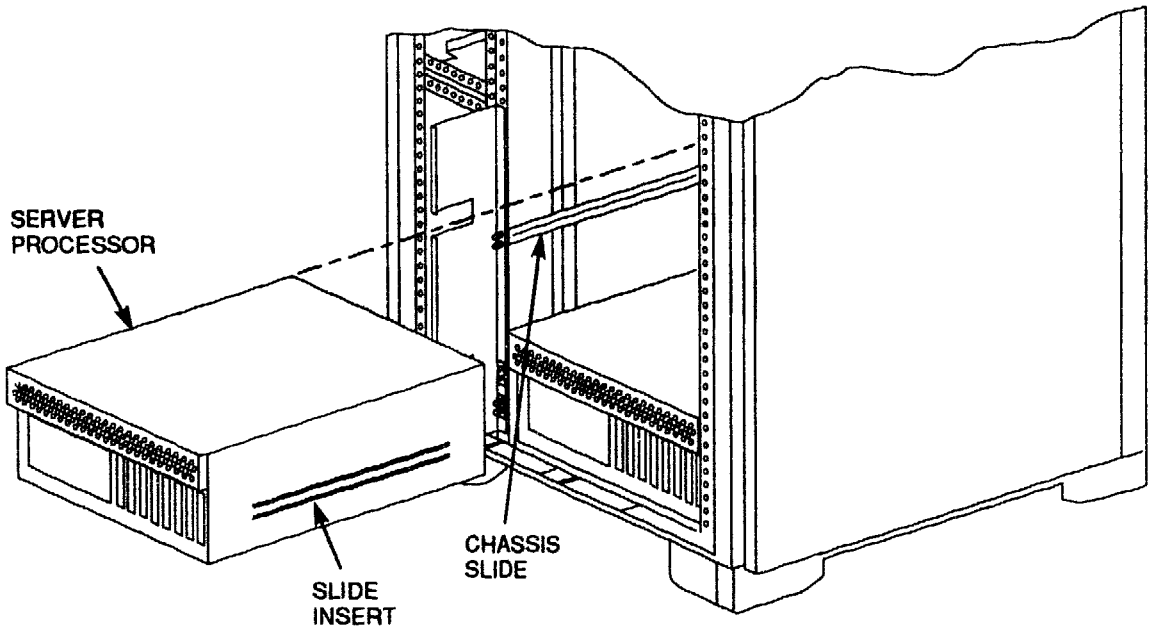
The last hole in the slide insert is not used.

4. Secure the slide insert to the server processor using 10–32 screws. To insure proper alignment, tighten the chassis slide holding screws after the server processor is reinstalled into the SW800 cabinet.

8.9 Reinstalling the Server Processor

Use the following procedure to reinstall the server processor into the SW800 cabinet:

Figure 8–17 Server Processor Installation

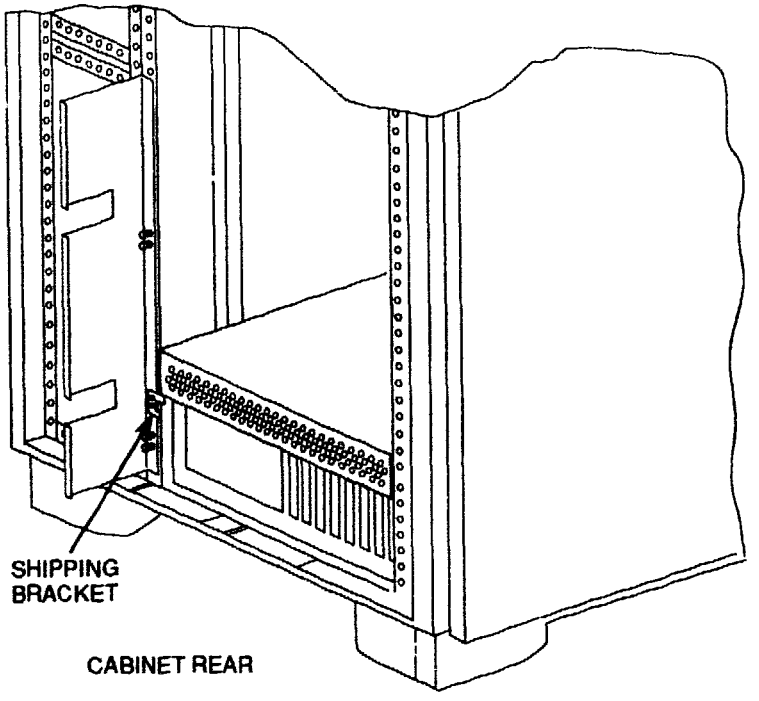


1. Reinstall front and top rear cover panels on the server processor.
2. Tighten all screws and quarter-turn fasteners on each cover panel.
3. With the help of a second person, install the server processor into the chassis slide from the *rear* end of the cabinet. Align the slide insert on the server processor to the chassis slides in the cabinet and slide the server processor into the cabinet until the insert locks on to the chassis slides (see Figure 8–17).
4. Slide the server processor back into the cabinet.
5. Reinstall the left and right shipping brackets for the upper (hole positions 42 and 43) or lower (hole positions 53 and 54) on both *rear* vertical rails (see Figure 8–18).

Note

The flat side of each shipping bracket must face out when installed. These shipping brackets are used to keep the server processor from sliding during shipping or when the cabinet is moved.

Figure 8–18 Shipping Brackets



6. Tighten all screws.

Note

When installed, the server processor should be able to slide easily in both directions.

7. Connect the power cord to the front of the server processor and route it neatly back to the CDU.

Use the cable clips to hold the power cord over the server processor and hold the cord. Make a service loop to ensure the power cords do not rub against the rails when the server processor is slid in and out of the cabinet. Secure the power cords to the vertical rail near the CDU with tie wraps.

8.10 Replacing the Cable Distribution Unit

The following procedure describes how to remove and replace the CDU.

Note

The first step may not be necessary if the cabinet is configured with the dual power option. The system can be left turned on while the defective CDU is replaced.

1. If the CDU to be replaced is installed in a single power system, spin down all disk drives and halt all tape drives in the cabinet. Refer to 8.5.2 for more detailed information on shutting down the system.

WARNING

Hazardous voltage is present within the cabinet. Use extreme caution when working around components to which power is applied.

2. Open the cabinet door to expose the CDU to be replaced.
3. Switch the circuit breaker on the front panel of the CDU to the ○ (OFF) position.
4. Disconnect the CDU's primary power cable from the power source.
5. While supporting the CDU, remove and retain the screws holding the CDU to the cabinet chassis rails.

WARNING

Cabinet rail edges may be sharp and can slice or abrade skin and cable insulation.

6. Gently pull the CDU from the cabinet while carefully guiding the cables attached to it through the cabinet openings.
7. Support the CDU while disconnecting the ac power cables from its rear panel.
8. Remove the CDU completely from the cabinet while guiding the power cable through its cabinet openings.
9. Guide the power cable of the replacement CDU back through the same cabinet openings and position the CDU near the previously removed power cables.
10. Connect the power cables to the connectors on the rear panel of the CDU, starting at the top-most connector and working downward.
11. Slide the CDU into its mounting position while guiding the cables through the openings in the cabinet.
12. Mount the CDU using the screws retained in step 5.
13. Connect the CDU's power cable to the appropriate power source.
14. Switch the front-panel circuit breaker on the CDU to the | (ON) position.
15. Close and secure the cabinet door.
16. Restart all disk and tape drives and check them for proper function.

8.11 Replacing FDDI Server Shelf Components

This section provides information on how to change FDDI Server shelves and shelf components. The following topics are covered:

- Warm swap procedure for storage devices
- BA350–M shelf replacement
- BA350–S shelf replacement
- Shelf blower replacement
- Shelf power supply replacement

8.11.1 Warm Swap Procedure

Warm swap is used to remove, replace, or add devices while the system remains on line.

The operator must quiesce SCSI bus operations during a device warm swap. Quiesce means to make a bus inactive or dormant. All activity on the bus where the device is being swapped must be halted for the duration of the removal or replacement.

CAUTION

Warm swap includes removal *and* replacement of *one device at a time*. Should another disk need to be removed, the entire warm swap procedure must be repeated. The warm swap procedure must be followed to protect data integrity for other shelf devices. Only one port can be quiesced at any time.

You also use the warm swap procedure to add a device to an empty shelf slot.

Warm swap is not applicable to service on unpowered StorageWorks shelves. Do not attempt to execute warm swap on an unpowered shelf.

The following sections identify the steps necessary for device removal and replacement.

8.11.1.1 Disk Drive SBB Warm Swap Removal

CAUTION

You must dismount the devices using your operating system's DISMOUNT command *before* beginning this procedure. Perform this warm swap procedure *exactly* as stated, or the HS1CP can perform unpredictably.

Refer to your operating system documentation for procedures necessary for dismounting a device.

Use the following procedure to remove disk drive SBBs safely:

1. Unlock and open the cabinet doors using a 5/32-inch Allen wrench.

2. Press and hold the port button for the disk SBB you wish to remove. Continue holding the button in until all amber channel device processor LEDs light.

Note

Only one port can be quiesced at at time.

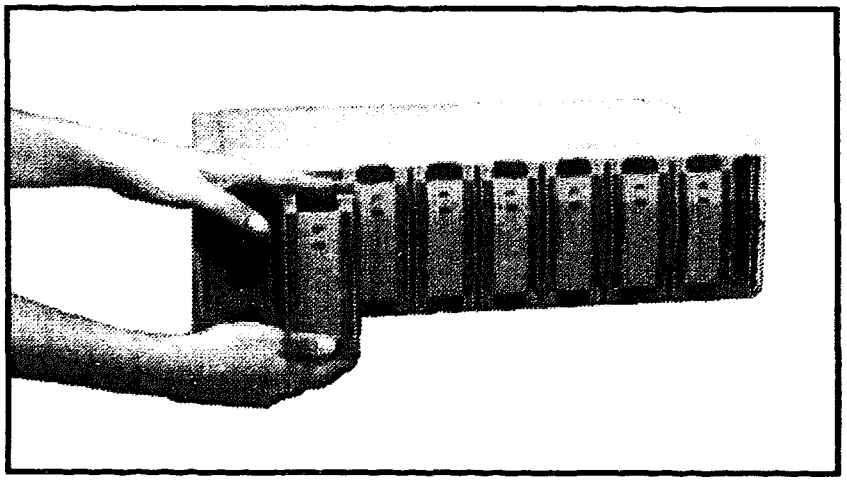
If the button is not held in long enough, or multiple buttons are pushed in quick succession, *all* button pushes are ignored and the port is not quiesced. You must press the button again to quiesce the port.

3. Wait until the chosen port LED flashes alternately with the other port LEDs (this indicates I/O has stopped). This alternating pattern flashes for approximately 30 seconds, during which time you can remove the SBB.

   If the pattern does not appear after a minute or two, another shelf is asserting a fault signal that prevents any quiesce function of this HS1CP. For tips to correct the problem, refer to Section 8.11.1.4.

4. Remove the disk SBB from its device shelf by squeezing both plastic tabs at the top and bottom of the SBB towards the center of the SBB. Slide the SBB out of the shelf using both hands, as shown in Figure 8–19.

Figure 8–19 SBB Warm Swap



CXO-3611B-PH

While the HS1CP LEDs are flashing their alternating pattern, the shelf SBB LEDs for all disk SBBs on that port also flash. (Although tape drives do not always have LEDs, remove them at this time if you are using warm swap to remove a tape drive.)

All SBB LEDs flash three times per second.

Note

The time to stop the I/O can vary from zero seconds to several minutes, depending on the load, device type, and cache status.

Once the disk SBB is removed, the flashing pattern on the HS1CP LED stops and normal operation on the ports resumes.

8.11.1.2 Disk Drive SBB Warm Swap Replacement

Use a replacement device of the same type as the removed device. Otherwise, subsystem failures such as the establishment of storage sets can occur. Disk drive SBBs can be replaced (inserted) safely using the following procedure:

1. Quiesce the SBB's port by pressing and holding the port button for the disk drive you wish to replace. Continue holding the button in until all amber HS1CP LEDs light.

   Note

   If the amber LEDs do not light to indicate the bus is quiesced, there could be a problem with a power supply in one of the shelves in the cabinet. Refer to Section 8.11.1.4 for tips on resolving this problem. If more information is needed, refer to the *StorageWorks Array Controllers HS Family of Array Controllers Service Manual* for troubleshooting procedures if this problem should occur.

2. Wait until the chosen port LED flashes alternately with the other port LEDs (this indicates I/O has stopped). This alternating pattern flashes for approximately 30 seconds, during which time you can insert the new SBB. While the OCP LEDs are flashing, any SBBs on the quiesced port that have status LEDs also flash.

   Note

   The time to stop the I/O can vary from zero seconds to several minutes, depending on the load and the type or device.

3. Replace (insert) the disk drive SBB using both hands to push it into the shelf until you hear the mounting tabs snap into place.

   After the SBB insertion is completed, the flashing pattern on the HS1CP LEDs stops and normal operation of the ports resumes. At this time, the port LEDs turn off.

4. If a *new* device is added in a previously *unused* slot, that port's LED remains on until the device is added by entering the following command:

   ```
   CLI> ADD DEVICE
   ```

5. If a tape drive SBB is placed in a slot where a disk drive SBB was installed previously, the port LED remains on until the tape device is added using the CLI ADD TAPE command. Delete the previously installed disk drive SBB from the list of CLI known devices using the CLI DELETE *device-name* command before you use the ADD TAPE command.

6. If the new disk drive SBB is to be part of a storageset, you must delete the storageset from the configuration and create (ADD) it again.

7. Initialize the inserted device by entering the following CLI command:

   ```
   CLI> INITIALIZE container-name
   ```

where *container* is either the disk drive SBB or group of disks linked as a storageset.

This initializes the metadata on each disk in the container, including the one that was just swapped.

CAUTION

If you failed to perform the warm swap exactly as stated, you must restart the HS1CP. Otherwise, the HS1CP can perform unpredictably.

8. Close and lock the cabinet doors.

8.11.1.3 Tape Drive SBB Warm Swap Removal and Replacement

The same procedure used to warm swap disk drive SBBs also is used for tape drive SBBs. Refer to Section 8.11.1.1 and Section 8.11.1.2.

8.11.1.4 When the Quiesce Bus State Is Not Displayed

When a power supply fails in a shelf with only one power supply, the shelf becomes inactive. A shelf with a failed power supply asserts a status signal, SWAP_L, to the device channel processor. This signal tells the HS1CP that the physical configuration of the shelf has changed during a normal warm swap operation.

The interrupt level coming into the HS1CP is the same as that of the HS1CP buttons. The shelf constantly asserts this interrupt to the HS1CP. Usually, the interrupt goes away rapidly during the warm swap operation. Having an interrupt remain active degrades the operation of the HS1CP. As a result, the HS1CP disables this interrupt. Disabling the interrupt means that the HS1CP does not respond to HS1CP button pushes, so the port cannot be quiesced.

If the bus quiesce LED pattern is not displayed after pushing the appropriate port button (as described in the warm swap procedure) within a minute or two, a failed power supply on another shelf within the cabinet could be causing the problem. This problem must be corrected before restarting the warm swap procedure.

To correct this problem, do one of the following:

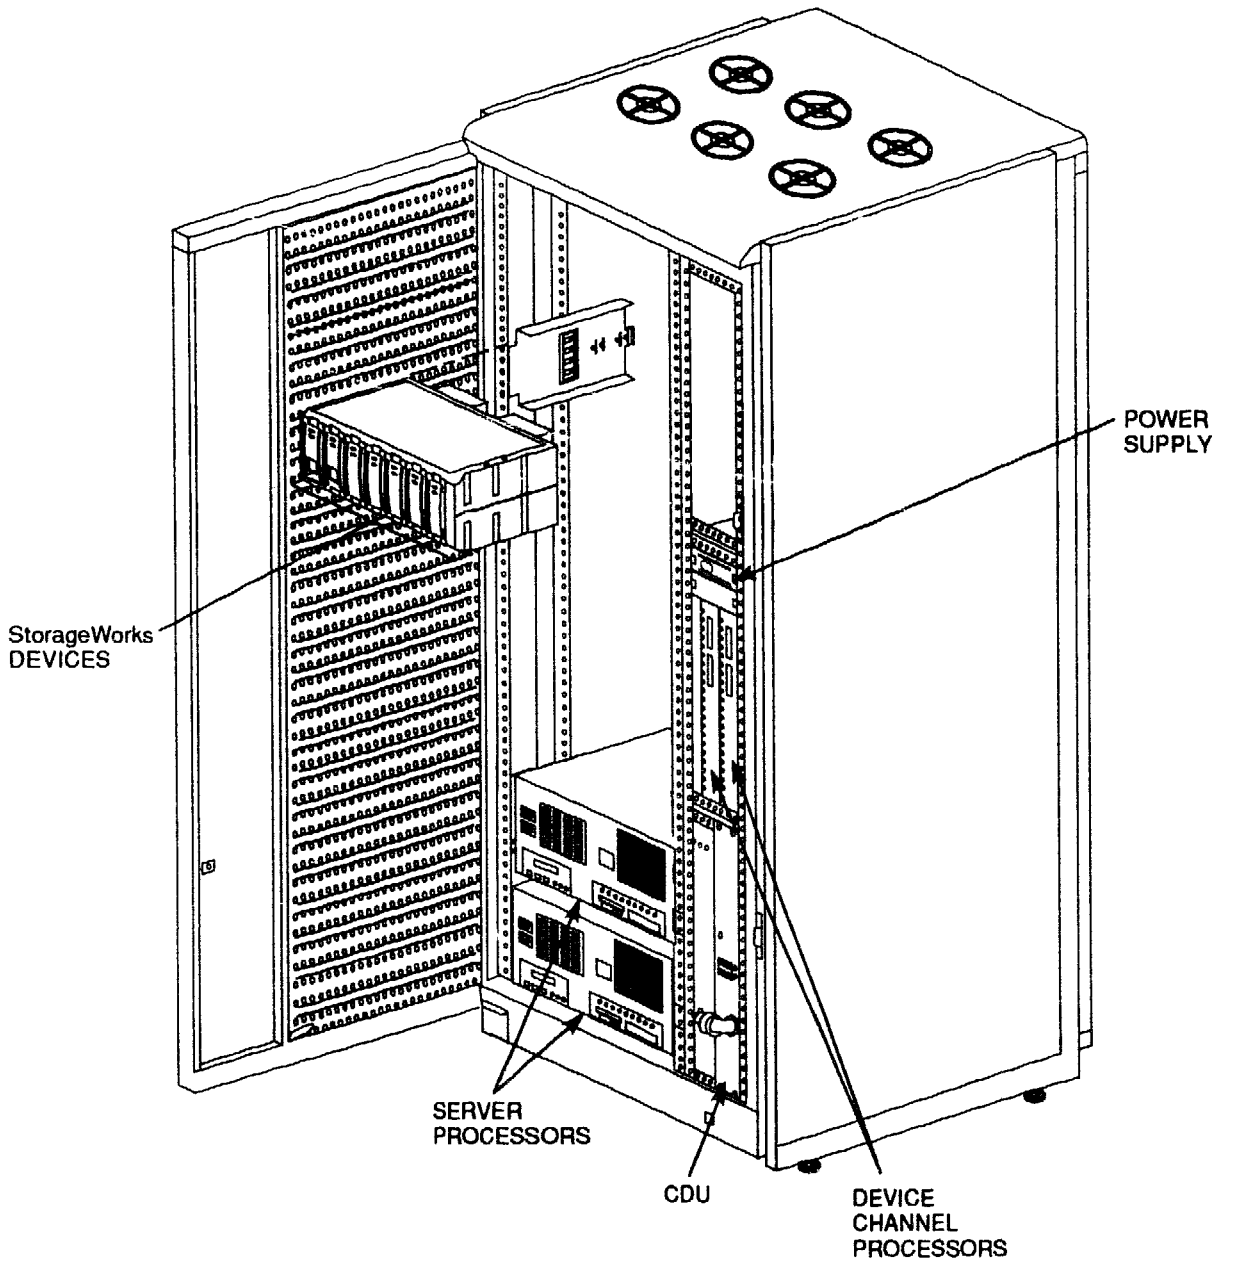
- Replace the power supply in the failed shelves.
- Remove all devices in the failed shelves.
- Unplug the SCSI cable within the failed shelves.

After correcting the problem, wait 2 to 3 minutes for the HS1CP to detect the removal of the SWAP_L signal and reenable the interrupt. You then can begin using the warm swap procedure as documented in Sections 8.11.1.1 and 8.11.1.2.

Note

The bus quiesce procedure for warm swap is NOT necessary when removing or adding devices to an unpowered shelf.

Figure 8–20 FDDI Server Component Locations



8.11.2 Removing and Replacing the BA350–M Shelf

All BA350–M shelves are inserted into a set of mounting brackets and secured in place with a front locking bracket. Use the following procedure to remove a shelf:

1. Turn off the power to the shelf and disconnect the power cords.
2. Record the location of each controller, power supply, and cache module in the shelf.
3. Record the location of each controller interface cable.
4. Record the location of each cable.

5. Remove the 5/16 nuts holding both of the front locking brackets.

CAUTION

Be sure to fully support the weight of the shelf with both hands at all times.

6. Note the shelf orientation and carefully slide it out of the mounting brackets.

To replace a BA350–M shelf, reverse the previous procedure.

8.11.3 Removing and Replacing a BA350–S Shelf

All BA350–S shelves are inserted into a set of mounting brackets and secured in place with a front locking bracket. Use the following procedures to remove and replace a BA350–S shelf:

Removing a BA350–S Shelf

Do the following to remove a BA350–S shelf:

1. Turn off the power to the shelf and disconnect the power cords.
2. Record the location of each storage device, controller, power supply, and cache module in the shelf.
3. Record the location of each controller interface cable.
4. Remove storage devices for access to the SCSI cable connectors.
5. Record the location of each SCSI cable.
6. Remove the nuts holding both of the front locking brackets.

CAUTION

Be sure to fully support the weight of the shelf with both hands at all times. Removing all the SBBs significantly reduce the shelf weight.

7. Note the shelf orientation and carefully slide it out of the mounting brackets.
8. Remove the two blowers from the back of the shelf (see Section 8.11.4).
9. Record the position of the jumper and terminator on the back of the BA350–S shelf.

Replacing a BA350–S Shelf

Do the following to replace a BA350–S shelf:

1. Check and reposition the jumper and terminator on the back of the new shelf to match the positions on the removed shelf.
2. Orient the replacement shelf and carefully slide it into the mounting brackets.
3. When the shelf is fully seated, install and secure the front locking brackets and secure the retaining nuts.
4. Connect each SCSI cable to the same connector from which it was removed.
5. Install each SBB, controller, and cache module in the same location from which it was removed.

6. Connect each controller interface cable to the same connector from which it was removed.
7. Insert the power SBB and connect the power cord.
8. Turn on power to the shelf and ensure that the shelf, the power supplies, and all devices are functioning properly.

8.11.4 Replacing a Shelf Blower

WARNING

To reduce the risk of electrical energy hazard, disconnect the power cables from the shelf power supplies before removing shelf blower assemblies or performing service in the backplane area, such as modifying the SCSI bus.

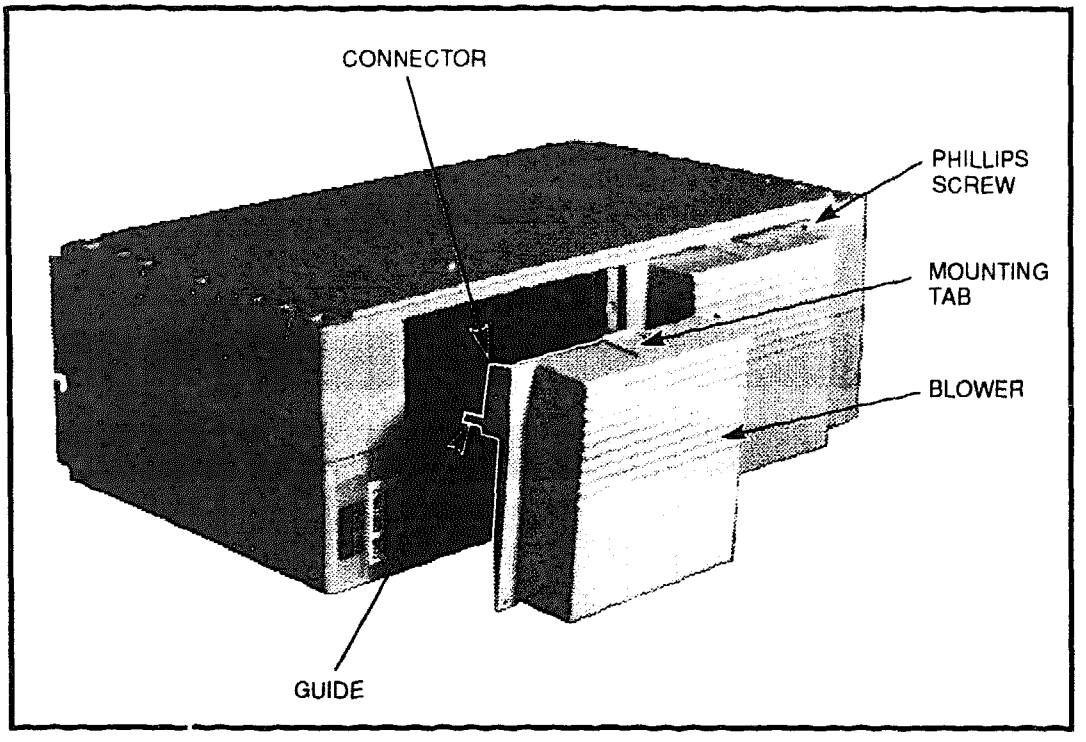
Use the following procedure to replace a shelf blower (see Figure 8–21).

1. If you cannot access the rear of the shelf, remove the shelf as described in Section 8.11.3.
2. Disconnect the power cables to the shelf power SBBs.
3. Use a Phillips screwdriver to remove the safety screw in the corner of the defective blower.
4. Press the upper and lower blower mounting tabs together to release the blower.
5. Pull the blower straight out to disconnect it from the shelf power connector.
6. Align the replacement blower connector and push the blower straight in, making sure that both mounting tabs lock in place.
7. Replace the safety screw.
8. Replace the shelf as described in Section 8.11.3.
9. Connect the shelf power cables and verify that the shelf and all SBBs are operating properly.
10. Check shelf blower for proper operation.

Note

If the upper power supply LED (shelf status) does not come on and all the shelf power supplies are operating, the second blower may have failed or the wrong blower was replaced.

Figure 8–21 Replacing Shelf Blowers



CXO-3659A-PH

8.11.5 Replacing a Shelf Power Supply

There are two methods for replacing a power supply: the **warm swap** method and the **cold swap** method.

- Use the warm swap method to replace power supplies *only* when there are two power supplies in a shelf. This method allows you to remove the defective power supply while the other supply furnishes the power.

  Note

  The warm swap method of power supply removal does not disable the controller shelf or the associated SCSI buses.

- The cold swap method normally is used during initial installation or when there is no operational shelf power supply. Should this occur, the controller shelf, the controller, the cache module, and all associated SCSI buses are disabled. None of the devices is operational until the power is restored.

  Note

  The procedures for removing power supplies and SBBs are basically the same.

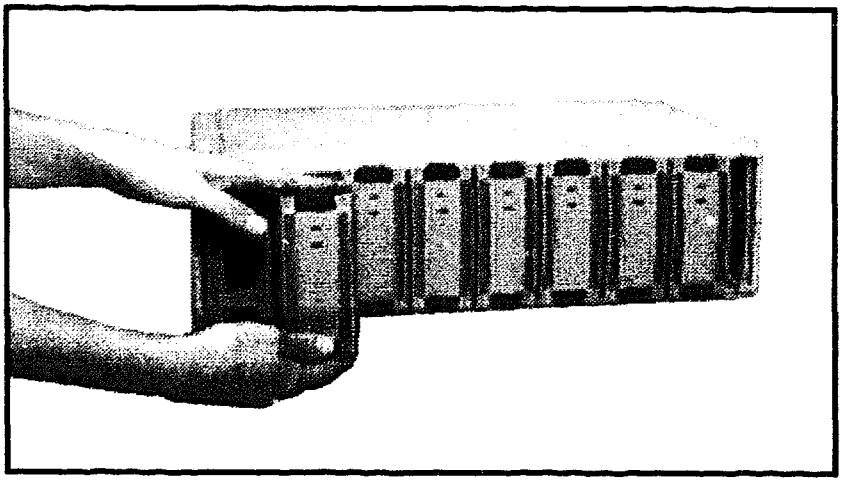
Use the following procedure to remove or replace a power unit.

CAUTION

The power supply is relatively heavy and can be damaged if dropped. Therefore, always use both hands to fully support the power supply during removal or installation.

1. As shown in Figure 8–22, press the two mounting tabs together to release the power supply from the shelf.
2. Use both hands and pull the power supply out of the shelf.
3. Insert the replacement power supply into the guide slots and push it in until it is fully seated and the mounting tabs engage the shelf.
4. After input power is applied, observe the power supply status LEDs to make sure the power supply is functioning properly. Both status LEDs should be on.

Figure 8–22 Removing a Power Supply



CXO-3611B-PH

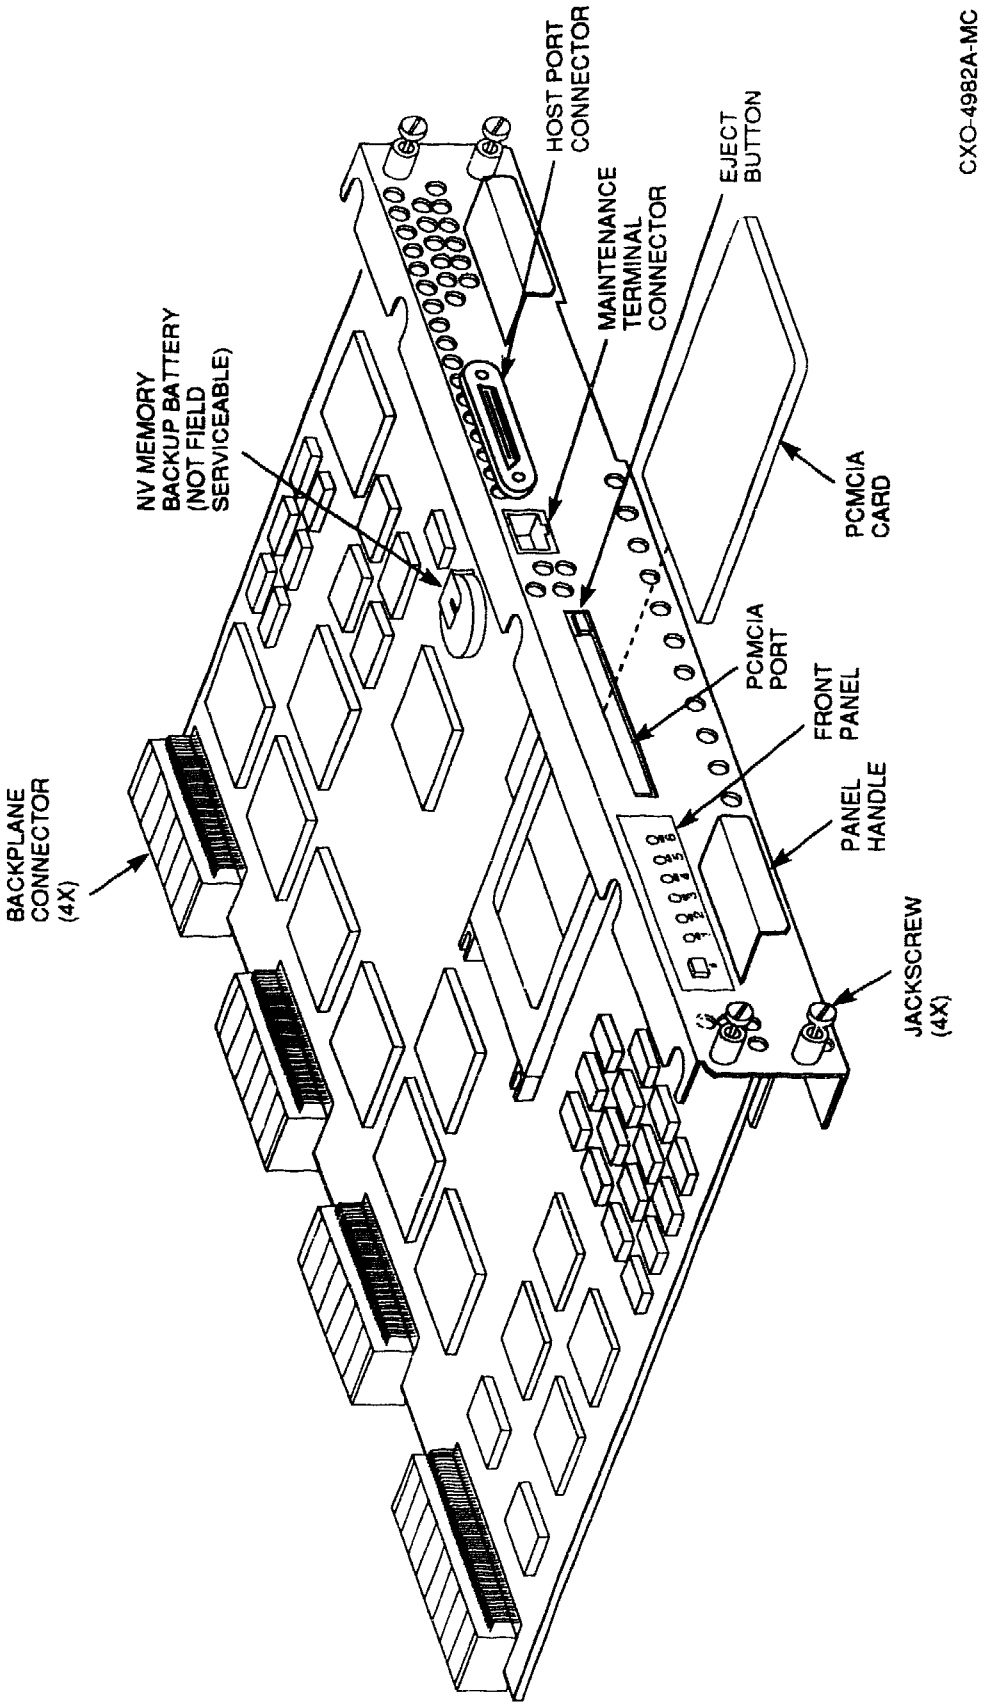
8.12 Removing and Replacing the HS1CP and HS1CP Components

The following sections describe the procedures for removing and replacing the HS1CP and its components.

Note

Depending on the situation, the HS1CP can be changed either while the system is operating (warm swap) or with power and data inputs removed. The decision as to which procedure to use depends on FDDI Server cabinet configuration and operational circumstances at each user site.

Figure 8–23 HS1CP and Cache Memory Boards



8.12.1 Reconciling Firmware in Different Version HS1CPs

When you add a REV C01 HS1CP device channel processor to an existing HS110/111/121 FDDI Server which has an older revision (REV A01) HS1CP installed, the SET FAILOVER COPY command used to copy the storage configuration from one HS1CP to another in the same dual-redundant pair will fail because of a difference in Product IDs between the two HS1CP revisions. The HS1CPs function without problem individually but will not work together as a dual-redundant pair.

HS1CPs at REV A01 were shipped with product ID "HSD40." With the introduction of Version 2.5 firmware for the HS1CP, the HS1CP had its Product ID modified to "HS1CP." This change started with REV C01.

Customers with REV A01 HS1CP device channel processors with Version 2.5 firmware should update the Product ID of the REV A01 HS1CPs to reflect "HS1CP". Follow the procedure below:

Note

This process can only be run on HS1CPs upgraded to Version 2.5 firmware. Version 2.0 firmware does not support the changing of Product IDs. To upgrade the firmware, obtain and install a PCMCIA card with the correct firmware version.

1. Contact the Customer Support Center for the specific steps for this update procedure. You will need the HS1CP serial number and security code which can be found by entering the SHOW CONFIG command when a terminal is attached to the HS1CP.
2. With this information, the Customer Support Center will run an encryption program to generate an update code. You will use this code to modify the Product ID of the HS1CP to reflect "HS1CP."

8.12.2 Using the Warm Swap Utility to Replace an HS1CP and HS1CP Components

The warm swap utility, referred to as C_SWAP, is used to replace controller and cache modules in dual-redundant controller configurations with power applied to the controller shelf. C_SWAP also is used for adding new controller and cache modules to preexisting nonredundant controller configurations to form dual-redundant configurations.

CAUTION

Ensure that your HS1CP's host port connector has mating guides installed before using the C_SWAP utility. Failing to do so risks shorting circuits that could blow fuses on all the devices on the bus.

C_SWAP fails over all I/O operations to the HS1CP running C_SWAP to allow the bus settling time in preparation for the remove/replace operation. When warm swapping an HS1CP, you are removing/replacing an HS1CP in the most transparent method available to the HS1CP subsystem.

Warm swap differs from executing the CLI SHUTDOWN command in that devices are not off line until a replacement is found, and you do not need to remove power from the HS1CP.

CAUTION

Warm swap only one HS1CP at a time. Never attempt to remove both HS1CPs at the same time in a dual-redundant configuration. The other (good) HS1CP must remain operational to support not only its own devices, but also those devices that failed over from the failed HS1CP.

8.12.2.1 When to Use C_SWAP

Use C_SWAP when you want to remove and replace a single failed HS1CP or cache module that is part of a redundant pair. You can either replace the module immediately while in C_SWAP, or exit C_SWAP and leave the failed-over HS1CP running until you replace the module at a later time.

Because C_SWAP performs a quiesce operation on all ports of the HS1CP, you should use C_SWAP even if failover has already occurred due to a fatal error of one HS1CP in the pair.

The operator must quiesce SCSI bus operations during a device warm swap. Quiesce means to make a bus inactive or dormant. All activity on the bus where the device is being swapped must be halted for the duration of the removal, replacement, or addition.

8.12.2.2 Functions of C_SWAP

C_SWAP performs the following functions during module removal or replacement procedures:

- Prompts you to identify the module to be removed
- Fails over I/O to the HS1CP that is running C_SWAP
- Performs a quiesce procedure for all HS1CP ports on the bus
- Prompts you to remove and replace the failed HS1CP or cache module
- Restarts the subsystem in either failed-over configuration or full redundant operation

You can abort the C_SWAP utility by entering either Ctrl/C or Ctrl/Y (followed by pressing the Return key) anytime during the utility's operation.

8.12.3 Removing an HS1CP or Cache Module During a C_SWAP Operation

Use the following procedure as an example for warm swapping an HS1CP:

1. Wear an ESD grounding strap and ground yourself to the cabinet grounding stud before servicing the HS1CP or cache module.
2. Connect the maintenance terminal to the HS1CP you are **not** removing.
3. If the HS1CP you are removing has totally failed, proceed to step 4.

   If the HS1CP you are removing has any I/O outstanding or devices online, shut down that HS1CP using the procedure in Section 8.5.2.
4. Enter the following command:

```
HS1CP> RUN C_SWAP
```

The system responds with:

```
Controller Warm Swap, Software Version -V2.0
*** Sequence to REMOVE other HSICP has begun. ***
Do you wish to REMOVE the other HSICP Y/N [N]? YES
```

5. Enter "Y" to continue the procedure.

```
Will its cache module also be removed Y/N [N]? YES
```

6. Enter "Y" only if you are removing the HS1CP's cache module as well.

After you answer the HS1CP/cache module identity prompts, C_SWAP fails over the I/O for the failed HS1CP to the HS1CP running C_SWAP. The failover operation is complete when the green indicator (reset LED) on the HS1CP to be replaced stays lit. C_SWAP then performs a bus quiesce operation and prompts you to remove the specified modules, as shown in the following example:

```
Killing other controller. Attempting to quiesce all ports.
Port 1 quiesced. Port 2 quiesced. Port 3 quiesced. Port 4
quiesced. Port 5 quiesced. Port 6 quiesced.
All ports quiesced.
```

7. Remove the other HS1CP (the one *without* a blinking green LED) within 5 minutes. Time remaining is 4 minutes, 50 seconds.

CAUTION

Do not remove the HS1CP with the blinking green LED Reset button.

You have 5 minutes to perform the remove operation for a HS1CP/cache module, and 2 minutes for an HS1CP only, following the steps described in Sections 8.12.5 and 8.12.6.

8. If you do not complete the removal in this time, the configuration that was running when you started is restored, all ports are restarted, and C_SWAP exits. All ports remain quiesced during this interval. If you have not inserted a replacement HS1CP or cache module, you are prompted as follows:

```
Do you have a replacement HSICP readily available [N]?
Return, No OTHER HSICP immediately available, program
exiting.
```

If you enter NO or press Return in response to this prompt, C_SWAP exits with the HS1CP in the failed-over configuration. If you enter YES, you are given an additional time interval to perform the remove/replace operation.

The replacement configuration must contain both a cache module and an HS1CP. The replacement HS1CP must have its program card removed for C_SWAP to restart the configuration properly. Once the second HS1CP is detected, C_SWAP displays a restart procedure and exits. You must enter the RESTART OTHER_CONTROLLER command to restart the configuration properly.

8.12.4 Replacing an HS1CP and/or Cache Module During a C_SWAP Operation

The following is an example of a C_SWAP session for an HS1CP/cache module swap in which the replacement HS1CP is *not* immediately available. After you enter the RUN C_SWAP command, C_SWAP prompts you as follows:

```
Do you have a replacement HS1CP readily available [N]?
NO
```

Try to have a replacement available. If you do not have one, you have to enter NO. Then, the warm swap sequence terminates, and you have to restart the routine later when you have a replacement module.

When you find a replacement, you can restart the sequence by entering the RUN C_SWAP command again. The system responds with the following:

```
Do you have a replacement HS1CP readily available [N]?
YES
```

1. **Enter "Y" if you have the replacement module.**

   ```
   *** Sequence to INSERT other HS1CP has begun. ***

   Do you wish to INSERT the other HS1CP [N]? YES
   ```

2. **Enter Y to insert the replacement module. (A message about the cache module appears only if you removed the cache module in the first place.)**

   Remember to reinsert the cache module, if necessary.

   ```
   Attempting to quiesce all ports.

   Port 1 quiesced. Port 2 quiesced. Port 3 quiesced. Port 4
   quiesced. Port 5 quiesced. Port 6 quiesced.

   All ports quiesced.

   Insert the cache module, then insert the other HS1CP, WITHOUT its
   program card, and press Return.
   ```

3. **Insert the HS1CP and cache module (if applicable) now. Follow the steps outlined in Section 8.12.7 to physically replace your HS1CP, or Section 8.12.8 to replace your cache module.**

   ```
   Restarting ALL ports.

   Port 1 restarted. Port 2 restarted. Port 3 restarted. Port 4
   restarted. Port 5 restarted. Port 6 restarted.

   The configuration has two HS1CPs.

   The Warm Swap program has terminated.
   ```

4. To restart the other HS1CP, do the following:

   a. Enter the RESTART OTHER_CONTROLLER command.

   b. Press and hold the Reset button while inserting the program card.

   c. Release the Reset button. The HS1CP initializes.

   d. Configure the new HS1CP by using the HS1CP Configure Storage procedure.

Configure Storage for the Second HS1CP

When you add a second HS1CP and/or StorageWorks storage devices to the FDDI Server, you must configure the storage system to meet your particular storage needs.

Appendix B in the *StorageWorks FDDI Server Installation Guide* provides a detailed description of stripesets and RAIDsets you can configure for the FDDI Server. It also provides rules for configuring the storage devices.

When you are configuring the storage devices connected to a second HS1CP for redundancy, both HS1CPs must have the same storage configuration. Follow the procedure below to copy the storage configuration from the original HS1CP1 to the newly installed HS1CP2.

1. Make sure you have a terminal connected to the newly installed HS1CP2 as described in Section 8.3.
2. Copy the configuration from HS1CP1 to the newly installed HS1CP2 with the following command:

   ```
   HS1CP2> SET FAILOVER COPY = other
   ```

3. Restart the newly installed HS1CP2 by entering the following command on the terminal connected to the HS1CP2:

   ```
   HS1CP2> RESTART THIS
   ```

   These messages 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
   ```

4. Enable the path from the HS1CP to the server processor by entering the following command:

   ```
   HS1CP2> SET THIS_CONTROLLER PATH
   ```

   You have completed the configuration procedure and are ready to save the storage configuration.

8.12.5 Removing an HS1CP

When running C_SWAP, you are prompted to remove the HS1CP and cache modules. Use the procedure in Table 8–3 to remove the HS1CP.

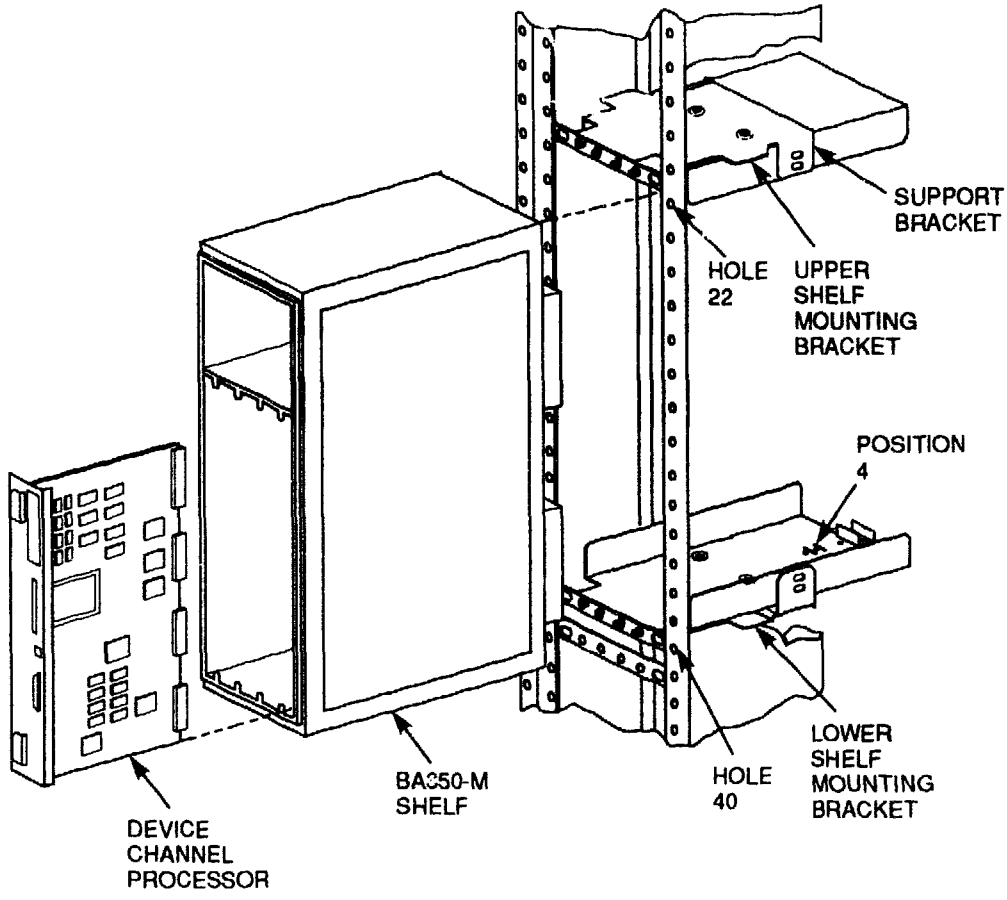
Note

Use proper ESD procedures when handling HS1CPs or cache modules, as described in Section 8.1.

Table 8–3 HS1CP Removal

Step	Description
1	Unlock and open the cabinet doors (SW800-series) using a 5/32-inch Allen wrench.
2	Ground yourself to the cabinet grounding stud.
3	Unsnap and remove the program card ESD shield that covers the program card.
4	Remove the program card by pushing the eject button next to the card. Pull the card out and save it for use in the replacement HS1CP.
5	With a small flathead screwdriver, loosen the captive screws on the trilink connector and remove the trilink from the front of the HS1CP. You need to work around any internal bus cable or terminator connections when removing the trilink. Do *not* remove the cables or terminators from the trilink or you interrupt the host bus.
6	Loosen the four screws (refer to Figure 8–24) on each side of the front bezel.
7	Use a gentle up-and-down rocking motion to loosen the module from the shelf backplane.
8	Slide the module out of the shelf (noting in which rails the module was seated) and place it on an approved ESD work surface or mat until it can be packaged for shipment to the Digital repair depot.
9	If necessary, you now can remove the cache module as described in Section 8.12.6.

Figure 8–24 HS1CP Installation



8.12.6 Removing a Cache Memory Module

Most HS1CPs have either a read or write-back cache memory module installed behind them in the BA350–M shelf.

Use the following procedure to remove the cache memory module:

WARNING

The write-back cache is a high-energy module. Do not allow the write-back cache to contact any conductive surface, or injury and/or equipment damage could result.

CAUTION

The write-back cache batteries add weight to the module. Grasp the module firmly by the side where the batteries are seated anytime you handle the module. If you need to change the batteries on your write-back cache module, refer to the *StorageWorks Array Controllers HS Family of Array Controllers Service Manual.*

Table 8–4 Cache Module Removal

Step	Description
1	Ground yourself to the cabinet grounding stud.
2	The HS1CP is seated in front of the cache module. Anytime you service a cache module, you must remove the HS1CP based on considerations of configuration, down time, and so on.
3	To access the cache module, remove its HS1CP. To remove the HS1CP, refer to Section 8.12.5.
4	Use a gentle up-and-down rocking motion to loosen the cache module from the shelf backplane.
5	Slide the cache module out of the shelf and place it on an approved ESD mat.

8.12.7 Replacing an HS1CP

Note

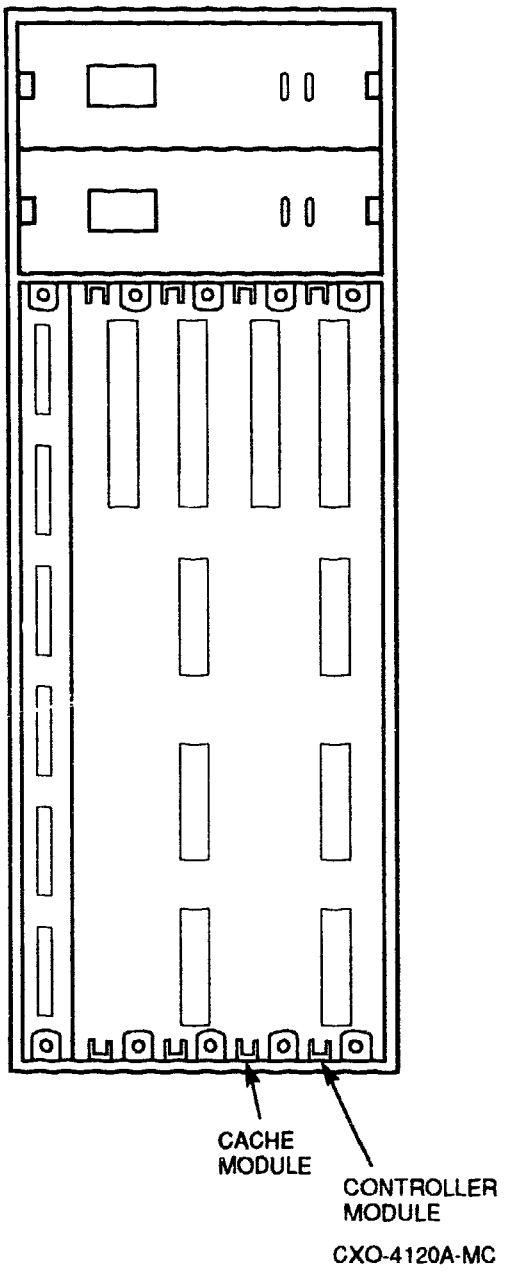
Use proper ESD procedures when handling HS1CPa or cache modules, as described in Section 8.1.

Use the procedure in Table 8–5 to replace your HS1CP during the warm swap replacement.

Table 8–5 HS1CP Replacement

Step	Description
1	Ground yourself to the cabinet grounding stud.
2	If the cache memory module was removed, replace it at this time. (Refer to Section 8.12.8).
4	Slide the HS1CP into the shelf using its slot's rightmost rails as guides (refer to Figure 8–25).
5	Use a gentle up-and-down rocking motion to help seat the module into the backplane. Press firmly on the module until it is seated.
6	Tighten the four hex screws on the front bezel using a 3/32-inch Allen wrench.
7	Connect the host interface cable (CI), and tighten its captive screws using a flathead screwdriver.
8	Replace the program card by inserting the card into the program card slot and pressing the card until the eject button is even with the edge of the card.
9	Snap the program card ESD shield into place by pressing the two plastic push pins on each side of the shield.
10	Remove the ESD grounding strap from the cabinet and your wrist.
11	Close and lock the cabinet door.
12	Connect a maintenance terminal to the MMJ of the other HS1CP (the one you did not replace) if one is not connected already.

Figure 8–25 HS1CP Shelf Rails



8.12.8 Replacing a Cache Module

Use the procedure in Table 8–6 to replace a cache module.

Table 8–6 Cache Module Replacement

Step	Description
1	The HS1CP is seated in front of the cache module. To service a cache module, you must remove the HS1CP.
2	To replace the cache module, slide it into the shelf using the cache module's slot's leftmost rails as guides (refer to Figure 8–25).
3	Press firmly and use gentle pressure and a rocking motion on the module until it is firmly seated.
4	Replace the HS1CP (refer to Table 8–5).

8.12.9 Removing, Disposing of, and Replacing a Write-Back Cache Battery

The following sections discuss how to remove, dispose of, and replace a write-back cache battery.

WARNING

The write-back cache batteries are high energy devices. Do not allow the battery contacts to touch any conductive surface, or injury and/or equipment damage could result.

Note

Although service described in this manual is limited to removing and replacing batteries, it is mandatory that Digital Multivendor Customer Services *never* repair the write-back cache module while the batteries are installed.

8.12.9.1 Removing a Write-Back Cache Battery

Use the following procedure to remove the write-back cache battery:

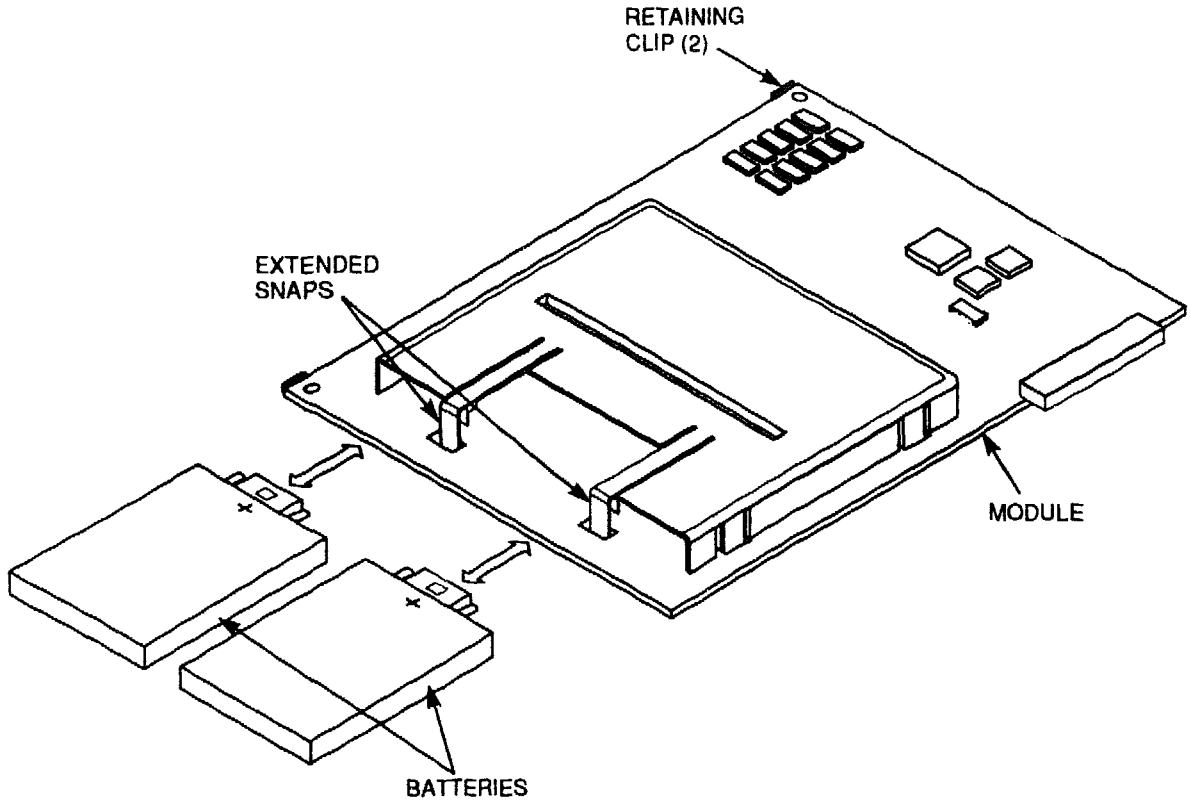
1. Locate the two write-back cache batteries seated under the plastic bracket at one end of the module.
2. Disengage and gently lift the extended snaps on the plastic bracket, and pull the batteries straight out of the module (see Figure 8–26). Place the batteries on an approved *nonconductive* ESD mat.

8.12.9.2 Disposing of a Write-Back Cache Battery

Dispose of batteries that are no longer capable of being recharged. Be aware that the write-back cache batteries contain lead. Therefore, follow disposal regulations applicable to their composition and marking in accordance with your local recycling laws.

Consult your service policies regarding hazardous materials handling for details. Do not simply discard spent batteries in the trash unless permitted by local regulations.

Figure 8–26 Write-Back Cache Batteries



8.12.9.3 Replacing a Write-Back Cache Battery

1. Disengage and gently lift the extended snaps on the plastic bracket and slide the batteries, positive (+) side up, straight into the module (refer to Figure 8–26).
2. Press the batteries firmly into the module so they are seated and the extended snaps close easily around the batteries.

   If you install a battery upside down accidentally, the battery does not seat firmly and the retaining clip does not close.

8.12.10 Verifying Firmware Versions and Write-Back Cache Size

CAUTION

In a redundant configuration, the following must apply:

- The firmware version and patch level must be the same for both HS1CPs in a dual-redundant pair for proper operation of the server.
- Both write-back cache modules must be the same cache size.

When the firmware versions are mismatched, the HS1CPs detect the mismatch and do not allow access to any devices.

1. Enter the following command at the HS1CP2> prompt:

   ```
   HS1CP2> SHOW THIS
   ```

   A display similar to the example below displays on your screen:

```
HS1CP2> SHOW THIS

Controller:
        HS1CP   (C) DEC ZG43100021 Firmware V25F-0, Hardware AX01
        Configured for dual-redundancy with ZG43100011
            In dual-redundant configuration
        SCSI address 7
        Time: 24-APR-1995 09:56:50

Host port:
        Node name: HS1CP2, valid DSSI node 2
        Host path is ON
        MSCP allocation class   13
        TMSCP allocation class  13

Cache:
        32 megabyte write cache, version 2
        Cache is GOOD
        Battery is GOOD
        No unflushed data in cache
        CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
```

2. The firmware version and patch level of the HS1CP displays on the first line of the text of the "Controller:" field, "V25F" in this case.

   The cache information displays in the "Cache:" field.

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.

3. Record the firmware version and cache information for later use.

4. Enter the following command at the HS1CP2> prompt:

   ```
   HS1CP2> SHOW OTHER
   ```

   A display similar to the following example displays on your terminal screen:

```
HS1CP2> SHOW OTHER

Controller:
        HS1CP   (C) DEC ZG43100021 Firmware V25F-0, Hardware AX01
        Configured for dual-redundancy with ZG43100011
            In dual-redundant configuration
        SCSI address 6
        Time: 24-APR-1995 09:56:50

Host port:
        Node name: HS1CP1, valid DSSI node 1
        Host path is ON
        MSCP allocation class   13
        TMSCP allocation class  13
```

```
Cache:
        32 megabyte write cache, version 2
        Cache is GOOD
        Battery is GOOD
        No unflushed data in cache
        CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
```

5. Compare the displayed values of the two HS1CPs and follow these instructions:
   - If the firmware version, patch level, and cache sizes match, you have completed the verification procedure.
   - If the firmware versions are different, you must disable the HS1CP that has the lower version of code by removing its PCMCIA card.
   - As soon as possible, obtain and install a PCMCIA card with the correct firmware version.
6. Compare the cache size on the two write-back cache modules and follow these instructions:
   - If the cache size of the two write-back cache modules matches, you have completed the verification procedure.
   - If the cache size of the two write-back cache modules does not match, obtain and install the correct cache module as soon as possible. Meanwhile, do not enable the write-back cache function on the module with the wrong memory size with the Firmware Licensing System (FLS) utility.
7. When you have the correct firmware version and write-back cache modules installed, repeat this verification procedure.

8.12.11 Saving the Storage Configuration

After defining the storage configuration, 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 storage configuration, perform the following steps:

1. Unplug the terminal cable from the newly installed HS1CP and insert the cable into the back of your server processor.
2. Connect the FDDI Server terminal logically to the HS1CP and record the results of the session in a file called CONFIG.LOG in the SYS$MANAGER directory on the server's system disk.

   Enter the following command at the DCL prompt to logically connect the FDDI Server to the HS1CP:

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

   where:

   node-name is the name assigned to the HS1CP (for example, HS1CP2).

   The HS1CP responds with a brief display and its prompt, as follows:

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

Last fail code: 018000A0

Press " ?" at any time for help.

HS1CP2>
```

3. At the HS1CP2> prompt, enter the following commands to capture a complete record of the storage configuration:

```
HS1CP2> SHOW THIS_CONTROLLER

HS1CP2> SHOW OTHER_CONTROLLER

HS1CP2> SHOW DEVICE FULL

HS1CP2> SHOW UNIT FULL
```

If you have created stripesets, raidsets, sparesets, mirrorsets, or passthrough containers, capture these configurations using the following command:

```
HS1CP2> SHOW STORAGESETS FULL
```

The information displayed on the terminal screen for each command is captured automatically in the CONFIG.LOG file in the SYS$MANAGER directory on the server's system disk, providing a record of the storage configuration.

4. Exit the display by entering the following command at the prompt:

```
HS1CP2> EXIT
```

8.12.12 Verifying Installation Procedures: Creating Redundant Data Paths

You must verify all the previous installation procedures, including:

- HS1AD installation
- HS1CP installation
- Internal bus cable installation

Use the following procedure to verify your installation:

Step1: Check for server processors and HS1CPs

- With the terminal connected to the server processor, enter the following command at the DCL prompt:

```
$ SHOW CLUSTER
```

- The names of all server processors and HS1CPs installed in your FDDI Server are displayed on your terminal screen in a table similar to the following example.

```
-------------------------------
|        SYSTEMS        | MEMBERS |
|-----------------------------|
|   NODE  |  SOFTWARE   | STATUS  |
|-----------------------------|
| FSERV1  |  VMS 6.2    | MEMBER  |
| FSERV2  |  VMS 6.2    | MEMBER  |
| HS1CP1  |  HSD V25F   |         |
| HS1CP2  |  HSD V25F   |         |
-------------------------------
```

Note

The previous table is representative of an HS121 FDDI Server with two server processors named FSERV1 and FSERV2. Your system displays the actual server processor names you selected during the Software Customization Procedure and the HS1CP number with operating firmware revision level.

- If you do not see the names of both server processors displayed for your system, one of the processors did not complete the booting process. Check the console terminal for messages and refer to Chapter 4 and Chapter 6 to correct the problem that is identified.
- If you do not see all HS1CPs displayed that are installed on your system, there is a problem with the equipment or configuration. Refer to Chapter 4 for the procedures to correct the problem that is identified.

Step 2: Check for attached storage devices

- Enter the following command at the DCL prompt:

```
$ SHOW DEVICE
```

- All attached storage devices, the HS1CPs to which the storage devices have been assigned, and the allocation class you chose during the Software Customization Procedure display at this time. An example of this display follows:

```
$ SHOW DEVICE

Device                  Device      Error    Volume      Free   Trans    Mnt
 Name                    Status     Count     Label      Blocks  Count    Cnt

$13$DKA0:      (FSERV1)  Mounted        0   FSRVER1$SYS
$13$DKA100:    (FSERV2)  Mounted        0   FSRVER2$SYS 3196736  191      1
$13$DKA400:    (FSERV2)  Online wrtlck   0   FSRVER2$SYS
$13$DAK500:    (FSERV1)  Online
$13$DKA100:    (HS1CP1)  Mounted        0   13$200      4109068    1      5
$13$DUA110:    (HS1CP2)  Mounted        0   13$110      4109068    1      5
$13$DUA120:    (HS1CP1)  Mounted        0   13$120      4109068    1      5
$13$DUA130:    (HS1CP2)  Mounted        0   13$130      4109068    1      5
.
.
.

Device                  Device      Error
 Name                    Status     Count
```

```
FRA0:                   Online          0
FRA1:                   Online          0
FRYA0:                  Offline         0
PAA0:                   Online          0
PAB0:                   Online          0
PKA0:                   Online          0
RMA0:                   Online          0
SWA0:                   Online          0
```

- If all the storage devices in the FDDI Server do not appear in your display, there is a problem. Refer to Chapters 4 and 6 to correct the problem that is identified.
- If all your storage devices and HS1CPs are displayed on your screen, close all the doors on the FDDI Server.

You have completed successful verification of your installation procedures.

Follow the steps on the display message to restart the other HS1CP. If the HS1CP initializes correctly, its green reset LED begins to flash at 1 Hz. If an error occurs during initialization, the front panel displays an error code. After the new HS1CP has initialized, perform the following steps:

1. Plug a maintenance terminal into the EIA–423 terminal port in the front bezel of the new HS1CP.
2. Snap the ESD shield into place over the program card.
3. Set the new HS1CP's initial parameters as follows. Be sure to use the same parameters from the removed HS1CP when installing a replacement.

   Use the following procedure to set the initial parameters of the HS1CP:

   a. Set the HS1CP node name:

   ```
   CLI> SET THIS_CONTROLLER SCS_NODENAME="HS1CP2"
   ```

   b. Set the HS1CP path:

   ```
   CLI> SET THIS_CONTROLLER PATH
   ```

   c. Set the HS1CP identification:

   ```
   CLI> SET THIS_CONTROLLER ID=2
   ```

   d. Set the prompt:

   ```
   CLI> SET THIS_CONTROLLER PROMPT="HS1CP2"
   ```

4. Unplug the maintenance terminal from the new HS1CP.

8.13 Removing and Replacing Fans

There are three types of cooling fans in the SW800 FDDI Server cabinet:

- Cabinet cover fan assembly
- BA350 shelf fans
- The internal server processor cooling fan

This section deals with removing and replacing the fans in the top cover assembly only.

CAUTION

Electrostatic discharge damages components. Always use proper ESD grounding procedures when handling conponents. Refer to Section 8.1 for proper grounding procedures.

WARNING

Cabinet rail edges may be sharp and can slice or abrade skin and cable insulation.

8.13.1 Removing the Cabinet Cover Fan Assembly

Use the following procedure to remove the cabinet cover fan assembly:

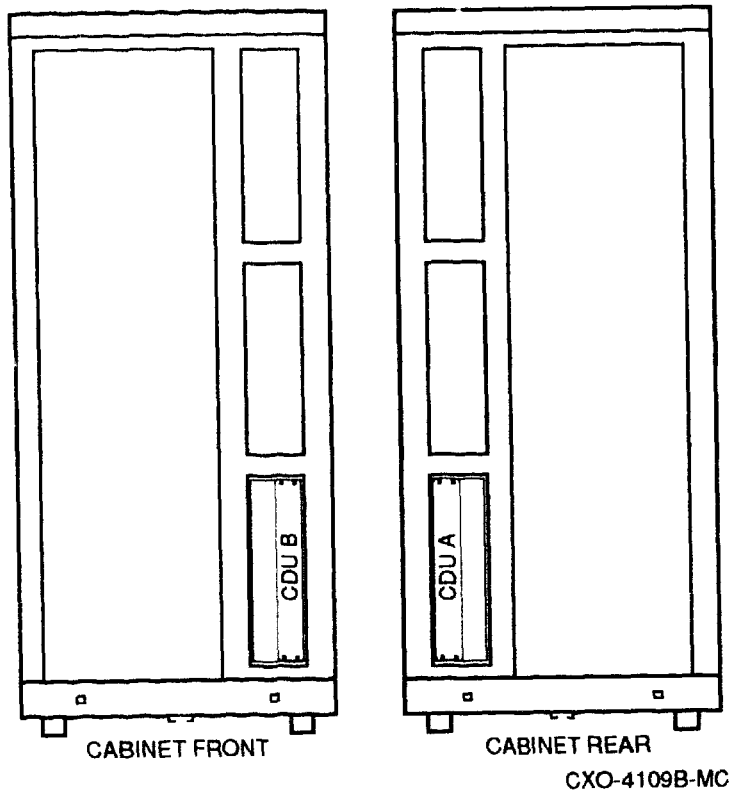
1. Remove the top cover of the SW800 cabinet.

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.

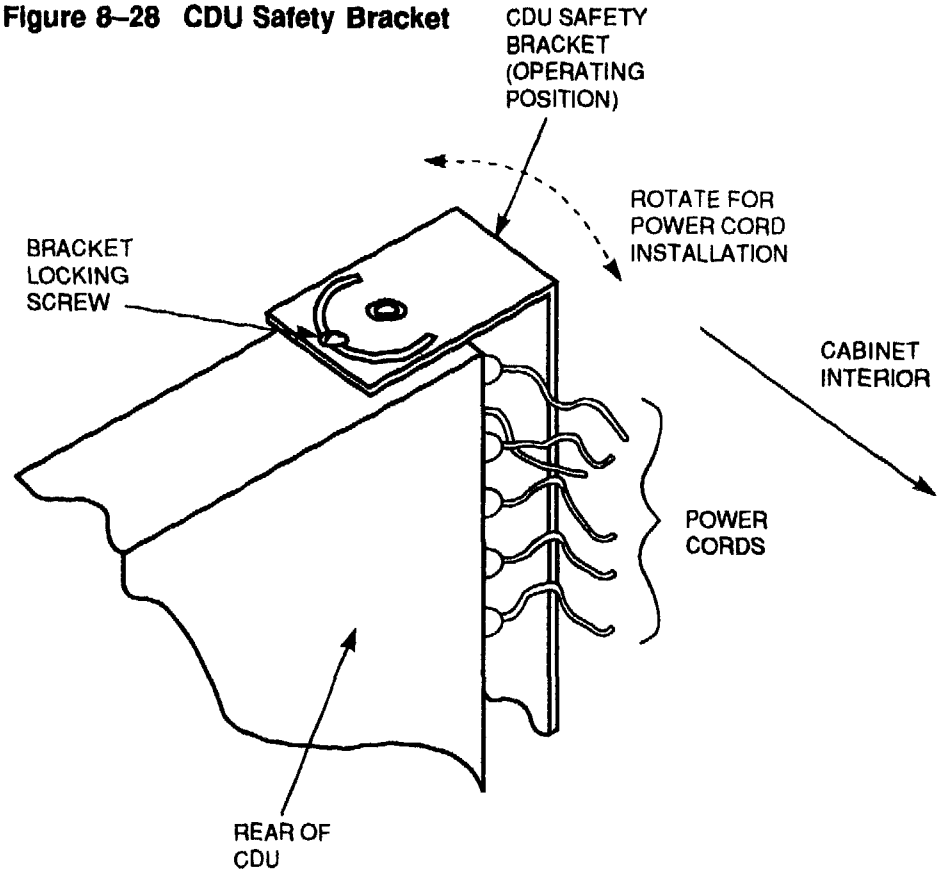
2. Loosen the top cover by pushing up on its front and rear edges until it snaps free of its fasteners.
3. Using two people, lift the top cover from the cabinet and set it aside.
4. Disconnect and remove any of the following components from the front of the cabinet (if needed) to access the cover fan assembly power cords:
   - Filler panel on the top right of the FDDI Server cabinet
   - BA350–S shelf
   - Both CDUs (see Figure 8–27)

Figure 8–27 FDDI Server Cabinet Component Locations



5. Trace the path of the black power cable and white power cable from the inside top of the power cabinet down to their respective CDU connections. Cut and discard the two cable ties holding the power cables.
6. Loosen the top and bottom bracket locking screw on the CDU safety bracket and swing the CDU safety bracket to the side (see Figure 8–28).

Figure 8–28 CDU Safety Bracket



7. Unplug the fan assembly power cable from the back of the CDU.
8. Repeat the previous two steps if two CDUs are installed.
9. Use a ladder or proper step stool to access the top of the FDDI Server cabinet.
10. Remove the twelve Phillips screws holding the fan assembly on the top cover of the FDDI Server.
11. Move the fan assembly to one side and pull up both power cables. Note that the power cables are routed through the larger hole at the top of the fan plenum.
12. Remove the defective fan assembly and replace with the new fan assembly.

CAUTION

Check the new fan assembly to make sure that it is configured for the proper voltage and frequency, which is 60Hz versus 50Hz.

13. Check each of the new cabinet fans to make sure that the blades turn freely and are not obstructed.
14. Visually inspect the cabinet to make sure that there are no obstructions to the airflow from the shelf blowers. (The side panels may need to be removed to check the shelf blowers.)

8.13.2 Installing/Replacing the Cover Fan

1. Reverse the previous procedure to install the new cabinet fan assembly.

Note

The black power cord should be connected to CDU A. If a single CDU is being used, plug both power cords into the one CDU. If a second CDU has been installed, plug the white power cord into it.

2. Install new cable ties to hold the power cords in their proper position.
3. Reconnect the FDDI Server power cords and data cables before applying system power.
4. Switch the circuit breakers on both CDUs to the | (ON) position.
5. Close and secure the cabinet doors.
6. 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.
7. Restart all disk and tape drives and check them for proper function.

8.14 Replacing Cabinet Cables and Cords

The following sections discuss SCSI cable, DSSI cable, and power cord replacement.

8.14.1 Replacing a SCSI Cable

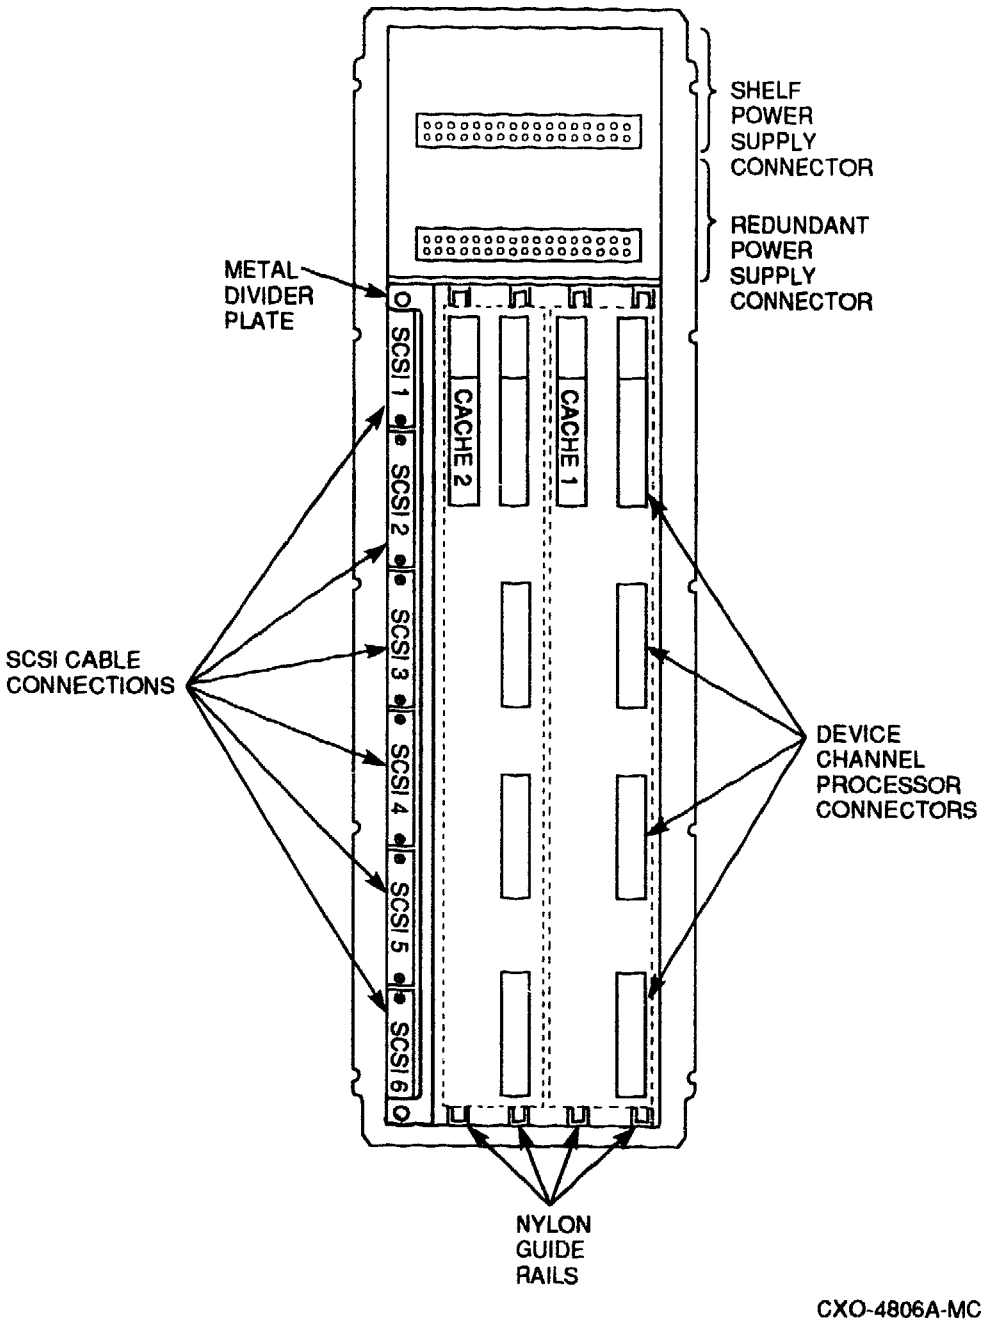
Use the following procedure to replace the SCSI cabling in the BA350–M shelf:

CAUTION

SCSI pins are bent easily. Make sure the SCSI connector pins are not bent or damaged when connecting SCSI cables.

1. Remove the metal divider plate inside the BA350–M shelf that separates the module area from the SCSI–2 port cable area in the shelf. Figure 8–29 shows the location of the metal divider plate. Removing the divider plate gives you the necessary space to access the SCSI–2 cable connectors.

Figure 8–29 BA350–M Shelf Layout (Front View)



2. Disconnect and remove the defective SCSI–2 cables from the backplane of the BA350–M shelf and clip the cable ties leading the cable to the BA350–S shelf.
3. Connect the SCSI–2 device cables to the BA350–M shelf backplane and route them out the front of the shelf and into the front of the storage shelves. Figure 8–29 shows the location of the SCSI connectors on the BA350–M backplane.
4. Plug the replacement SCSI–2 port cable into the storage shelve.
5. Replace the metal divider plate that separates the SCSI–2 port cable area from the module area in the BA350–M shelf, then tighten the two screws to hold it in place.

8.14.2 Replacing a DSSI Cable

This section describes two procedures for replacing the DSSI cable between an HS1CP device channel processor and a server processor:

- Hot swap procedure for a dual server processor FDDI Server
- Standard procedure

8.14.2.1 Hot Swap Procedure for a Dual Server Processor FDDI Server

Digital recommends that you should power everything off prior to servicing equipment such as replacing a DSSI cable. However, if the FDDI Server must remain running while servicing one of the server processors in a dual server processor FDDI Server, then the following procedure provides a means to hot swap a failed DSSI cable.

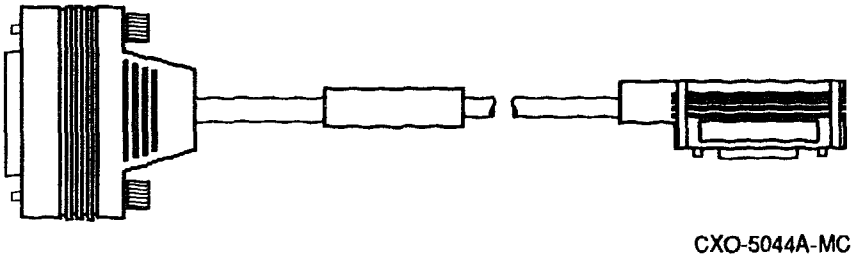
1. Detach the internal bus cable from the HS1CP and the device bus adapter in the server processor as described in Steps 1-7 of Section 8.4.1.
2. Reattach the new internal bus cable to the HS1CP and the device bus adapter in the server processor as described in Steps 6-10 and 13-17 of Section 8.4.3.

8.14.2.2 Standard Procedure

Use the following procedure to replace a DSSI cable from the HS1CP to the server processor:

1. Remove power from all units on the bus connected to the defective DSSI cable.
2. Remove straight-angle end of the DSSI cable from the back of the HS1AD device bus adapter in the server processor and the right-angle end from the trilink connector on the HS1CP. (See Figure 8–30 for a schematic of the internal bus cable connectors).

Figure 8–30 Internal Bus Cable Connectors



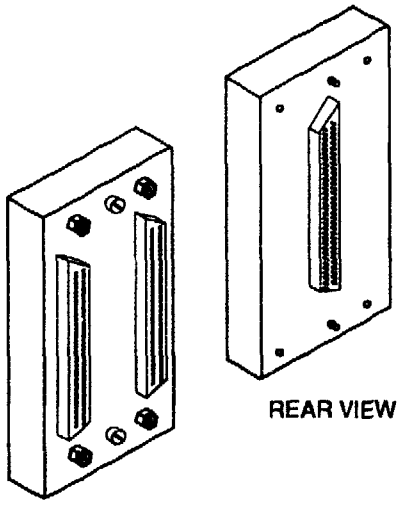
CXO-5044A-MC

3. Clip the cable ties holding the cable between the two units and remove the defective cable.
4. Reroute the new cable between the two units and fasten with new cable ties.

CAUTION

Currently, component damage can result if internal bus cables are connected or disconnected with power applied, *unless* the mating guide (Digital part number 74–49066–01) is installed around the outside edge of your HS1CP port connector. HS1CPs are shipped from the factory with the mating guide and the trilink connector block premounted on the HS1CP port connector.

Figure 8–31 Trilink Connector



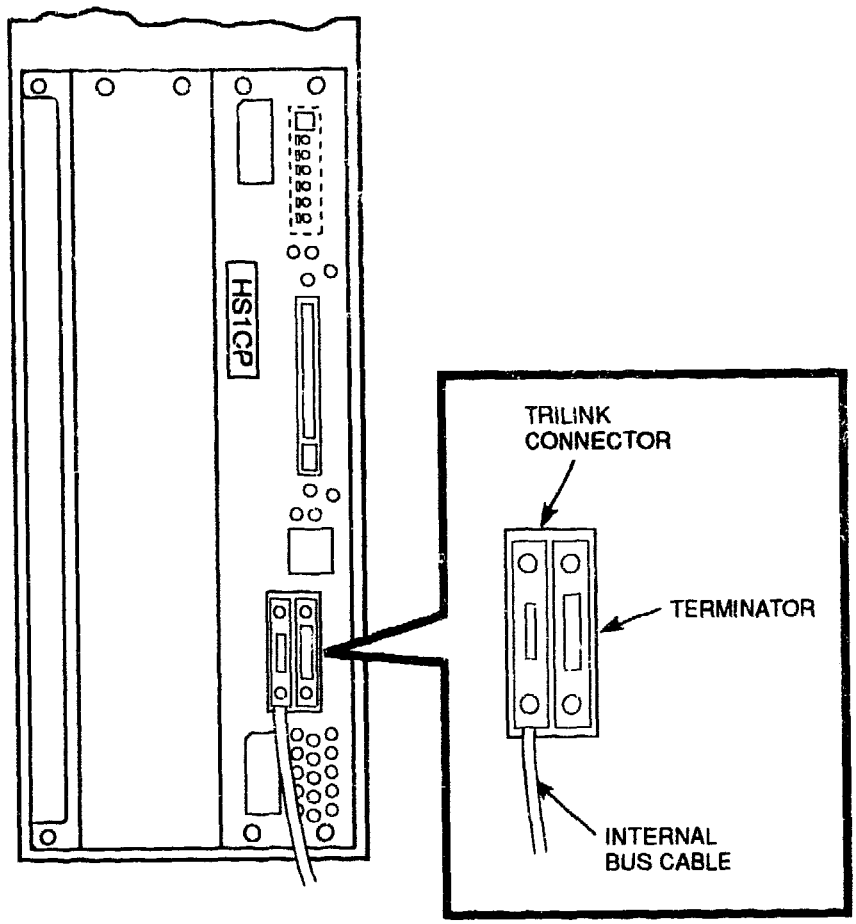
Be aware that the trilink connector block (see Figure 8–31 is the interface between the HS1CP and other internal bus cable connections. It is the *trilink* that you can disconnect and connect safely when you have a mating guide installed. In a power-on situation, you must work around any internal bus cable or terminator connections to the trilink *without* disconnecting them.

Some internal bus cable and terminator connectors do not provide enough access to the trilink screws for you to disconnect the trilink (without first disconnecting the cable and/or terminator). In these cases, you must disconnect power from all bus members and disconnect cables and terminators before disconnecting the trilink connector block.

---

5. Connect the right-angle end of the internal bus cable to the trilink connector (see Figures 8–30 and 8–32) and connect the straight-angle end to the HS1AD connector on the rear of the server processor.

Figure 8–32 Trilink Connector with Cable and Terminator



Note

If you are replacing a nonterminated internal bus cable (Digital Part Number 17–03855–*xx*) with a terminated internal bus cable (Digital Part Number 17–04288–*xx*), you must remove the three SIP terminators from the device bus adapter (inside the server processor).

6. Switch the circuit breakers on the front panels of the CDUs to the ON (|) position.
7. Depress the power button on the front of the server processor.

8.14.3 Replacing the Power Cord

The procedure used to replace a power cord varies greatly from case to case, depending on which power cord is defective and its routing from the component to the back of the CDU.

CAUTION

Electrostatic discharge damages components. Always use proper ESD grounding procedures when handling components. Refer to Section 8.1 for proper grounding procedures.

WARNING

Cabinet rail edges may be sharp and can slice or abrade skin and cable insulation.

Use the following procedure to replace a defective power cord:

1. Disconnect the defective power cord from the affected component.
2. Remove the components blocking access to the power cord cable routing from the affected component and the back of the CDU (refer to Figure 8–27).
3. Remove one or both CDUs (refer to Section 8.10).
4. Trace the path of the power cable to its respective CDU connection. Cut and discard the cable ties holding the power cable.
5. Loosen the top and bottom bracket locking screws on the CDU safety bracket and swing the CDU safety bracket to the side.
6. Remove the defective power cable.

Reverse this procedure to reinstall the new power cable.

A

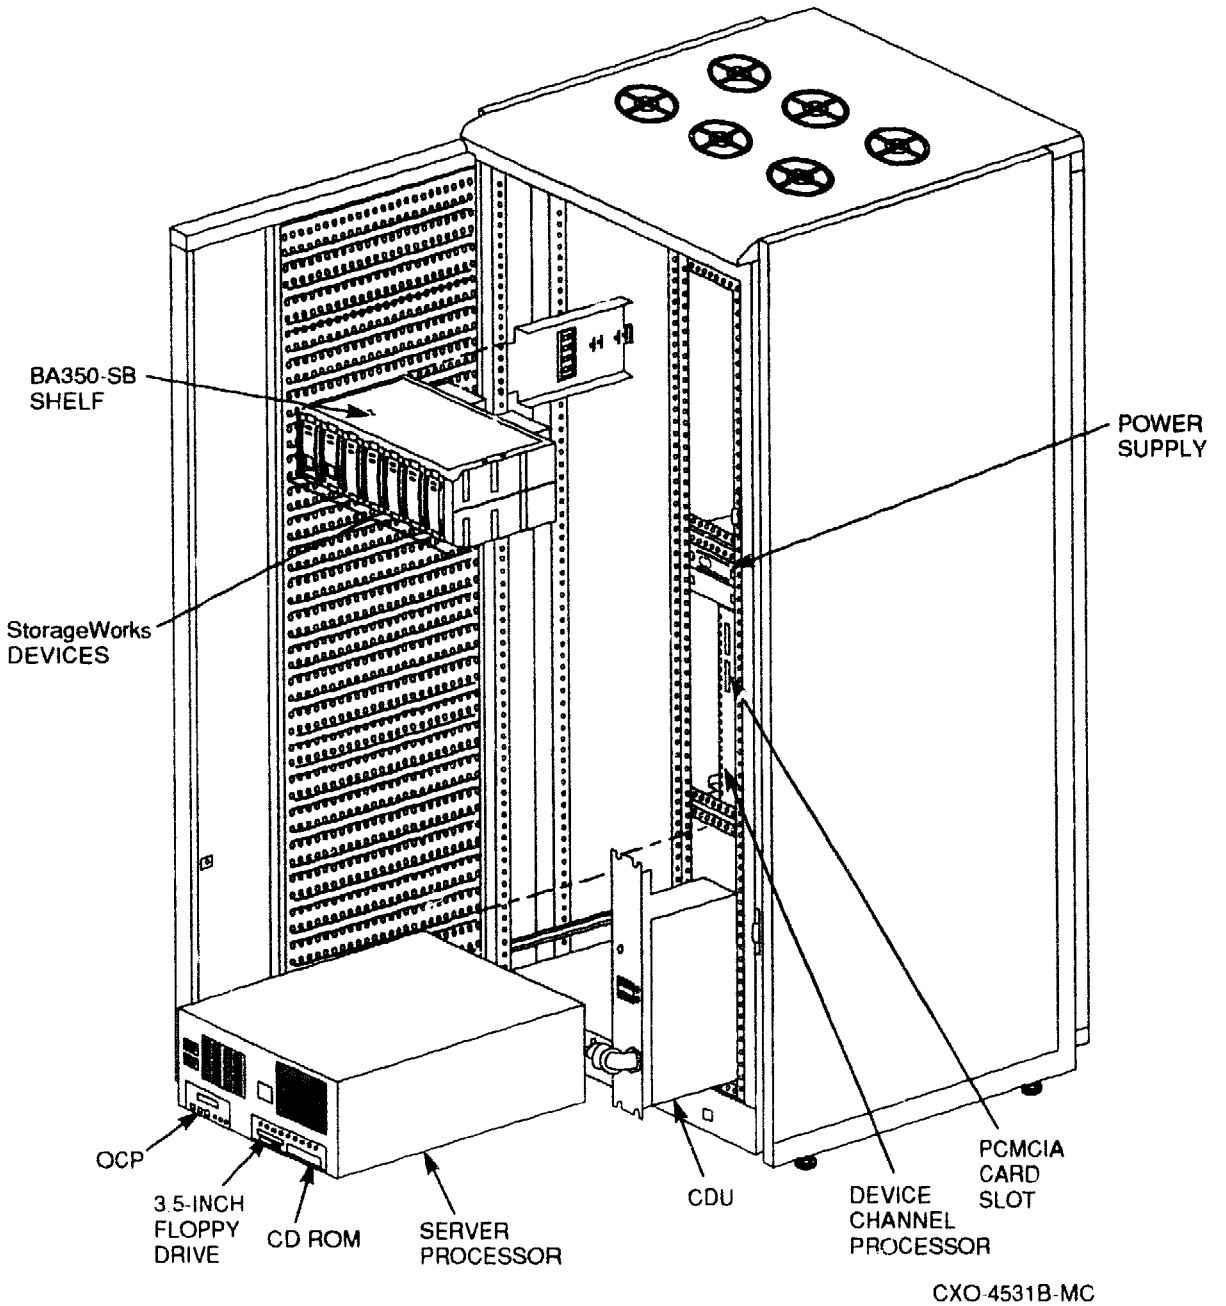
Functional Description

This appendix provides a detailed functional description of the hardware and firmware of the HS111 StorageWorks FDDI Server and the HS121 StorageWorks FDDI Server, and provides information on the following major components:

- Cable distribution unit (CDU)
- Storage devices
- Server processor components accessible from the front
- Server processor components accessible from the rear
- The HS1CP

See Figure A–1 for a view of the major components of the FDDI Server.

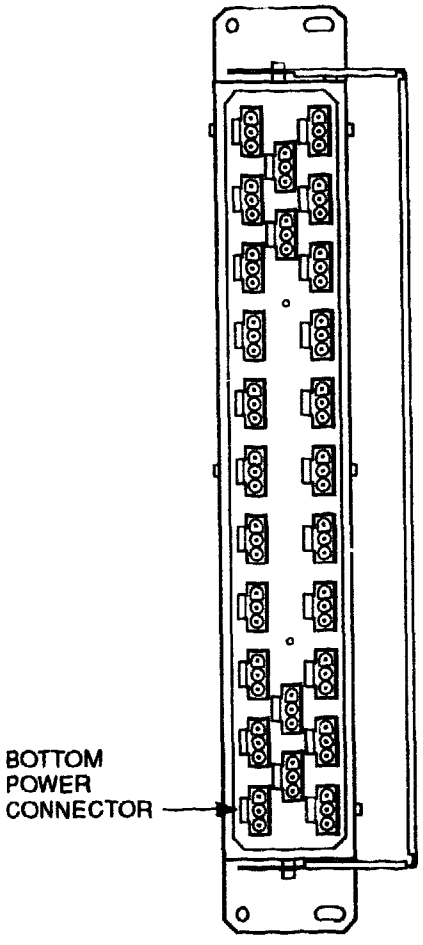
Figure A–1 StorageWorks FDDI Server–Major Components



A.1 Cable Distribution Unit (CDU)

The cable distribution unit (CDU) is the ac power entry device for StorageWorks cabinets. The unit provides the connections necessary to distribute ac power to cabinets. Figure A–2 shows the power cord connectors on the CDU. Figure A–1 shows the position of the CDU in the FDDI Server.

Figure A–2 CDU Power Cord Connectors



CXO-4671A-MC

A.2 Storage Devices

The StorageWorks FDDI Server supports a wide range of StorageWorks solid state disks, magnetic disks, tape, optical, and loader devices. The storage devices fit into the BA350–SB shelf, shown in Figure A–1.

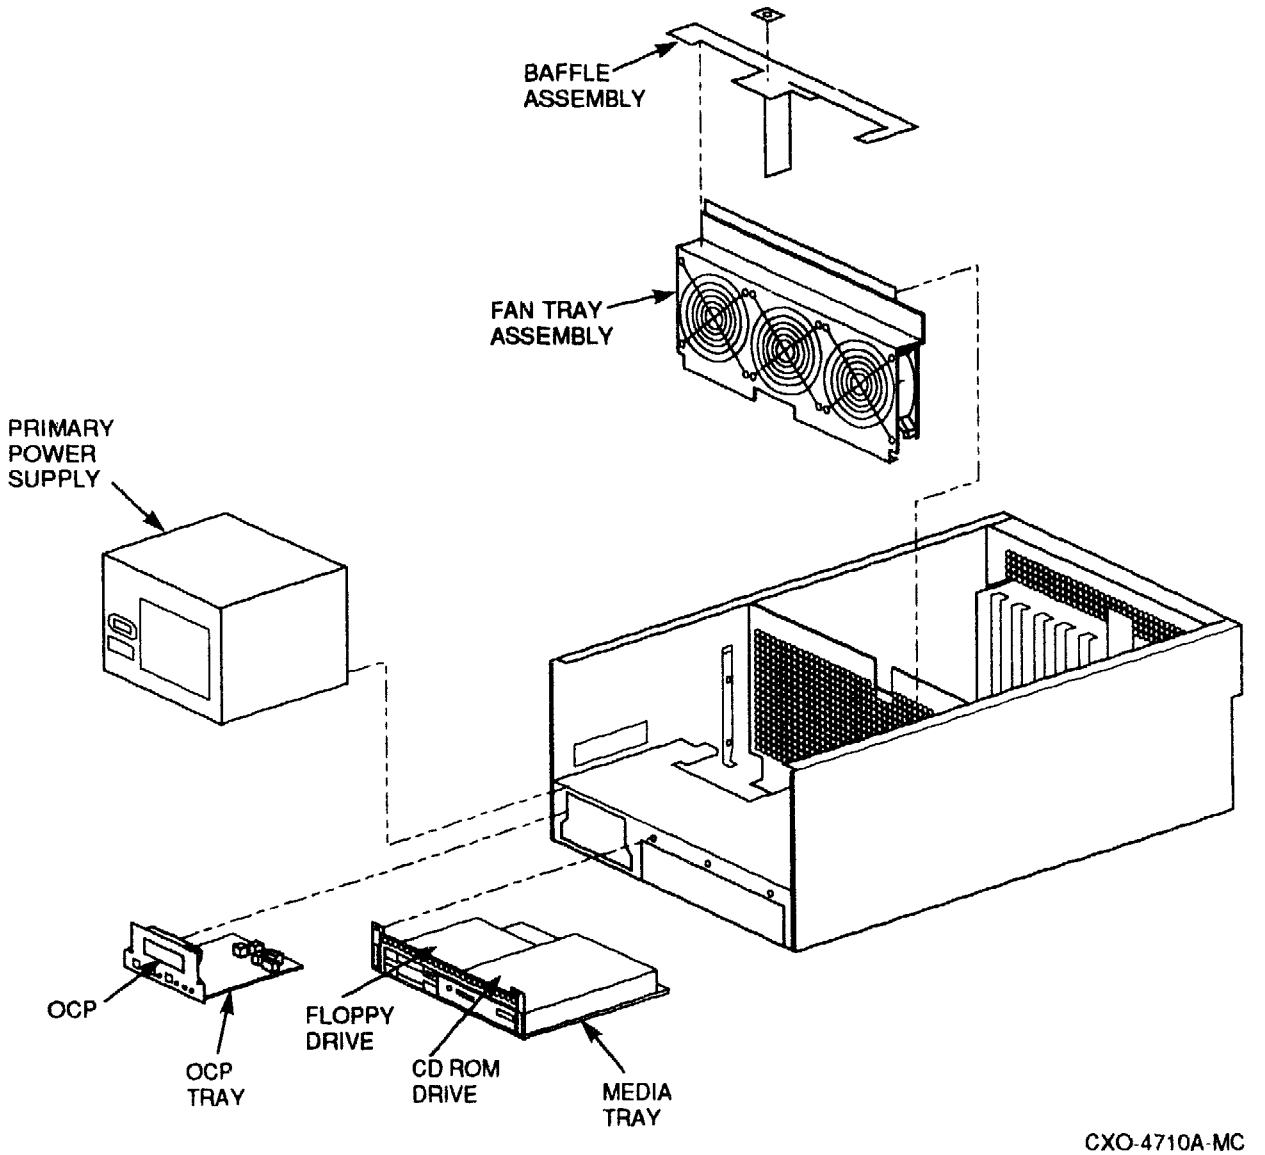
A.3 Server Processor Components Accessible from the Front

The following components can be accessed from the front of the server processor. Figure A–3 is an exploded front view of server processor components.

- OCP
- 3.5-inch floppy drive
- CD–ROM

- Power supplies
- Fan assembly

Figure A–3 Server Processor Front View—Exploded



A.3.1 Operator Control Panel Controls and Indicators

The OCP module consists of a 16-character liquid crystal display (LCD), DC power-enable switch, momentary Halt switch, momentary Reset switch, and two light-emitting diodes (LEDs), indicating DC POWER ON and SYSTEM HALT. The controls and indicators on the OCP are shown in Figure A–4 and described in Table A–1.

Figure A–4 Operator Control Panel

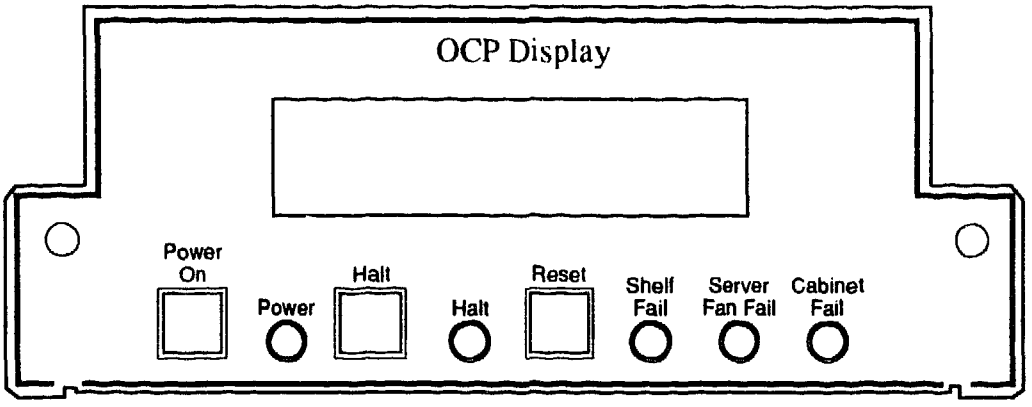


Table A–1 OCP Controls and Indicators Description

Control or Indicator	Function
Power On	Switches the system unit on and off.
Power indicator	Lights when the system unit is turned on. In a redundant system, the light displays amber if one power supply is failed.
Halt switch	Halts an OpenVMS system, returning it to console mode control.
Halt indicator	Lights when the system is halted.
Reset switch	Reinitializes the system and performs startup tests.
Shelf fail indicator	Not used.
Server fan fail indicator	Lights when any system fan is not working.
Cabinet fail indicator	Not used.
Operator control panel display	Displays diagnostic and error messages.

A.3.2 Media Tray

The media tray is a subassembly that houses both the floppy drive and CD–ROM. It is located on the front of the server processor directly underneath the redundant power supply (see Figure A–3).

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, firmware update, or loading console code.

CD–ROM Drive

The CD–ROM drive 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 the OpenVMS Alpha operating system.

A.3.3 Power Supplies

The server processor power supplies support the following two modes of operation:

- Single power supply

  This mode of operation provides all the power your system needs.

- Dual (redundant) power supply

  To increase reliability, you can add a redundant power supply. In redundant mode, the failure of one power supply does not cause the system to shut down. The current share cable connects the two power supplies so that failover occurs immediately when power to the primary power supply is lost. Figure A–3 shows the position of the primary power supply. A redundant power supply resides in the space to the right of the primary power supply in the server processor.

Voltage selection is not required. The system is intended for use at all rate ac-input voltages. Table A–2 lists the power supply ratings for systems using one or two power supplies.

Table A–2 Power Supply Ratings

Specification	Range
Voltage	100-120/220-240 volts ac
Frequency	50-60 Hz
Current	3.3/1.8 amperes (one power cord)
	1.8/1.0 amperes (two power cords)

A.3.4 Fan Assembly

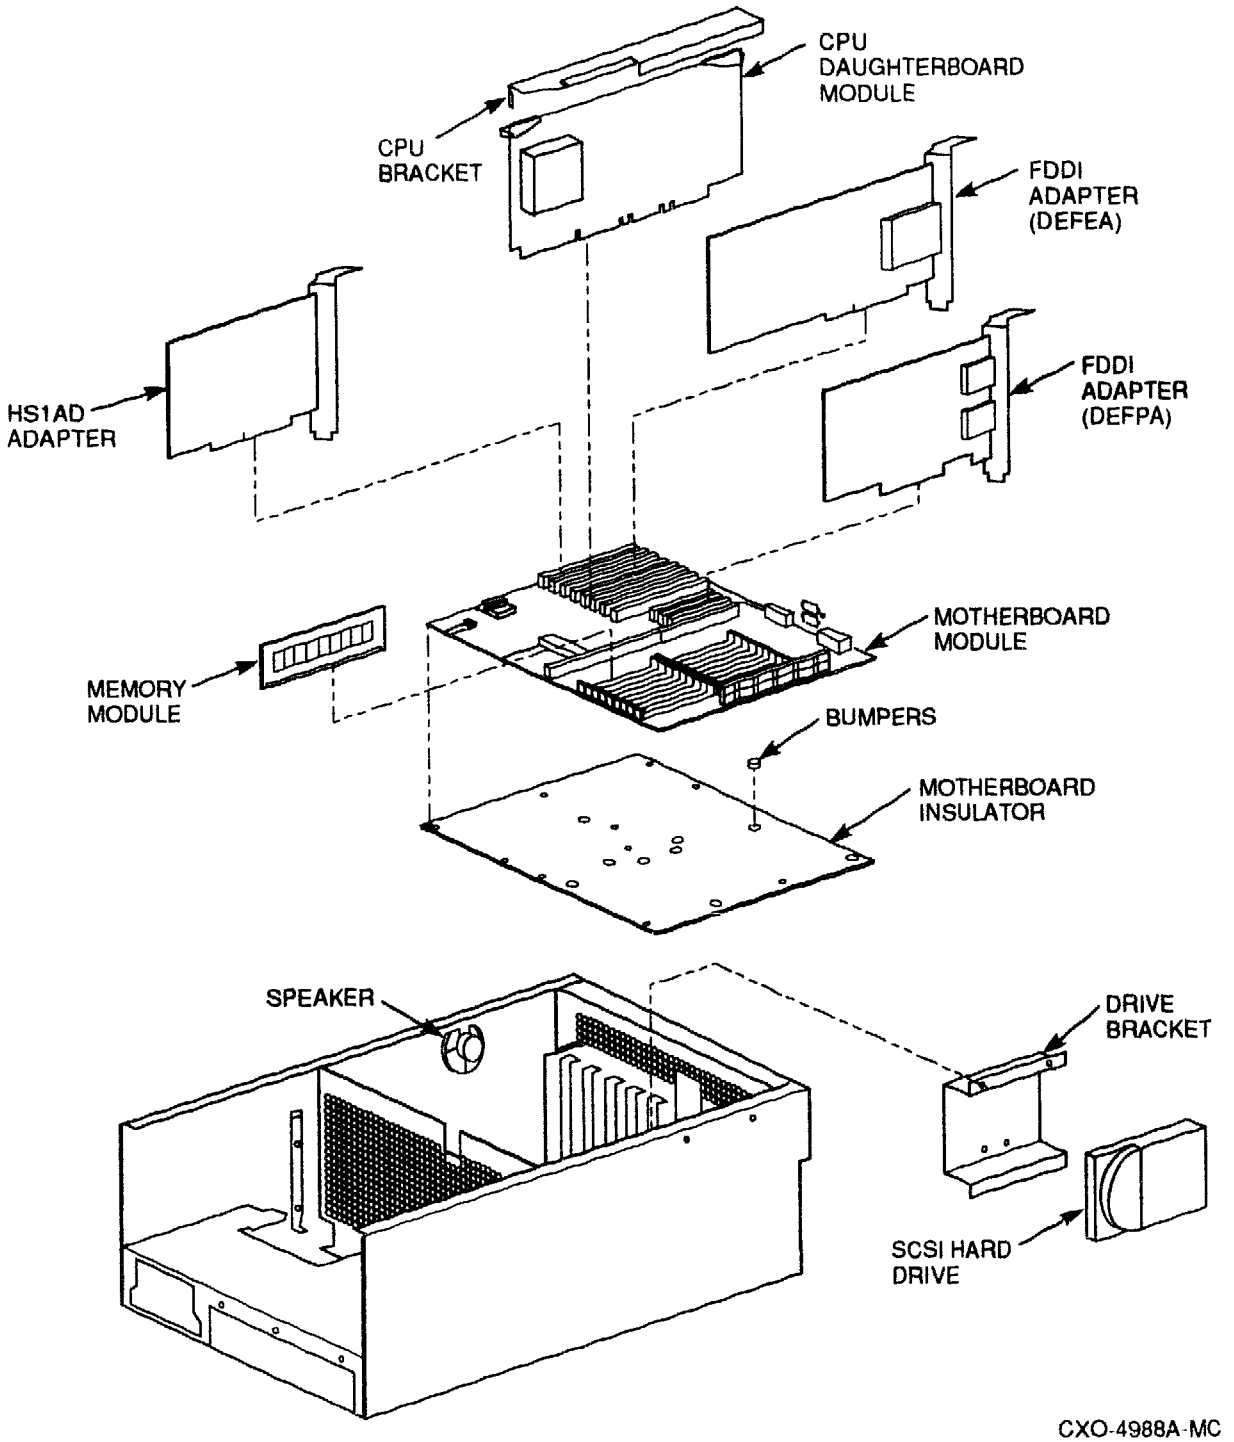
The fan assembly separates the front and rear compartments of the server processor. It provides air flow and cooling for the entire enclosure. See Figure A–3 for the position of the fan assembly.

A.4 Server Processor Components Accessible from the Rear

The following components and options can be accessed from the rear of the server processor. Figure A–5 is an exploded rear view of server processor components.

- Rear panel ports and slots
- Motherboard
- CPU daughterboard
- Memory modules
- SCSI hard drive
- HS1AD device bus adapter
- FDDI adapter (DEFEA or DEFPA module)

Figure A–5 Server Processor Rear View—Exploded



A.4.1 Rear Panel Ports and Slots

The ports and slots on the rear of the system unit are described in Table A–3.

Table A–3 Rear Panel Ports and Slots

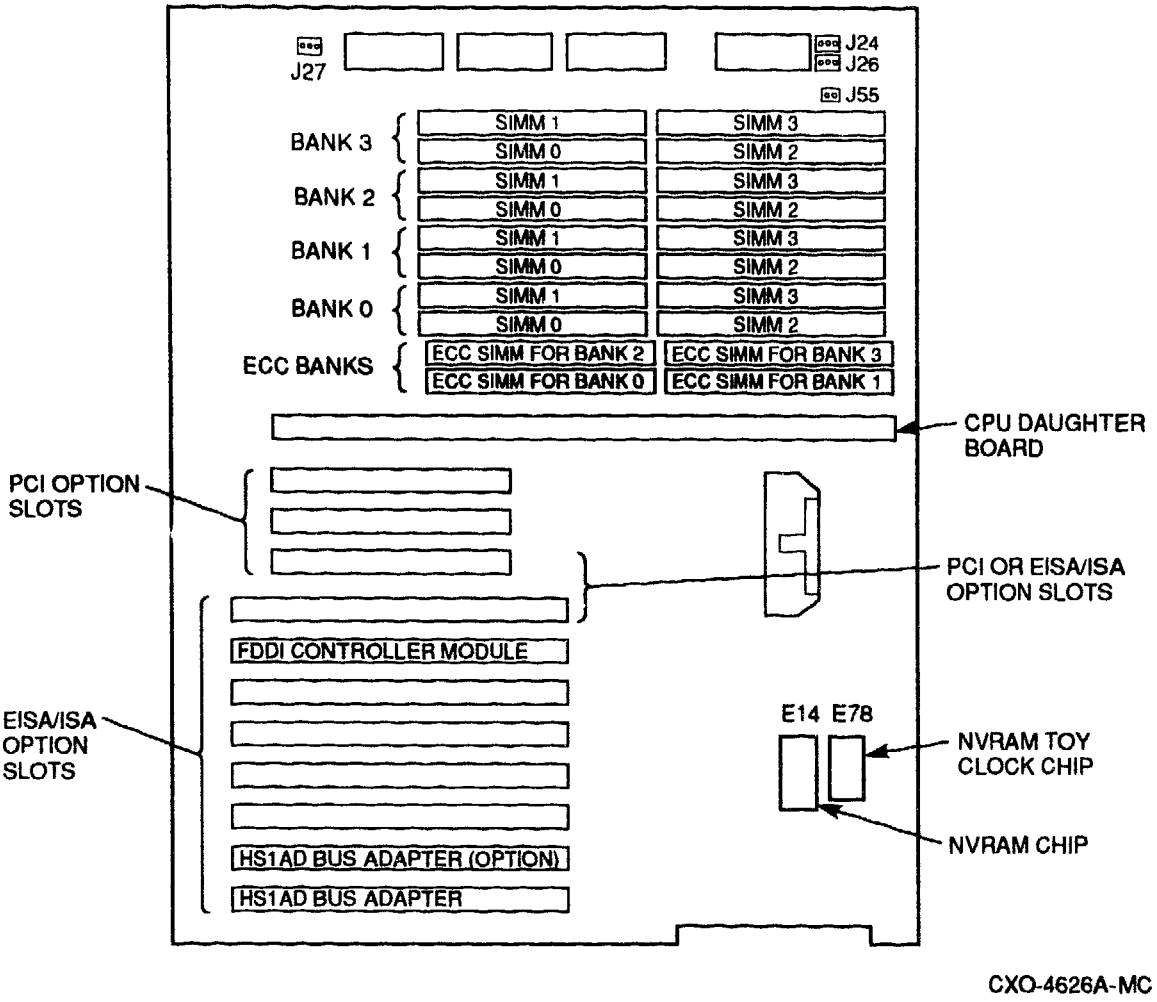
Port or Slot	Connects...
Up to eight EISA slots	Option cards for network or disk controllers
Parallel port	Parallel devices such as a printer
Serial port/terminal port (COM1)	Console terminal or serial-line peripherals such as a modem
Mouse port	PS/2-compatible mouse
VGA port	VGA monitor
Keyboard port	PS/2-compatible keyboard
Serial port (COM2)	Serial-line peripherals such as a modem

A.4.2 Motherboard

The motherboard contains the following components:

- Floppy disk controller (FDC)
- FAST SCSI–2 controller
- Integral SVGA graphics interface
- Serial ports (2)
- Parallel port
- PS/2 compatible keyboard/mouse controller
- Memory subsystem that supports up to 20 single in-line memory modules (SIMMs)
- Connection for the CPU module (daughterboard)
- Interconnects for up to eight EISA option modules
- Connection for the operator control panel (OCP)
- Interconnects to a speaker

Figure A–6 Motherboard Layout



CXO-4626A-MC

A.4.3 Daughterboard

The daughterboard contains the DECchip 21064 Alpha processor chip, cache, data and control ASICSs, serial ROM and interface, and configuration jumpers. The daughterboard is installed into two edge connectors (J1 & J2) on the motherboard.

A.4.4 Memory Modules

Each system can support from 64 MB to 512 MB of memory. Memory options consist of five SIMMs and are available in the following variations:

- HS1MM 64 MB memory option consists of five 16 MB SIMMs. An HS1*xx* server processor fully populated with 16 MB SIMMS contains 256 MB of memory capacity.
- HS2MM 128 MB memory option consists of five 32 MB SIMMs. An HS1CP server processor fully populated with 32 MB SIMMS contains 512 MB of memory capacity.

Note

When upgrading to the HS121 server processor, you must upgrade the memory of each server processor equally.

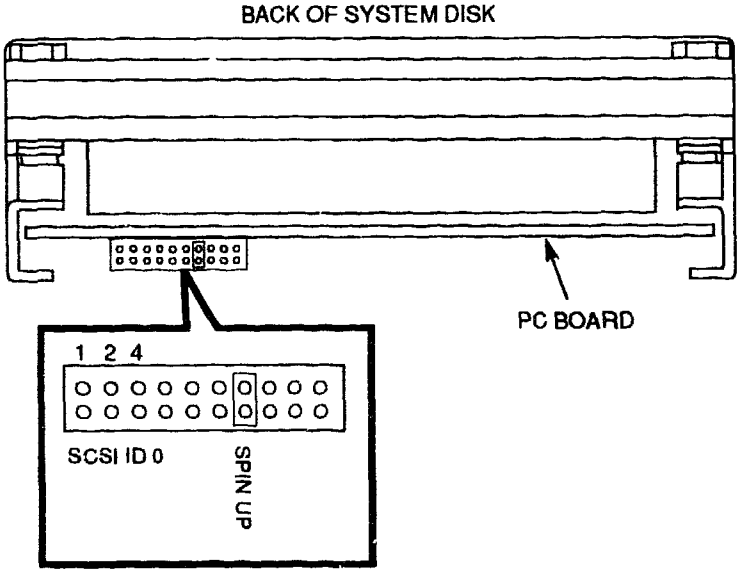
The motherboard has 20 SIMM connectors. The SIMM connectors are grouped in four memory banks (0, 1, 2, and 3) and one bank (for each bank of memory) for ECC (Error Correction Code) memory. Memory layout is show in Figure A–6.

A.4.5 System Hard Disk Drive

The SCSI hard disk drive is used as the system disk within each server processor. The system disk contains the OpenVMS Alpha operating system, the Software Customization Procedure (SCP), various relevant documents, and required layered software components.

The system disk is mounted on the side wall of the rear compartment. For a detailed view of the system disk, see Figure A–7.

Figure A–7 SCSI System Disk and Jumper Locations



LED Indicators

The disk drive has two surface-mounted LED indicators: Busy (BSY), which is green, and Fault (FLT), which is amber.

The BSY LED indicates that the drive is working on a SCSI command. It is not equivalent to the SCSI BSY signal. The FLT LED indicates a drive fault condition.

It is normal for both LEDs to light briefly when applying power.

A.4.6 HS1AD Device Bus Adapter

The HS1AD provides an internal bus for EISA-based systems, and connects a server processor with an HS1CP. Up to two HS1ADs can be installed in a single server processor, with each HS1AD configured as an end-node with a single host on a bus. Figure A–6 shows the position of the HS1AD.

A.4.7 FDDI Adapter (DEFEA or DEFPA Module)

The FDDI adapter connects the HS1CP to the FDDI ring as a single attachment station (SAS). The FDDI adapter module connects the server processor with other nodes on the VMScluster. See Figure A–6 for the position of the DEFEA or DEFPA adapter on the motherboard.

Physical Description

The adapter is a full-length IBM AT® form-factor card that plugs into the EISA bus. It is a one-card, SAS adapter with multimode optics and ANSI MIC connectors.

The major components on the adapter are as follows:

- 1 MB of packet memory
- 68000 onboard processor
- DMA control gate array
- EISA bus interface control logic
- FDDI interface chipset
- FDDI optical or TP–PMD interface
- IEEE address ROM
- Onboard, nonvolatile memory for firmware storage
- Multimode physical layer medium dependent

FDDI Adapter Onboard Diagnostics

The adapter contains onboard diagnostics that execute when power is applied. A two-color (green or red) LED on the adapter mounting bracket indicates the operating status of the adapter and its PHY port.

The adapter LEDs indicate the status of the adapter and the FDDI port. Table A–4 lists and describes possible LED states.

Table A–4 FDDI Adapter LED States

Color	Normal	Description
Green	On	On—PHY connection complete.
Green	Flashing	Flashing—PHY connection in progress (or no cable attached).
Red		On—If on after system boots, indicates broken port or Link Confidence Test (LCT) failure; retry loop.
		If on before system boots, indicates self-test failure.
Green/Red		Off–Port disabled via management, or LED or adapter broken.

A.5 The HS1CP

The HS1CP provides a connection between a host computer and an array of SCSI–2 compatible storage devices. The HS1CP can provide nonredundant connections between a server processor and up to 42 StorageWorks devices. The HS1CP supports StorageWorks magnetic, solid state, and optical disks, tapes, and media loader devices. Each HS1CP is capable of processing up to 1000 I/O requests per second, and includes a 32 MB nonvolatile write-back cache.

The HS1CP firmware resides on a PCMCIA card that ships with the StorageWorks FDDI Server. The PCMCIA card plugs into a slot located on the front of the HS1CP.

The following sections describe the subcomponents of the HS1CP and the functions of each subcomponent.

A.5.1 Maintenance Terminal Port

The HS1CP has a modified modular jack (MMJ) that can support an EIA–423 compatible maintenance terminal.

You must connect a maintenance terminal during controller installation to set initial controller parameters. During normal operation, you can use a maintenance terminal to add devices and storagesets, or to perform other storage configuration tasks.

Note

If you connect a maintenance terminal to one server processor in a dual-redundant configuration, and both server processors are functioning, you can change the configuration of both server processors.

A.5.2 Program Card

The program card is a PCMCIA standard program card device containing the firmware for operating the controller. The firmware is loaded from the program card into shared memory each time the controller initializes. Special error correction code corrects program card errors when possible.

A.5.3 Policy Processor

The policy processor consists of microprocessor hardware necessary for running the HS1CP. The heart of the policy processor is an Intel 80960CA (i960) processor chip. This processor chip runs the firmware from the program card and provides a consistent instruction execution rate of 25 MIPs. The i960 chip controls all but low-level device and host port operations.

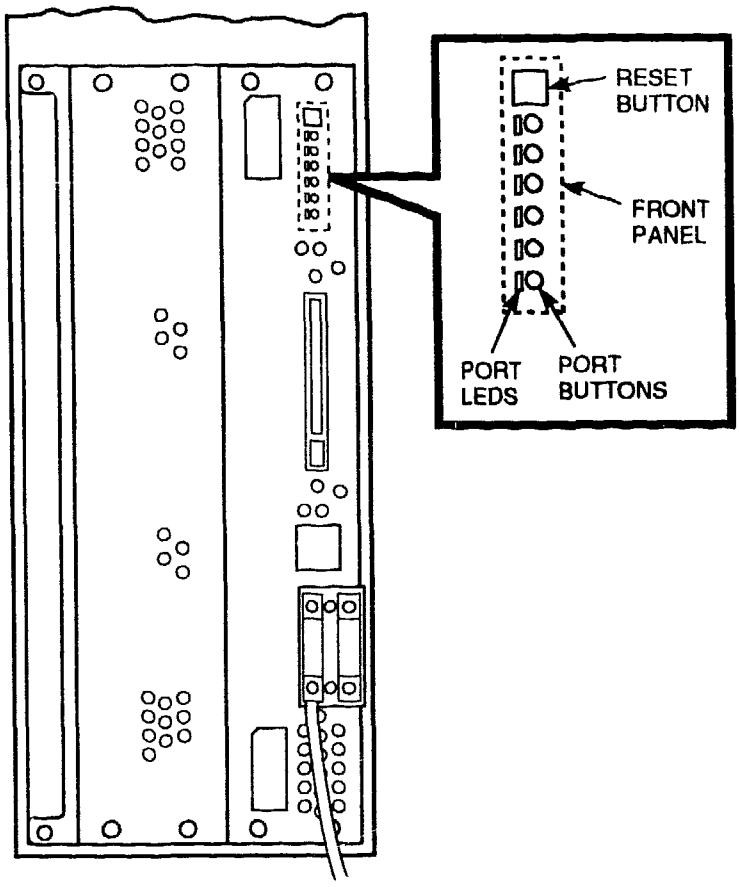
A.5.4 Diagnostic Registers

The HS1CP has two write and two read diagnostic registers. Diagnostic and functional firmware use the write diagnostic registers to manipulate controller and device operations. Certain bits in the registers activate test modes for forcing errors in the HS1CP. Other bits control the front panel LEDs. The policy processor examines the read diagnostic registers to determine the cause of an interrupt when an interrupt occurs, and at other times.

A.5.5 Front Panel

The front panel for the HS1CP consists of one Reset button with an embedded green LED and six amber port LEDs located directly next to the Reset button. The Reset button restarts (initializes) the controller. The six amber LEDs report fault and normal conditions. Figure A–8 shows the front panel on the HS1CP.

Figure A–8 HS1CP Front Panel



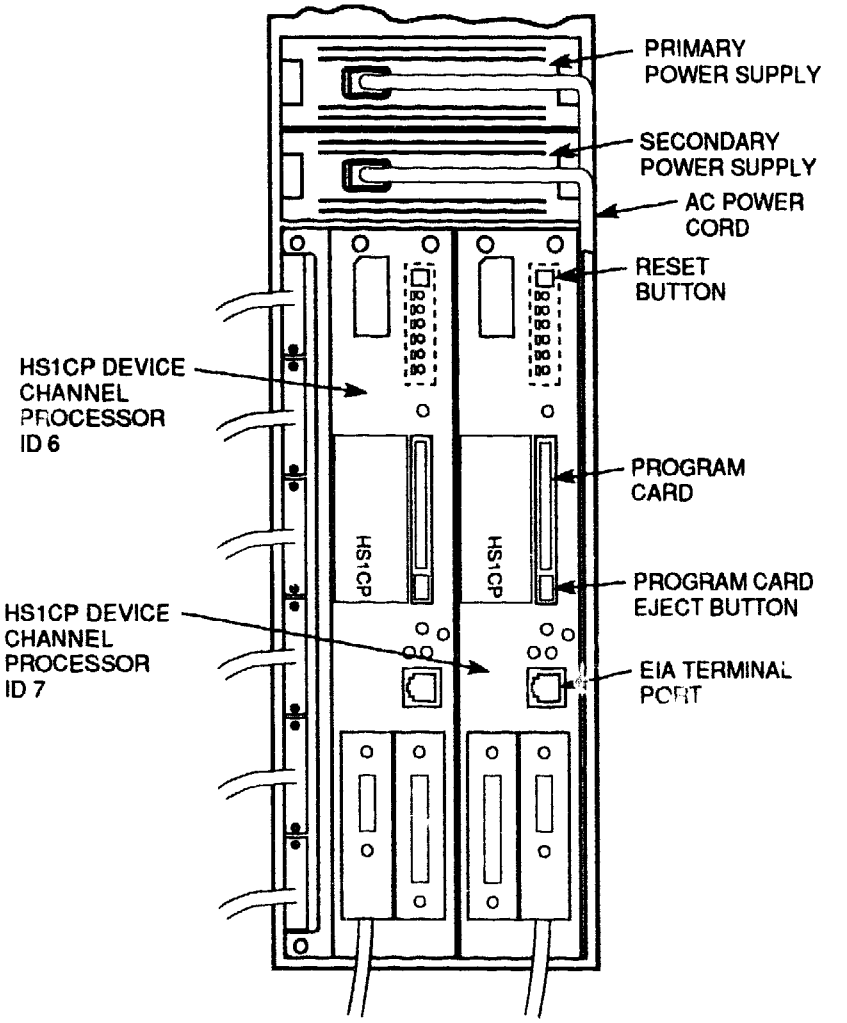
CXO-4461B-MC

See Chapter 6 for further information on using the HS1CP LEDs.

Reset Button

The HS1CP front panel includes a Reset button with an embedded green LED. The green reset LED indicates normal operation and fault conditions. Pressing the Reset button halts all activity on the controller and causes it to reboot. Figure A–9 shows the location of the Reset button on the HS1CP. See Chapter 6 for further information on using the Reset button.

Figure A–9 HS1CP Device Channel Processor Reset Button and LEDs



A.5.6 Dual Controller Port

The HS1CP has an internal serial port for communication with a second controller of the same model. The second controller must be mounted in the same controller shelf, with communication passing through the shelf backplane. A dual-redundant configuration allows one controller to take over for another (failed) controller. The takeover process is called *failover*. During failover, the surviving controller supports the SCSI–2 devices linked to the failed controller. See Chapter 5 for more information on failover.

A.5.7 Nonvolatile Memory (NVMEM)

The HS1CP has 32 KB of nonvolatile memory (NVMEM). NVMEM is implemented using battery backed up SRAM. This memory serves two purposes:

- NVMEM stores parameter and configuration information such as device and unit number assignments entered by the user and by the controller firmware.
- NVMEM stores cache module metadata, which allows the controller to check for the correct cache module.

A.5.8 Bus Exchangers

Bus exchange devices allow high-speed communication between bus devices and shared memory. One bus exchanger handles address lines while the other bus exchanger handles data lines. The bus exchangers are classified as four-way cross-point switches, which means the bus exchangers allow connections between any two ports on the switch.

A.5.9 Shared Memory

Shared memory consists of a dynamic RAM controller and arbitration engine (DRAB) gate array controller and 8 MB of associated dynamic RAM (DRAM). Shared memory uses parity-protected 9-bit error correction code (ECC) and error detection code (EDC) for improved error detection and data recovery. The shared memory also stores the controller firmware and is shared between bus devices for data structures as well as data buffers.

One portion of shared memory contains instructions for the i960 chip, firmware variables, and data structures, including the look-up table for the i960 chip. In the absence of the cache module, another portion of shared memory acts as a cache. Otherwise, this portion contains cache module context for cache look-ups when a cache module is in place.

A.5.10 Value-Added Functions

This circuit acts as the accelerator for RAID XOR operations as well as disk and tape compare operations. To support recovery of a lost member of a RAID5 storageset, the controller performs an XOR of parity data with the surviving RAIDset members' data to reconstruct the lost member.

Note

You must have the write-back cache option to execute any RAID functions.

A.5.11 Device Ports

The HS1CP SCSI–2 device ports are a combination of NCR® 53C710 SCSI port processors and SCSI transceivers. The 53C710 processors perform operations in 8-bit, single-ended normal or fast mode. The 53C710 processors execute scripts read from shared memory and under control of the policy processor.

Each SCSI–2 port can have up to six or seven attached devices depending on controller configuration (dual-redundant or nonredundant, respectively). In a dual-redundant configuration, device availability improves, because each controller has access to the other controller's devices.

A.5.12 Write-Back Cache Module

The HS1CP works with a write-back cache module on a companion board in the BA350–M shelf. The write-back cache module increases subsystem performance as well as preserves data integrity under power failure situations.

In write-back caching, data is not always written to storage and cache simultaneously (as in write-through caching). Instead, data intended for storage can remain in the cache until the optimum time to write to a device occurs. When data is suspended in this way it is referred to as unwritten cache data. A power failure when there is unwritten cache data causes the loss of the data. In RAIDset configurations, the impact of power failure is worsened by the possibility of write hole data loss as well.

For this reason, the write-back cache differs from the read cache by using onboard rechargeable batteries. The batteries power the memory to retain data when shelf backplane power is interrupted. The battery circuit automatically detects loss of shelf backplane power and switches to battery power.

Battery Discharging

During a power failure, the write-back cache batteries discharge very slowly, providing standby power to retain cache data. Digital guarantees that fully charged batteries maintain cache data during a power outage for a minimum of 100 hours.

Studies performed by utility companies note that over 99 percent of all power outages last less than 1 minute. However, once power fails for more than 3 minutes (32 MB write-back cache) or 5 minutes (16 MB write-back cache) the controller marks the cache batteries as "not fully charged" and modifies its operation upon restart to maximize data preservation as follows:

- Stripesets and disk-based units with write-back caching enabled are accessed in write-through (read cache) mode until the cache batteries are fully recharged. Once the batteries are recharged, write-back caching resumes.

- RAIDsets and mirrorsets are available either in write-through mode, or unavailable until the batteries are fully recharged from a long power failure, depending on the cache policy set for the controller. See Chapter 6 for a full explanation of write-back cache battery policy.

A battery failure in one cache module causes RAIDsets to fail over to the companion cache in a dual-redundant configuration, provided the other cache's batteries are fully charged. Although this feature does not help during power outages (because both controllers are down), it keeps your mirrorsets and RAIDsets on line if one battery becomes faulty.

Battery Charging

Under normal conditions (power restored), a battery charge circuit senses the battery voltage and automatically activates a charger to achieve and maintain full battery charge. The relationship of power down time versus battery recharge time is given by the following two ratios:

16 MB write-back cache—24:1 (power down time versus time to fully charge)
32 MB write-back cache—12:1

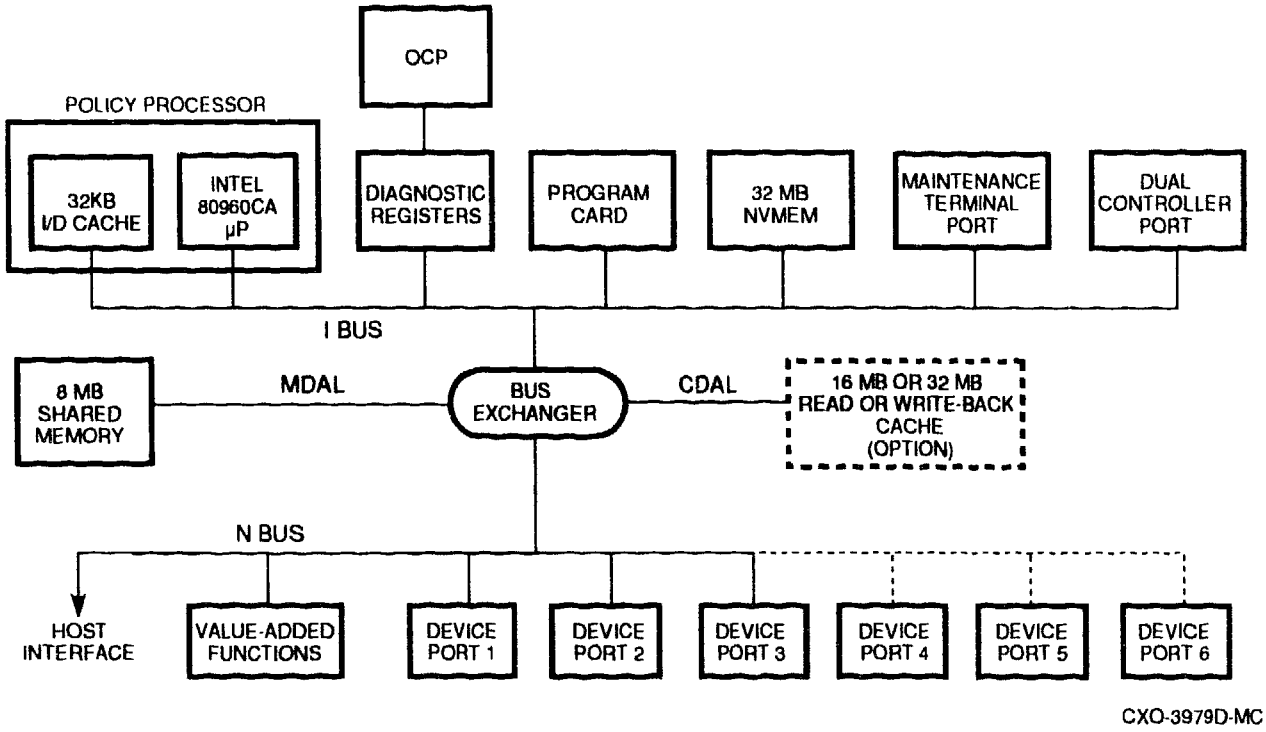
For example, a two hour power outage results in approximately 10 minutes of recharge time for a 32 MB write-back cache. The recharge time is a safety precaution, taken to preserve the integrity of the controller subsystem for any and all power outages that can occur thereafter.

Note

Digital recommends replacing the write-back cache batteries at 5 year intervals.

Figure A–10 shows a block diagram of the HS1CP hardware.

Figure A–10 Common Hardware Block Diagram



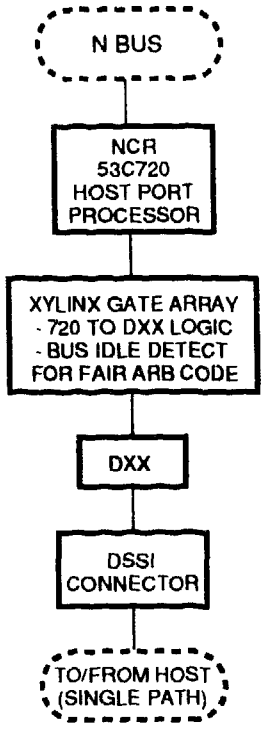
A.5.13 Host Interface

The following sections describe the host interface hardware for the HS1CP:

DSSI Interface

Figure A–11 shows a block diagram of the DSSI host interface hardware.

Figure A–11 DSSI Host Interface Hardware Block Diagram

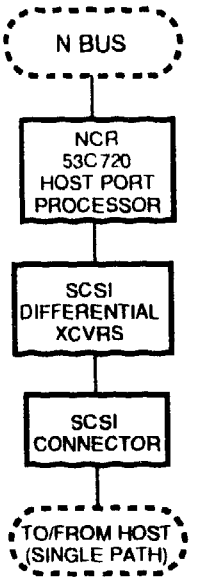


The DSSI interface consists of an NCR 53C720 chip plus specific DSSI logic and transceivers. The NCR 53C720 chip reads and runs scripts from controller shared memory to perform command and DMA operations on the DSSI interface. The policy processor sets up and maintains the operation of the NCR 53C720 chip.

SCSI–2 Interface

Figure A–12 shows a block diagram of the SCSI–2 host interface hardware.

Figure A–12 SCSI–2 Host Interface Hardware Block Diagram



The controller interfaces with two types of host SCSI buses:

- A fast, wide, differential 16-bit SCSI bus (FWD SCSI)
- A fast, differential 8-bit SCSI bus (FD SCSI)

The hardware consists of the NCR 53C720 chip and transceivers, and functions in much the same way as the DSSI interface (refer to Section A.5.13).

CAUTION

Do not connect the HS1CP host trilink connector accidentally to a SCSI bus. Connecting to the wrong bus causes equipment damage.

A.5.14 Firmware

The controller firmware consists of functional code, diagnostics, utilities, and exercisers.

The following sections describe each function and service of the firmware.

A.5.14.1 Core Functions

HS1CP operating firmware provides the following core functions, in the order they are executed after turning on the controller:

Tests and Diagnostics

Controller tests and diagnostics are integrated in a controller self-test procedure performed when the controller is switched on. The output of the self-test procedure is a simple go/nogo status of the controller subsystem. The self-test procedure includes a test of the cache module.

Init Functions

After successful completion of the tests and diagnostics, the controller data structures are initialized and memory is allocated for controller functions.

Executive Functions

Firmware executive (EXEC) functions act as the operating system kernel for the HS1CP. EXEC functions control firmware execution with respect to interrupts, thread control, queuing support, timers, and so forth. The EXEC functions establish the controller environment as a nonpreemptive interrupt-driven process.

A.5.14.2 Host Interconnect Functions

HS1CP operating firmware supports two different host interconnections: DSSI and SCSI. The following list identifies the protocols used for host access of controller storage:

- **DSSI**—SCS; MSCP and/or TMSCP and DUP
- **SCSI**—SCSI–2 protocol with SCSI passthrough software to the CLI, tagged command queuing on the host and device side, and mode select/sense support for SCSI

A.5.14.3 Operator Interface and Subsystem Management Functions

The operator interface and subsystem management functions support the user interface, subsystem management, subsystem verification, and error logging/fault management. These functions are presented in the following sections.

Command Line Interpreter

The command line interpreter (CLI) is the primary user interface for the HS1CP. You can access the CLI through the maintenance terminal on the front bulkhead of the HS1CP. The CLI contains firmware for responding to most management functions plus local program execution. Briefly, the CLI provides the following two types of commands:

- SET/SHOW commands for the controller itself.

  This command includes setting and showing the controller ID, name, path controls, and other vital information.

- Configuration commands to add/delete devices, storagesets, and logical units.

See Appendix E for a complete list of CLI commands.

Diagnostic Utility Protocol

Diagnostics and Utilities Protocol (DUP) from the host is supported over the DSSI bus. DUP allows you to access the CLI and local programs through a remote host "virtual" terminal in much the same way as using a maintenance terminal. See Chapter 5 for more information on DUP.

Local Programs

There are several local exercisers and utilities available for controller subsystem management/verification, as follows:

- **DILX** and **TILX** are disk and tape exercisers that enable you to test and verify operation of the controller with attached SCSI–2 storage under a high or low I/O load. These utilities place the load on the controller, bypassing the host port. Chapter 7 provides a full description of DILX and TILX.
- The **VTDPY** utility enables you to display current controller state and performance data, including processor utilization, host port activity and status, device state, logical unit state, and cache and I/O performance. See Chapter 7 for detailed information on this utility.

- The **controller warm swap** procedure (C_SWAP) enables you to remove and replace one controller safely in a dual-redundant configuration. When you perform a controller warm swap, you are replacing a controller in the most transparent method available to the controller subsystem. There is minimal system and device impact, as explained in Chapter 8.

  You also can use the warm swap procedure to upgrade from a single controller to a dual-redundant configuration, and to add or upgrade a cache module.

- The **Firmware Licensing System** (FLS) allows you to activate licensed features within HS operating firmware. Any feature turned on without entering a valid license code (key) causes a message to be sent to the host event log and to a connected terminal every hour. These notifications stop when a valid license key is entered. See Chapter 7 for more information on FLS.

- The **Configuration menu** (CFMENU) allows you to configure quickly the storage devices attached to the controller. CFMENU presents configuration commands normally entered at the CLI in a menu-driven format. See Chapter 7 for more information on CFMENU.

- The **Code Patch** utility (CLCP) allows you to install new firmware or repair small firmware areas without installing a new program card. Updated firmware and patches become active once you restart the controller. See Chapter 7 for more information on CLCP.

- The **Fault Management Utility** (FMU) allows you to control display of some spontaneous event, controller last failure, and memory system failure information. You also can use FMU to view recent last failure and event messages. See Chapter 7 for more information on FMU.

- The **Configure** utility (CONFIG) checks the SCSI device ports for any device not previously added. You can use this utility to add and name these devices. See Chapter 7 for more information on the CONFIG utility.

- The **Clone** utility (CLONE) creates a duplicate of any stripeset, mirrorset, or single-disk unit. You can use this utility to perform a backup of a unit (for example, a stripeset) without significantly impacting performance for users by first cloning the unit and then performing the backup from the clone. See Chapter 7 for more information on the CLONE utility.

Event Reporting and Fault Management

Event Reporting and Fault Management are integrated functions that control reporting of significant events detected during controller operation. Fault Management firmware recommends repair actions when faults are detected. See Chapter 7 for more information on these subjects.

A.5.14.4 Device Services Functions

SCSI–2 device service firmware includes device port drivers, support for disks, tapes, tape loaders, CD–ROMs, and other devices on one controller, and physical device addressing and access. Device service consists of normal functions such as read and write operations, plus error recovery code. It also contains firmware for controlling and observing the BA350–S shelf and StorageWorks building blocks (SBBs), such as LED, power, and blower monitoring. Specific features include:

- Normal SCSI–2, 8-bit, single-ended support.
- FAST, synchronous, 8-bit, single-ended device support.
- Tagged queuing for SCSI–2 devices.

- Support of specified devices per HS1CP operating firmware release. Refer to your HS1CP operating firmware release notes for a list of supported devices.
- Mixed disk and tape support. You can mix disk and tape storage on one controller. Furthermore, you can place disks and tapes together on any of the controller's SCSI–2 ports.
- Device warm swap. You can remove and replace most devices without taking the subsystem offline (see Chapter 8). See your firmware release notes for any restrictions on devices that cannot be warm swapped.
- Device shelf and SBB observation and control. This service monitors SHELF_OK signals and alerts you of blower and power supply failures. This firmware also controls the fault LEDs on the SBBs for use in warm swap procedures and identifying device failures or configuration mismatches.
- Device error recovery. This service performs error recovery and read and write retries directly, making several attempts to serve data to and from the host before declaring an unrecoverable error or marking a device as failed.
- Command disks. Command disks allow jukeboxes and other generic SCSI devices to be attached to MSCP-based hosts. See Chapter 5 for more information on command disks.

A.5.14.5 Value-Added Functions

HS1CP operating firmware contains value-added functions to enhance availability, performance, subsystem management and maintenance, and connectivity features of the controller subsystem. The failover function is presented in the following section.

Failover

The two controllers in a dual-redundant configuration are linked together. The controllers exchange status signals and configuration information. When one controller fails, the surviving controller takes over service to the failed controller's devices and units. See Chapter 5 for more information on failover.

A.5.14.6 Addressing Storage Within the Subsystem

This section provides an overview of how storage is addressed in a controller subsystem. Storage is seen in two different ways, depending on your perspective:

- At the physical device level from the HS1CP SCSI device interface
- At the virtual device level from the host interface

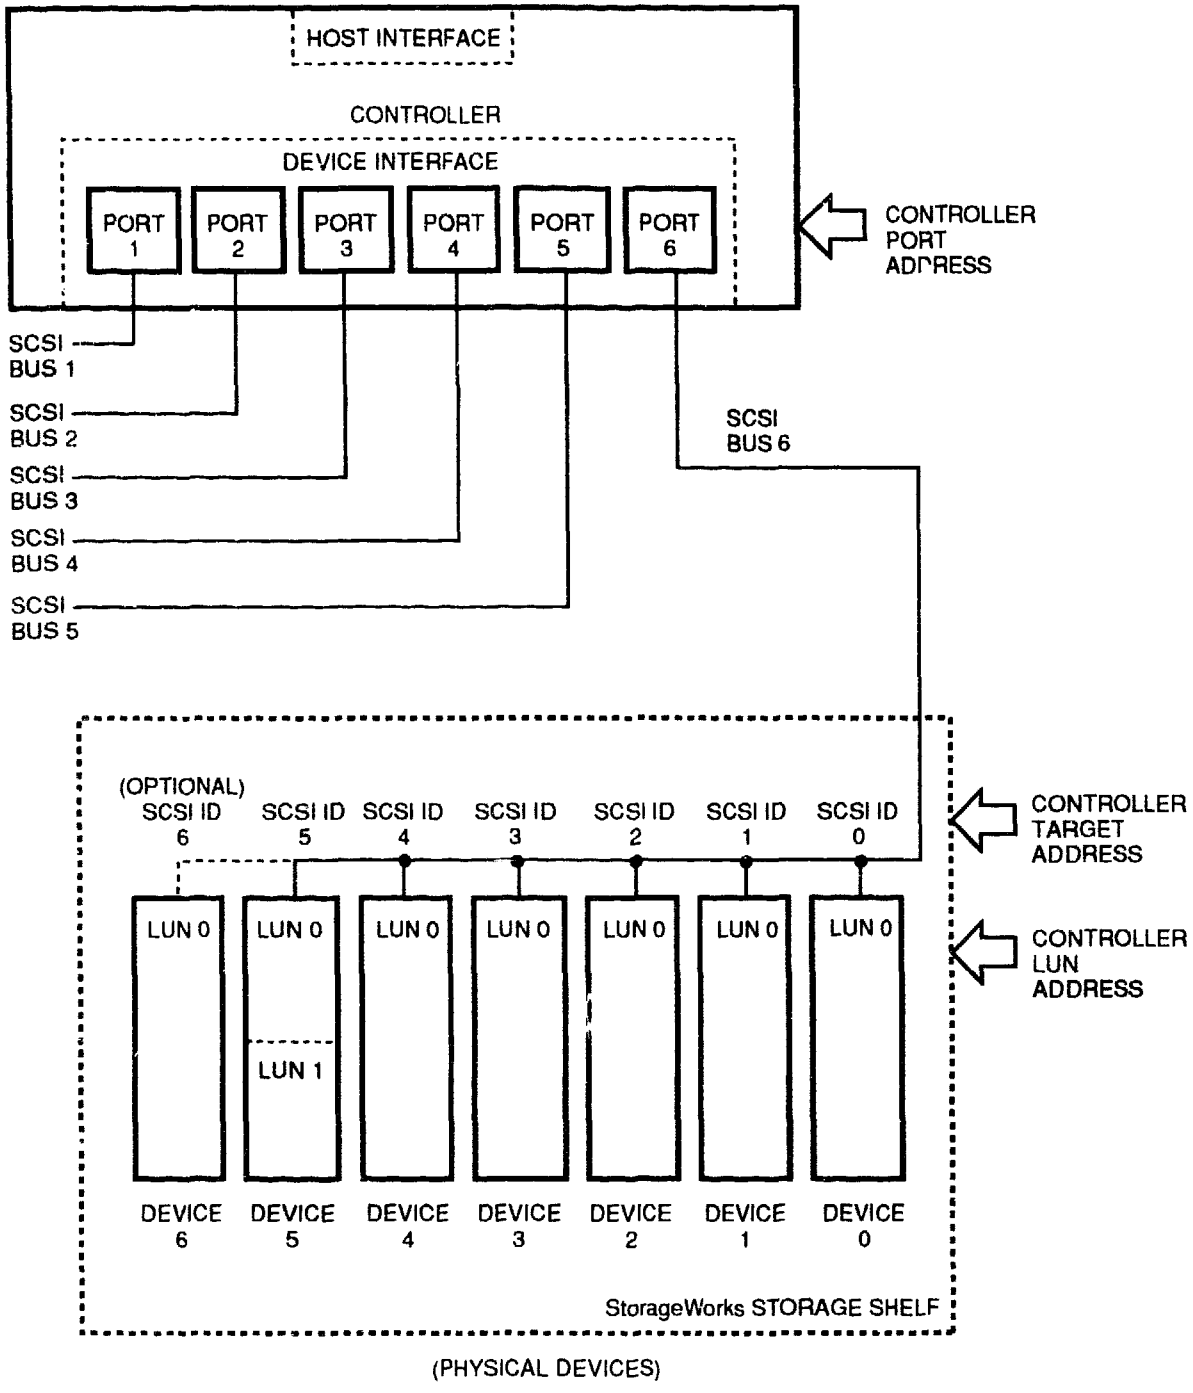
The following sections describe both levels of storage addressing.

Controller Storage Addressing

Figure A–13 shows a typical physical storage device interface for an HS1CP. Each of the HS1CP's six device ports supports a SCSI bus connected with up to six devices. The devices typically reside in a StorageWorks BA350–S storage shelf.

HS1CPs support both single-LUN and multi-LUN devices. Most devices, including disks drives, are single-LUN devices and use the default address of LUN 0. Certain devices, such as some tape loaders, require multiple LUNs so that the host operating system can address both the loader mechanism and the drive as separate entities.

Figure A–13 Controller Storage Addressing



CXO-4596A-MC

Controller Port/Target/LUN Addressing Controller port/target/LUN (PTL) addressing is the process by which the controller selects storage space within a specific, physical storage device. The process takes place in three steps:

1. The port selection—The controller selects the SCSI bus port connected to a particular device.
2. The target selection—The controller selects the device's SCSI ID (that is, the target) on that port.

3. The LUN selection—The controller selects the desired LUN within that physical device, usually LUN 0.

Note that controller PTL addressing always is tied to a physical storage device.

Host Storage Addressing

A typical host device interface consists of a number of host ports, each connected to a bus containing devices. From the host perspective, the HS1CP is one of these devices.

To support certain high-level storage subsystem functions such as RAID, the HS1CP presents the devices (from Figure A–13) to the host as *host logical units*. A host logical unit can consist of a single device or a storageset comprised of more than one physical device. For this reason, a host logical unit is often called a virtual device. You configure host logical units using the CLI.

B

Field Replaceable Unit Part Numbers

This appendix lists the field replaceable units (FRUs), required tools and equipment, and related FRUs.

B.1 Required Tools

The following tools may be required to perform maintenance activities on the FDDI Server:

- Portable antistatic kit
- ESD mat
- Antistatic bag
- Allen wrench (3/8-inch)
- Allen wrench (5/32-inch)
- Flathead screwdriver, 3/8-inch
- Flathead screwdriver, 1/8-inch
- Phillips screwdriver, #2
- Phillips screwdriver, #0
- Nut driver, 5/16-inch
- Wire cutters, diagonal
- MMJ adapter
- Maintenance terminal

B.2 Field Replaceable Units

The following component parts may needed to be replaced during the course of normal maintenance activities.

Table B–1 FDDI Server Replaceable Parts

Item	Description	Part Number
1	Device Bus Adapter (HS1AD)	54–23236–01
2	Internal Bus Cable	BC29S–09
3	Cable Distribution Unit (CDU)	SW8XP
4	Shelf, controller (BA350–M)	BA350–MB
5	Cable, internal bus, 9 foot	BC295–09
6	HS1CP device channel processor with 32 MB Cache Module	HS1CP–AF
7	Adapter, 9-pin, MMJ (H8571–J)	12–23599–08
8	Terminator, internal bus	12–31281–01
9	Slide, chassis, solid bearing, 2-way	12–43233–04
10	Harness wiring, 10A, (Gray), 7 foot	17–03718–05
11	Harness wiring, 10A, (Black), 7 foot	17–03718–06
12	Adapter, EISA to FDDI	DEFEA–AA
12	Adapter, PCI to FDDI	DEFPA–AA
13	Bracket, slide mounting	74–48875–01
14	Baffle, horizontal air	74–49860–01
15	Baffle, vertical air	74–49861–01
16	Bracket, shipping, right hand	74–49862–01
17	Bracket, shipping, left hand	74–49862–02
18	StorageWorks FDDI Server V*x.x* CD–ROM (AG–QGPM*x*–BE)	AG–QGPMA–BB
20	EISA Configuration Utility (ECU)V*x.x* Diskette, 3.5-inch (AK–Q2CRG–CA)	AK–Q2CRG–CA

Table B–2 CDU Parts List

Item	Description	Part Number
1	Distribution unit, cable, 120/208V, 60 Hz	SW8XP–AA
2	Cord, ac power, gray, 5 feet	17–03718–03
3	Cord, ac power, gray, 7 feet	17–03718–05
4	Cord, ac power, gray, 9 feet	17–03718–07
5	Internal Bus Cable	BC29S–09

Table B–3 BA350–S Shelf Assembly

Item	Description	Part Number
1	BA350–S shelf assembly	70–29760–01
2	BA350 modular storage shelf subsys (Users Guide)	EK–BA350–UG
3	StorageWorks, power supply (150 watt)	70–29764–05
4	BA350 modular storage shelf SBB (Users Guide)	EK–SBB35–UG
5	Label, bar, code, thermal paper	36–27585–01
6	Bag, poly reclose 2 mil thk clear	99–06557–30
7	Envelope, polyethylene, front load	99–07034–02
8	Rack mounting kit for BA350 in met	BA35X–RD
9	Bracket, shelf mounting	74–47506–01
10	Kit, shelf mounting - univ. assemb	70–30092–01
11	Bracket, locking rear	74–46439–01
12	Bracket, locking front	74–46441–01
13	Bracket, stop	74–46440–01
14	350 SBB PS (A/C)	70–29764–02

Table B–4 BA350–M Shelf Assembly

Item	Description	Part Number
1	BA350–M shelf assembly	70–29760–03
2	StorageWorks, power supply (150 watt)	70–29764–05
3	Bezel, blank	74–45644–01
4	BA350-M modular storage shelf use (Users Guide)	EK–350MA–UG
5	Bulkhead, shelf-M (blank)	74–47029–02
6	Bracket, shelf mounting	74–46442–01
7	Nut, u-nut	90–07786–00
8	Bracket, locking rear	74–46439–01
9	Bracket, locking front	74–46441–01
10	Bracket, stop	74–46440–01
11	Screw, thrl	90–00063–39
12	StorageWorks solutions screw/mount	EK–35XR[illegible]–IG
13	Nuts, keps	90–06565–00

Table B–5 Server Processor Replaceable Parts

Item	Description	Part Number
1	OCP Tray Assembly	70–31767–01
2	Fan Tray Assembly	70–31769–01
3	3.5" Diskette Drive	RX26–AA 2.8
4	Half Height Dual Speed 600MB Drive	RDD43–AA
5	2GB Formatted Capacity SCSI Fixed Disk	RZ28–EA
6	400 Watt Fanless Power Supply	30–41976–02
7	Mikasa Alpha Processor Daughterboard	54–23297–01
8	Mikasa System Board	54–23308–01
9	2.5" Speaker, 80 OHM Round	12–15050–00
10	Harness Assembly—2 wire	17–03655–03
11	Wire Harness Assy—3 wire	17–03964–01
12	Wire Harness Assy—24 wire	17–03966–01
13	Wire Harness Assy—20 wire	17–03968–01
14	Cable Assy, Flat—34 wire	17–03970–01
15	Cable Assy, Flat—10 wire	17–03971–01
16	Wire Harness Assy—12 wire	17–04069–01
17	Wire Harness Assy—21 wire	17–04070–01
18	Cable Assy, 50 Flat SCSI—50 wire	17–04094–01

B.3 Server Processor Cabling Interconnection

Figure B–1 shows a cable connection diagram for the server processor. Figure B–2 shows a cable routing diagram for the server processor.

Figure B–1 Cable Connection Diagram

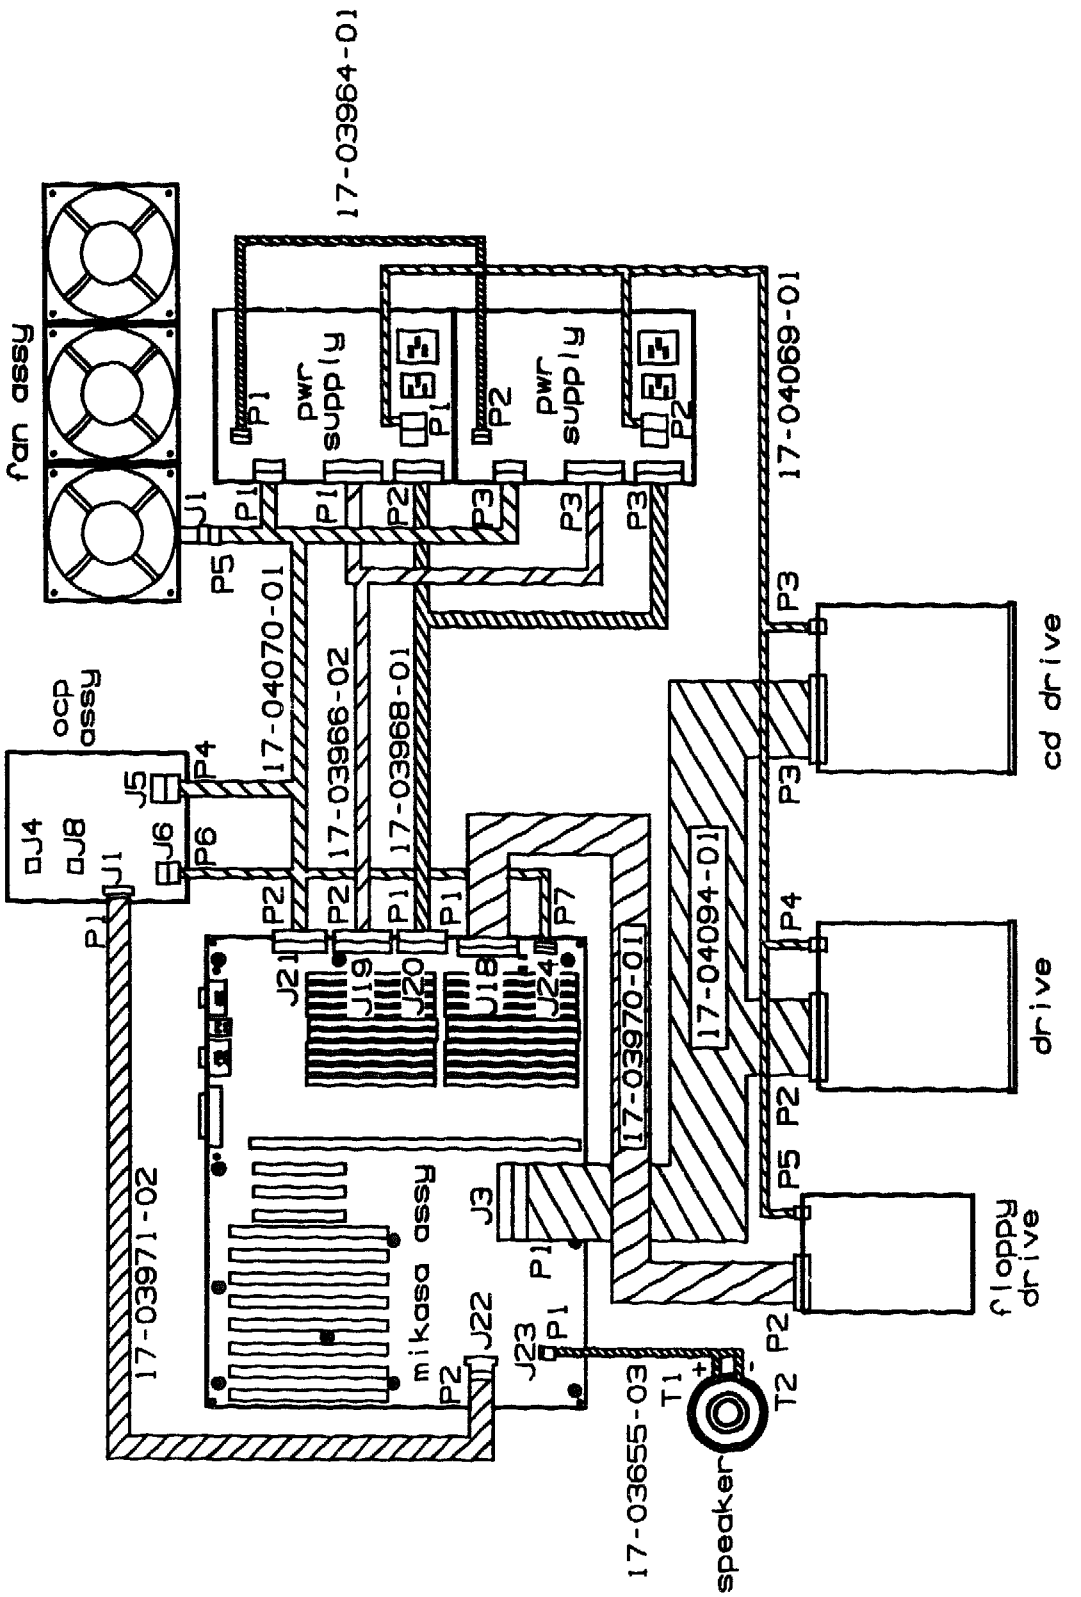
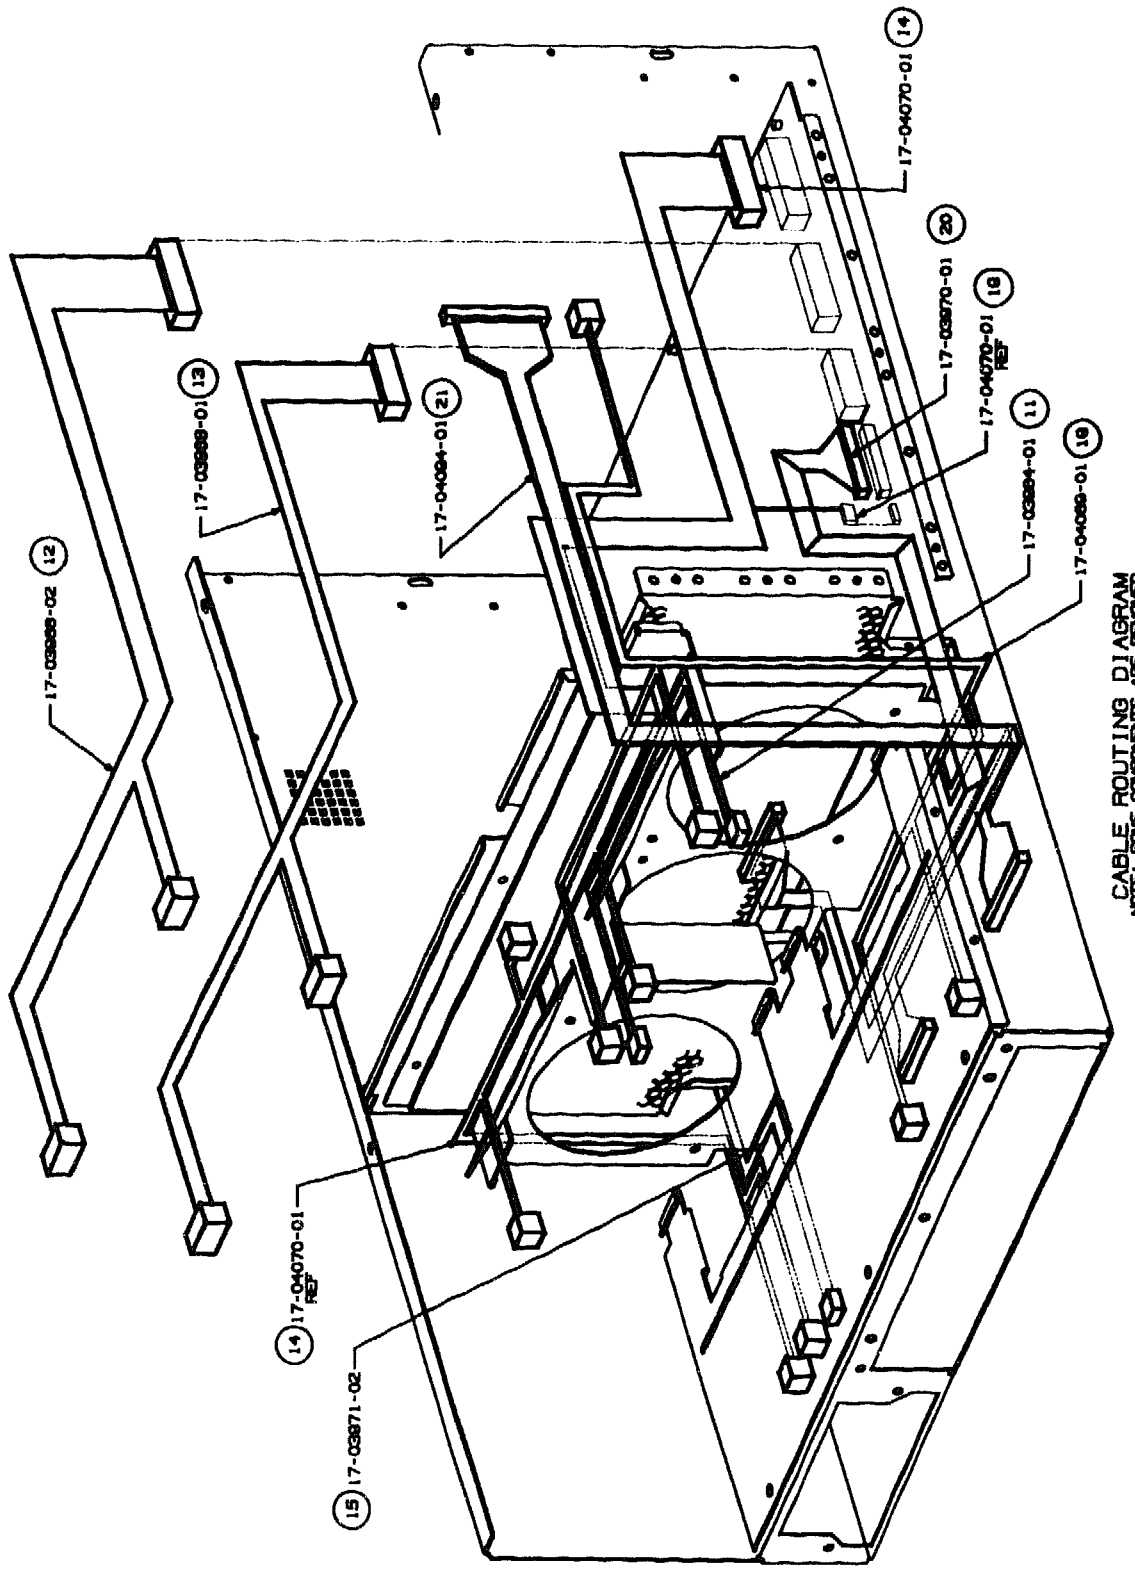


Figure B–2 Cable Routing Diagram



C

HS1CP Error Logging

This appendix details errors the HS1CP reports in its host error logs under the OpenVMS operating system, as well as how to extract the information from the logs.

Note

Host error log translations are correct as of the date of publication of this manual. However, log information may change with firmware updates. Refer to your firmware release notes for error log information updates.

C.1 Reading an HS1CP Error Log (ERF)

To understand the error logs, use the following guidelines:

- Each error log contains an "MSLG$B_FORMAT" field (in the upper portion of the log), plus a "CONTROLLER DEPENDENT INFORMATION" area (in the lower portion of the log). "CONTROLLER DEPENDENT INFORMATION" will vary according to the "MSLG$B_FORMAT" field.

  Example C–1 shows an example of an ERF translated host error log (a Disk Transfer Event log). See Example C–1 to find "MSLG$B_FORMAT" and "CONTROLLER DEPENDENT INFORMATION."

- The key to interpreting error logs is a 32 bit instance code located in the "CONTROLLER DEPENDENT INFORMATION" area. The instance code uniquely identifies the following:
  - The error or condition
  - The component reporting the condition
  - The recommended repair action
  - The threshold when the repair action should be taken

Note

The instance code is the single, most important part of interpreting the error log.

Example C–1 Disk Transfer Error Event Log

```
V A X / V M S        SYSTEM ERROR REPORT        COMPILED 16-MAR-1993 11:05:04
                                                                 PAGE 146.

******************************** ENTRY      12. ********************************
ERROR SEQUENCE 2832.                              LOGGED ON:        SID 05903914
DATE/TIME 16-MAR-1993 10:27:58.95                                 SYS_TYPE 00000000
SYSTEM UPTIME: 4 DAYS 02:11:34
SCS NODE: CNOTE                                                   VAX/VMS V5.5-2

ERL$LOGMESSAGE ENTRY  KA825  HW REV# B  PATCH REV# 28.  UCODE REV# 20.
                      BI NODE # 2.

I/O SUB-SYSTEM, UNIT _FRED$DUA115:

      MESSAGE TYPE        0001
                                        DISK MSCP MESSAGE

      MSLG$L_CMD_REF  9DB30013
      MSLG$W_UNIT         0073
                                        UNIT #115.

      MSLG$W_SEQ_NUM      0002
                                        SEQUENCE #2.

      MSLG$B_FORMAT         02
                                        DISK TRANSFER LOG

      MSLG$B_FLAGS          00
                                        UNRECOVERABLE ERROR

      MSLG$W_EVENT        000B
                                        DRIVE ERROR
                                        UNKNOWN SUBCODE #0000(X)

      MSLG$Q_CNT_ID   00134534
                      01280001
                                        UNIQUE IDENTIFIER, 000100134534(X)
                                        MASS STORAGE CONTROLLER
                                        MODEL = 40.

      MSLG$B_CNT_SVR        FF
                                        CONTROLLER SOFTWARE VERSION #255.

      MSLG$B_CNT_HVR        00
                                        CONTROLLER HARDWARE REVISION #0.

      MSLG$W_MULT_UNT     0005
      MSLG$Q_UNIT_ID  00000001
                      02FF0000
                                        UNIQUE IDENTIFIER, 000000000001(X)
                                        DISK CLASS DEVICE (166)
                                        MODEL = 255.

      MSLG$B_UNIT_SVR       0B
                                        UNIT SOFTWARE VERSION #11.

      MSLG$B_UNIT_HVR       0C
                                        UNIT HARDWARE REVISION #12.

      MSLG$B_LEVEL          01
      MSLG$B_RETRY          00
      MSLG$L_VOL_SER  00001492
                                        VOLUME SERIAL #5266.

      MSLG$L_HDR_CODE 000659B6
                                        LOGICAL BLOCK #416182.
                                        GOOD LOGICAL SECTOR
```

(continued on next page)

Example C–1 (Cont.) Disk Transfer Error Event Log

```
CONTROLLER DEPENDENT INFORMATION
      LONGWORD 1.     03094002
                                  /.@../
      LONGWORD 2.     00003C51
                                  /Q<../
      LONGWORD 3.     00000000
                                  /..../
      LONGWORD 4.     000016D4
                                  /Ô.../
      LONGWORD 5.     00000000
                                  /..../
      LONGWORD 6.     00030002
                                  /..../
      LONGWORD 7.     56415246
                                  /CNOT/
      LONGWORD 8.     20205355
                                  /E  /
      LONGWORD 9.     00000501
                                  /..../
      LONGWORD 10.    36325A52
                                  /RZ26/
      LONGWORD 11.    20202020
                                  /    /
      LONGWORD 12.    29432820
                                  / (C)/
      LONGWORD 13.    43454420
                                  / DEC/
      LONGWORD 14.    20202020
                                  /    /
      LONGWORD 15.    31202020
                                  /   1/
      LONGWORD 16.    i00F0002A
                                  /*.ð./
      LONGWORD 17.    59060004
                                  /...Y/
      LONGWORD 18.    000016B6
                                  /¶.../
      LONGWORD 19.    01030000
                                  /..../
      LONGWORD 20.    000A8001
                                  /..../
```

The 32-bit instance code always appears in "LONGWORD 1" of "CONTROLLER DEPENDENT INFORMATION," with the following exceptions:

- When MSLG$B_FORMAT reads "09 BAD BLOCK REPLACEMENT ATTEMPT," the instance code does not appear, because ERF does not provide "CONTROLLER DEPENDENT INFORMATION."
- When MSLG$B_FORMAT reads "0A MEDIA LOADER LOG," the instance code appears in "LONGWORD 2."
- When MSLG$B_FORMAT reads "00 CONTROLLER LOG," the instance code appears in part of both "LONGWORD 1" and "LONGWORD 2."

  For this "MSLG$B_FORMAT," the code is skewed and not directly readable as a longword. (The code's low-order bytes appear in the two high-order bytes of "LONGWORD 1," and the code's high-order bytes appear in the two low-order bytes of "LONGWORD 2.") For example:

```
CONTROLLER DEPENDENT INFORMATION
     LONGWORD 1.      030A0000
                                   /..../
     LONGWORD 2.      24010102
                                   /...$/
```

  In this case, the instance code is 0102030A.

  An OpenVMS DCL command procedure is provided at the end of this appendix (see Section C.7) for deskewing this particular instance code. Running the command procedure will make the error log directly readable when used in conjunction with the other information supplied in this appendix.

- Once you locate and identify the instance code, see the following sections for further information:
  - Section C.4 contains the Event Log Code tables, Tables C–2 through C–51. These tables list specific code descriptions.
  - Section C.3 contains detailed error packet descriptions, based on template type.
  - Section C.5 contains error threshold values.
  - Section C.6 contains recommended repair actions.
- When you look up a specific instance code, you will notice that each error belongs to one of fifteen template types. Each template type has a one byte value identifying it, which also is located in the "CONTROLLER DEPENDENT INFORMATION" area longwords, as shown in Table C–1.

  You may be able to use Table C–1 to quickly identify the template type, after examining the longwords in the "CONTROLLER DEPENDENT INFORMATION" area. However, because the location of the value identifying the template varies, the safest way to determine the template is to use the instance code. The template type is always the very next byte after the instance code.

Table C–1 Template Types

Description	Template	Longword	Value	Deskewed Value
Last Failure Event Log	01†	2	2401xxxx	00002401
Failover Event Log	05†	2	0005xxxx	00000005
Host buffer Access Error Event Log	10	2	00000C10	
Nonvolatile Parameter Memory Component Event Log	11	2	00000811	
Backup Battery Failure Event Log	12	2	00000012	
Subsystem Built-In Self Test Failure Event Log	13†	2	2413xxxx	00002413
Cache Memory Failure Event Log	14	2	00002414	
Internal Bus Port/Port Driver Event Log	32†	2	1032xxxx	00001032
Internal Bus System Communication Services Event Log	33†	2	2C33xxxx	00002C33
Device Services Nontransfer Event Error Log	41†	2	0441xxxx	00000441
Disk Transfer Error Event Log	51	2	00003C51	
Disk Bad Block Replacement (BBR) Attempt Event Log	57	No Longwords		
Tape Transfer Error Event Log	61	2	00003C61	
Media Loader Error Event Log	71	3	00003C71	

†The MSLG$B_FORMAT field for these templates will read "00 CONTROLLER LOG," so you may want to run the OpenVMS DCL command procedure provided at the end of this appendix (Section C.7) for deskewing the longwords.

- You should use the template type to learn even more from the error log. Information available in longwords, other than the instance code, includes the following:
  - Template type
  - Template information size
  - Event time
  - Drive sense data
  - Other information specific to the template

  Knowing the template type allows you to better use Section C.3 to obtain a complete description of each template and determine where information is located within the associated "CONTROLLER DEPENDENT INFORMATION."

C.2 Reading an HS1CP Error Log (DECevent)

DECevent-generated error reports, while containing basically the same information as ERF-generated reports, are far easier to interpret. This is true because more of the binary log is translated. Some of the information directly available from the ASCII report output includes the following:

- Instance code
- Recommended repair action
- Recovery threshold
- PTL of the faulty device
- ASC/ASCQ code values
- Template type
- MSCP event code
- Controller model

The following two examples illustrate the difference between event logs generated by ERF and DECevent. Example C–2 shows an ERF-generated disk transfer error log. Example C–2 shows the same error event translated by DECevent.

When using DECevent to translate event information, you should not need to refer as frequently to the information and tables included in the remaining sections of this appendix. However, familiarity with the ERF error log format, template types, and other elements of the event logs will help you understand the wide variety of events that might be reported, regardless of the translation utility you are using.

Example C–2 ERF-Generated Error Event Log

```
V M S                SYSTEM ERROR REPORT          COMPILED  9-AUG-1994 13:41:37
                                                                    PAGE 758.

******************************** ENTRY    1122. ********************************
ERROR SEQUENCE 5.                                 LOGGED ON:  CPU_TYPE 00000002
DATE/TIME 28-APR-1994 11:39:40.33                              SYS_TYPE 00000003
SYSTEM UPTIME: 0 DAYS 00:01:41
SCS NODE: MTX2                                                 OpenVMS AXP X6.1-FT7

HW_MODEL: 00000401 Hardware Model = 1025.

ERL$LOGMESSAGE ENTRY DEC 7000 Model 610

I/O SUB-SYSTEM, UNIT _MAT$DUA450:

      MESSAGE TYPE        0001
                                      DISK MSCP MESSAGE
      MSLG$L_CMD_REF  00000000
      MSLG$W_UNIT         01C2
                                      UNIT #450.
      MSLG$W_SEQ_NUM      0015
                                      SEQUENCE #21.
      MSLG$B_FORMAT         02
                                      DISK TRANSFER LOG
      MSLG$B_FLAGS          00
                                      UNRECOVERABLE ERROR
      MSLG$W_EVENT        014B
                                      DRIVE ERROR
                                      CNTRLR DETECTED PROTOCOL ERROR
      MSLG$Q_CNT_ID   01280009 40802576
                                      UNIQUE IDENTIFIER, 000940802576(X)
                                      MASS STORAGE CONTROLLER
                                      MODEL = 40.
      MSLG$B_CNT_SVR        14
                                      CONTROLLER SOFTWARE VERSION #20.
      MSLG$B_CNT_HVR        49
                                      CONTROLLER HARDWARE REVISION #73.
      MSLG$W_MULT_UNT     0035
      MSLG$Q_UNIT_ID  02FF0000 00000022
                                      UNIQUE IDENTIFIER, 000000000022(X)
                                      DISK CLASS DEVICE (166)
                                      HSXnn
      MSLG$B_UNIT_SVR       01
                                      UNIT SOFTWARE VERSION #1.
      MSLG$B_UNIT_HVR       43
                                      UNIT HARDWARE REVISION #67.
      MSLG$B_LEVEL          01
      MSLG$B_RETRY          00
      MSLG$L_VOL_SER  00000000
                                      VOLUME SERIAL #0.
      MSLG$L_HDR_CODE 00000000
                                      LOGICAL BLOCK #0.
                                      GOOD LOGICAL SECTOR

CONTROLLER DEPENDENT INFORMATION
```

(continued on next page)

Example C–2 (Cont.) ERF-Generated Error Event Log

```
LONGWORD 1.     030C4002
                                /.@../
LONGWORD 2.     00003C51
                                /Q<../
LONGWORD 3.     00000000
                                /..../
LONGWORD 4.     000B9331
                                /1.../
LONGWORD 5.     00000000
                                /..../
LONGWORD 6.     00000000
                                /..../
LONGWORD 7.     00000000
                                /..../
LONGWORD 8.     00000000
                                /..../
LONGWORD 9.     1F000504
                                /..../
LONGWORD 10.    36325A52
                                /RZ26/
LONGWORD 11.    20202020
                                /    /
LONGWORD 12.    29432820
                                / (C)/
LONGWORD 13.    43454420
                                / DEC/
LONGWORD 14.    34333533
                                /3534/
LONGWORD 15.    37313739
                                /9717/
LONGWORD 16.    00000000
                                /..../
LONGWORD 17.    00000004
                                /..../
LONGWORD 18.    00000000
                                /..../
LONGWORD 19.    853F0000
                                /..?./
LONGWORD 20.    00000000
```

Example C–3 DECevent-Generated Error Event Log

```
********************************* ENTRY 1122 *********************************

Logging OS                          3. OpenVMS AXP
OS version                             X6.1-FT7
Event sequence number               5.
Timestamp of occurrence                28-APR-1994 11:39:40
System uptime in seconds          101.
VMS error mask                 x00000000
VMS flags                          x0001  Dynamic Device Recognition present
Host name                              MTX2

AXP HW model                           DEC 7000 Model 610
System type register           x00000003  DEC 7000
Unique CPU ID                  x00000002
mpnum                          x000000FF
mperr                          x000000FF

Event validity                     -1. Unknown validity code
Event severity                     -1. Unknown severity code
Entry type                        100.
Major Event class                   3. IO Subsystem
AXP Device Type                     0.

IO Minor Class                      1. MSCP
IO Minor Sub Class                  5. Logged Message

---- Device Profile ----
Vendor
Product Name                           HSX00 MSCP basic disk
Unit Name                              MAT$DUA
Unit Number                       450.
Device Class                     x0001  Disk

---- IO SW Profile ----
VMS DC$_CLASS                       1.
VMS DT$_TYPE                      141.

---- MSCP Logged Msg ----

Command Reference number       x00000000
Unit Number                       450.
MSCP Sequence number               21.
Logged Message Format               2. Disk Transfer Error
MSCP Flags                       x00  No MSCP Flags indicated

MSCP Unique Controller-ID x0000000940802576
MSCP Controller Model              40. HS1CP HS Array Controller
MSCP Controller Class               1. Mass Storage Controller class
Controller SW version              20.
Controller HW version              73.

MSCP Unique Unit-ID       x0000000000000022
MSCP Unit Model                    -1. HSXOn MSCP basic/virtual disk
MSCP Unit Class                     2. Disk class - DEC Std 166 disk
Unit SW version                     1.
Unit HW version                    67.

MSCP Event Code                 x014B  Major Event =  Drive Error
                                       Sub-event =  Controller Detected Protocol
                                                    Error
```

(continued on next page)

Example C–3 (Cont.) DECevent-Generated Error Event Log

```
Multiunit code                  x0035
Error recovery Level               1.
Retry count                        0.
Volume Serial Number               0.
Header code                x00000000  Flags=  Good LBN
                                      LBN  =   0.

------- HSAC Data -------
Instance Code              x030C4002  A Drive failed because a Test Unit Ready
                                      command or a Read Capacity command failed.

                                      Component ID =   Device Services.
                                      Event Number =   x0000000C
                                      Repair Action =   x00000040
                                      NR Threshold =   x00000002
Template Type                   x51  Disk Transfer Error.
Power On Time Value       x00000000000B9331

Completed Byte Count               0.
Starting LBN                       0.
Device Locator             x00000504  Port    =   4.
                                      Target  =   5.
                                      LUN     =   0.
SCSI Device Type                 x1F  Device Type not decoded.
Drive Product Name                    RZ26     (C) DEC
Drive Serial Number                   35349717
Command Opcode                   x00  Test Unit Ready
Sense Data Qualifier             x00  Buf Mode =   The target shall not report
                                                   GOOD status on write commands
                                                   until the data blocks are
                                                   actually written on the
                                                   medium.
                                      UWEUO =    zero, not defined
                                      MSBD =    zero, not defined
                                      FBW =    zero, not defined
                                      DSSD =    Sense Data fields were generated
                                                by the HSAC controller on behalf
                                                of the target devices because the
                                                Sense Data could not be obtained
                                                from that device.
--- HSAC Sense Data ---

Error Code                       x00  Error Code no decoded
Segment #                        x00
Information Bytes          x00000000
CMD Specific Info          x00000000
Sense Key                        x04  Hardware Error

ASC & ASCQ                     x3F85  ASC  =   x003F
                                      ASCQ =   x0085
                                      Test Unit Ready or Read Capacity Command
                                      failed

FRU Code                         x00
Sense Key Specific Data    x00000000  Sense Key Data NOT Valid
                                      Byte 1 =   x00000000
                                      Byte 2 =   x00000000
                                      Byte 3 =   x00000000
```

C.3 Event Log Formats

Note

The numeric code values discussed in the figures and tables of this appendix are hexadecimal, unless otherwise stated.

The HS1CP controller reports significant events that occur during normal controller operation using the following standard MSCP and TMSCP error log message formats:

- Controller errors
- Memory errors
- Disk transfer errors
- Bad block replacement attempts
- Tape errors
- Media loader errors
- Disk copy data correlation

To more fully use the remainder of this appendix, you should become familiar with MSCP and TMSCP protocols, especially in the area of error log message formats.

Implementation Dependent Information Area

With the exception of the disk copy data correlation error log message format, each of the error log message formats listed in Section C.3 provides an "implementation dependent information" area located at the end of the message. For HS1CP controller specific event logs that area is formatted as shown in Figure C–1.

Note that the fields shown in Figure C–1 always begin on a longword boundary within HS1CP controller specific event logs. If the "implementation dependent information" area of a particular MSCP error log message format does not begin on a longword boundary, a "reserved" field containing the appropriate number of bytes is appended to the format to provide the necessary alignment (such as see offset 16 in Figure C–15).

Implementation Dependent Information Fields:

instance code

A number that uniquely identifies the event being reported. The format of this field is shown in Figure C–2.

Figure C–1 Implementation Dependent Information Format

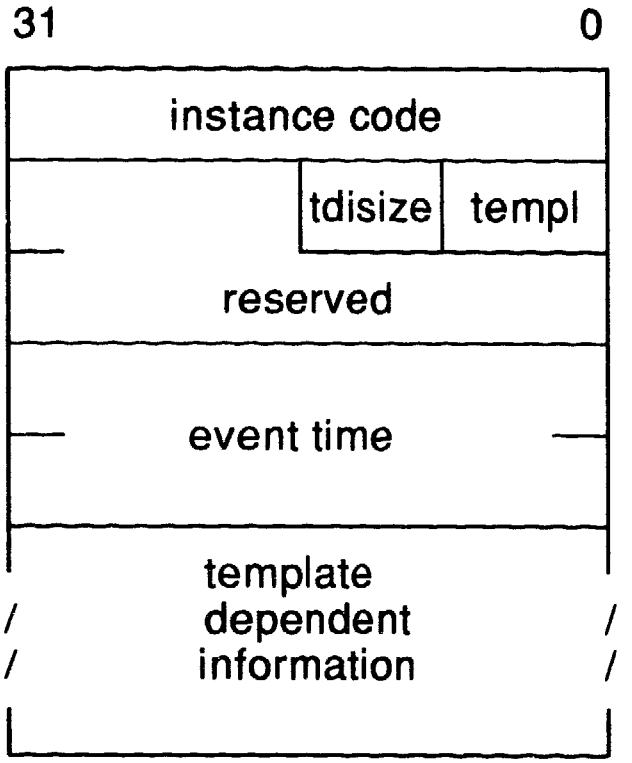
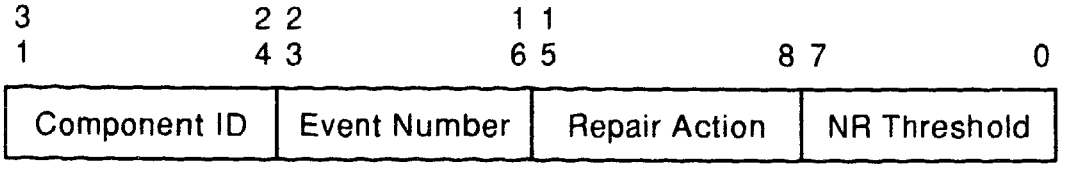


Figure C–2 Instance Code Format



Instance Code Specific Subfields:

NR Threshold

The notification/recovery threshold assigned to the event. This value is used during Symptom Directed Diagnosis procedures to determine when notification/recovery action should be taken. See Section C.5 for more detail.

Repair Action

The recommended repair action code assigned to the event. This value is used during Symptom Directed Diagnosis procedures to determine what notification/recovery action should be taken when the NR Threshold is reached. See Section C.6 for more detail.

Event Number

A number, when combined with the value contained in the Component ID subfield, uniquely identifies the event.

Component ID

A number that uniquely identifies the firmware component that detected the event as shown in Table C–2.

templ

A number that uniquely describes the format of the "template dependent information" field.

tdisize

The number of bytes contained in the "template dependent information" field.

reserved

Reserved for future use.

event time

The time the event occurred according to the power on time value maintained by the HS1CP controller operational firmware.

The power on time value is a 64-bit unsigned integer that represents the total number of seconds HS1CP controller operational firmware has executed on the HS1CP controller board. Note that the time expended during controller restarts, power-on diagnostics, and system initialization is not accounted for by this value.

template dependent information

A variable length field containing information specific to the event being reported.

This field is divided into separate fields specific to the template identified in the "templ" field. The template specific fields common to multiple event logs are described in separate subsections of Section C.3.1 to avoid duplication of the field descriptions in Section C.3.2.

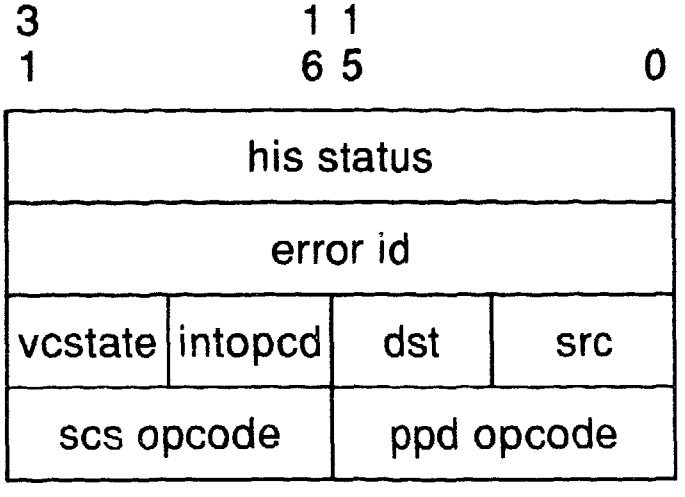
C.3.1 Common Event Log Fields

Common fields are generated across certain event logs. These common fields are described in Sections C.3.1.1 through C.3.1.5.

C.3.1.1 CI Host Interconnect Services Common Event Log Fields

The fields common to certain event logs generated by the CI Host Interconnect Services firmware component are shown in Figure C–3.

Figure C–3 CI Host Interconnect Services Common Event Log Fields



CI Host Interconnect Services Common Fields:

his status

The Host Interconnect Services status code as shown in Table C–3.

error id

The address of the Host Interconnect Services routine that detected the event.

src

The CI source node address.

dst

The CI destination node address.

intopcd

The CI message opcode as shown in Table C–4.

vcstate

The virtual circuit state code as shown in Table C–5.

NOTE

The setting of the high order bit (Bit 7) in this field indicates the state of ID polling for the virtual circuit.

If Bit 7 is set, ID polling is complete. Otherwise, ID polling is incomplete.

ppd opcode

The Port/Port Driver layer opcode as shown in Table C–6.

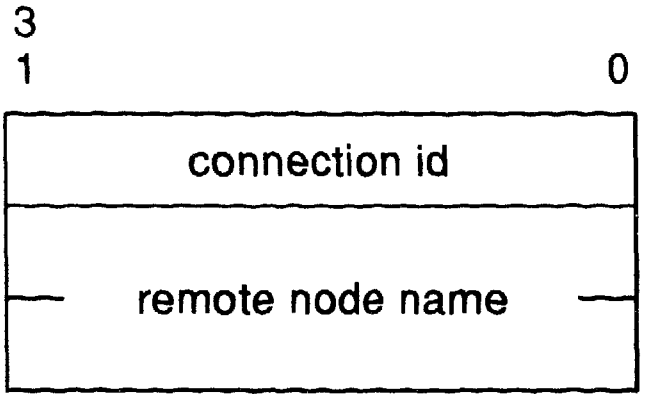
scs opcode

The System Communication Services layer opcode as shown in Table C–7.

C.3.1.2 Host/Server Connection Common Fields

The fields common to certain event logs generated by the Disk and Tape MSCP Server, CI Host Interconnect Services, Device Services, and Value Added firmware components are shown in Figure C–4.

Figure C–4 Host/Server Connection Common Fields



Host/Server Connection Common Fields:

connection id

Identifies the host/server connection associated with the event being reported. If this value is zero, the host/server connection information was invalidated before the event could be reported.

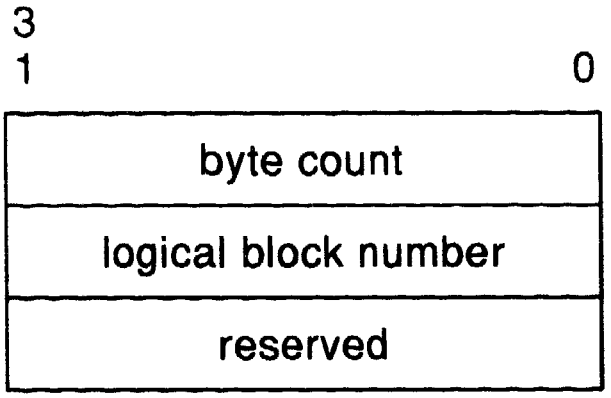
remote node name

An 8-byte ASCII string that represents the node name associated with the host/server connection identified in the "connection id" field. If the "connection id" field is zero, the content of this field is undefined.

C.3.1.3 Byte Count/Logical Block Number Common Fields

The fields common to certain event logs generated by the Device Services and Value Added firmware components are shown in Figure C–5.

Figure C–5 Byte Count/Logical Block Number Common Fields



Byte Count/Logical Block Number Common Fields:

byte count

Number of bytes of the HS1CP controller firmware component initiated transfer successfully transferred.

logical block number

Starting logical block number of the HS1CP controller firmware component initiated transfer.

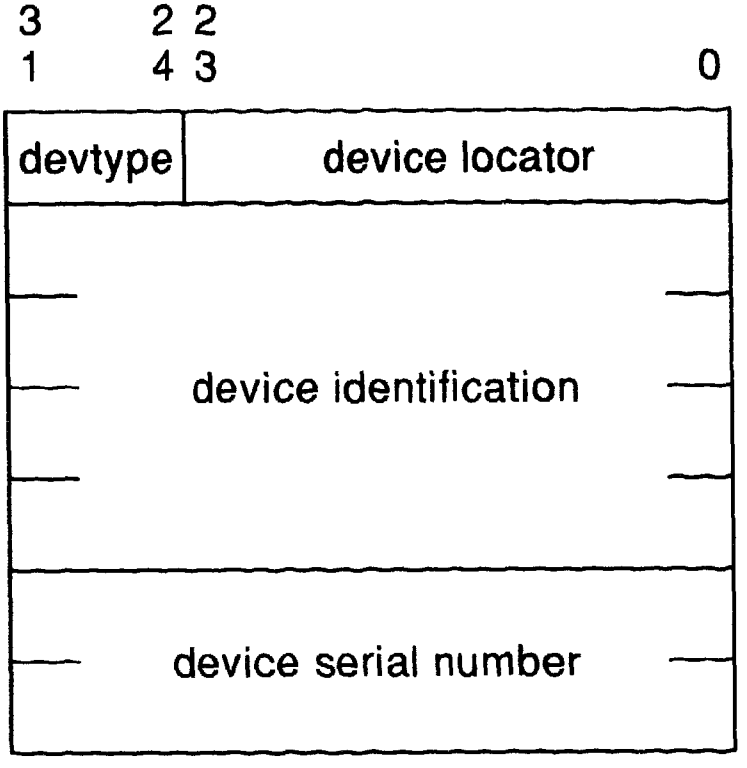
reserved

Reserved for future use, currently contains the value 0.

C.3.1.4 Device Location/Identification Common Fields

The fields common to certain event logs generated by the Device Services and Value Added firmware components are shown in Figure C–6.

Figure C–6 Device Location/Identification Common Fields

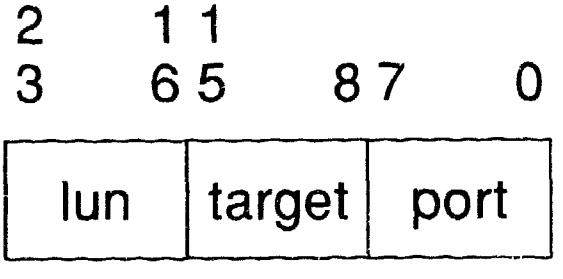


Device Location/Identification Common Fields:

device locator

The location within the HS1CP controller's subsystem of the target device involved in the event being reported. This field is formatted as shown in Figure C–7.

Figure C–7 Device Locator Field Format



Device Locator Specific Subfields:

port

The SCSI bus number to which the target device is connected.

target

The SCSI ID of the device connected to the port.

lun

The logical unit number on the "target" by which the target device is logically addressed.

devtype

The SCSI device type of the device. The various SCSI device types supported by the HS1CP controller are shown in Table C–9.

device identification

Sixteen bytes of ASCII data as defined by the device vendor in the Product Identification field of the SCSI INQUIRY command data.

The most significant character of the product identification data will appear in the low order byte of the first longword of this field while the least significant character appears in the high order byte of the last long word.

device serial number

Eight bytes of ASCII data as defined by the device vendor in the Product Serial Number field of the SCSI Unit Serial Number Page data.

The most significant character of the serial number data will appear in the low order byte of the first longword of this field while the least significant character appears in the high order byte of the last longword.

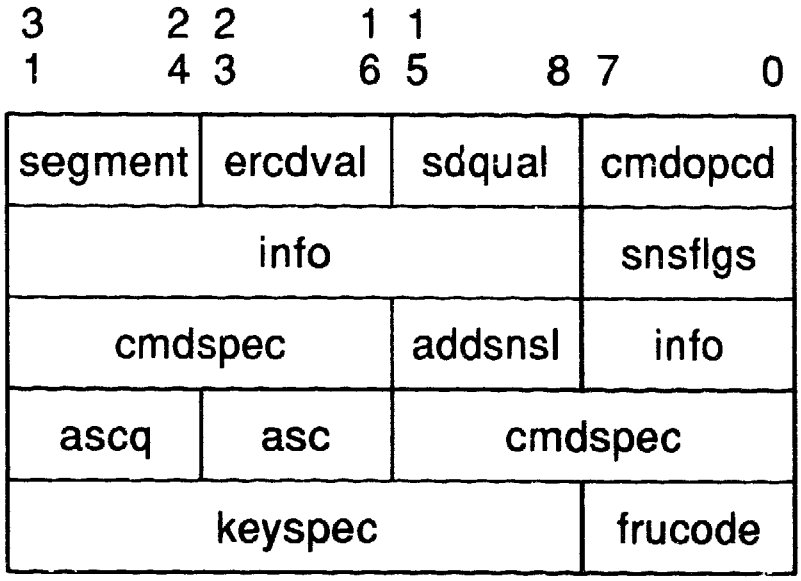
Note that the number of characters of serial number data supplied may vary from vendor to vendor as well as from device to device. If the serial number data supplied is less than eight characters, this field is ASCII space filled from the lowest order byte (relative to the low order byte of the first longword) containing a serial number character through the high order byte of the last longword. If the serial number data supplied is greater than eight characters, the serial number data is truncated at eight bytes (that is, the least significant characters of the serial number data is lost). If the serial number data is not available at all, this field is ASCII space filled.

C.3.1.5 SCSI Device Sense Data Common Fields

The fields common to certain event logs generated by the Device Services and Value Added firmware components are shown in Figure C–8.

The first two fields shown in Figure C–8, the "cmdopcd" and "sdqual" fields, are supplied by the HS1CP controller to provide qualifying information required to interpret the other SCSI Sense Data Common fields. The other fields, "ercdval" through "keyspec," contain standard Sense Data, returned in the response of a SCSI REQUEST SENSE command issued to the target device or generated by the HS1CP controller on the target device's behalf.

Figure C–8 SCSI Device Sense Data Common Fields



SCSI Device Sense Data Common Fields:

cmdopcd

The operation code of the SCSI command issued to the target device. SCSI command operation codes vary according to device type (see Table C–10) so the content of this field depends on the content of the "devtype" field.

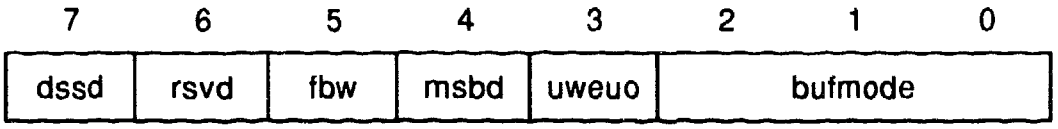
See the description of the "ercdval" field for information regarding the validity of this field.

sdqual

This field contains information necessary to determine whether the Sense Data contained in the "ercdval" through "keyspec" fields is supplied by an attached device or generated by the HS1CP controller

itself and to qualify the content of the "info" field. This field is formatted as shown in Figure C–9.

Figure C–9 Sense Data Qualifier Field Format



Sense Data Qualifier Specific Subfields:

bufmode

The SCSI buffered mode selected on the device. The various SCSI Buffered Modes are shown in Table C–11.

uweuo

This bit is set to one if and only if an unrecoverable write error was detected while unwritten objects (that is, data blocks, filemarks, or setmarks) remain in the buffer.

msbd

This bit is set to one if and only if the MODE SENSE block descriptor is nonzero.

fbw

This bit is set to one if and only if the Fixed bit of the WRITE command is set to one.

rsvd

Reserved for future use.

dssd

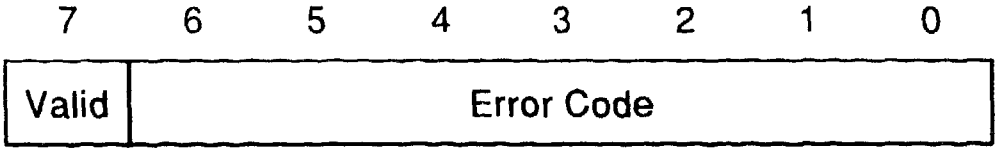
This bit is set to one if and only if the Sense Data contained in the "ercdval" through "keyspec" fields is supplied by the target device.

If this bit is zero, the Sense Data contained in the "ercdval" through "keyspec" fields is generated by the HS1CP controller on behalf of the target device because the Sense Data could not be obtained from that device.

ercdval

This field contains byte 0 of the Sense Data returned in the response of a SCSI REQUEST SENSE command. This field is formatted as shown in Figure C–10.

Figure C–10 SCSI Sense Data Byte Zero ("ercdval") Field Format



SCSI Sense Data Byte Zero ("ercdval") Specific Subfields:

Error Code

An error code of 70 indicates that the event being reported occurred during the execution of the current command, identified in the "cmdopcd" field.

An error code of 71 indicates that the event being reported occurred during execution of a previous command for which GOOD status has already been returned. The "cmdopcd" field is undefined in this case.

For error codes 70 and 71 the remaining fields of the event log (such as segment, snsflgs, info, and so forth) will contain the standard SCSI Sense Data fields (bytes 1 through 17) returned in the response of a SCSI REQUEST SENSE command.

An error code of 7F indicates that the Sense Data fields are in a vendor-specific format so the content of the remaining event log fields can only be determined from documentation provided by the vendor of the target device.

The SCSI specification states that error code values 72 through 7E are currently reserved for future use and that error codes 00 through 6F are not defined. Should this field contain any of those codes the remaining event log fields are undefined.

Valid

If this bit is set to one, the content of the Sense Data Information field (bytes 3 through 6) is valid and its content is as defined by the SCSI specification (see the description of the "info" field for the SCSI definition of the Sense Data Information field). Otherwise, the Sense Data Information field is not as defined by the SCSI specification (refer to documentation provided by the device vendor for their definition of the field).

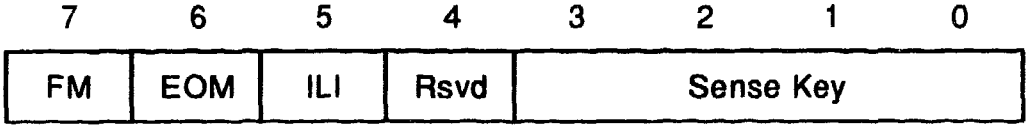
segment

This field contains byte 1 (Segment field) of the Sense Data returned in the response of a SCSI REQUEST SENSE command. If the "cmdopcd" is an 18 (COPY), 39 (COMPARE), or 3A (COPY AND VERIFY), this field contains the number of the current segment descriptor.

snsflgs

This field contains byte 2 of the Sense Data returned in the response of a SCSI REQUEST SENSE command. This field is formatted as shown in Figure C–11.

Figure C–11 SCSI Sense Data Byte Two ("snsflgs") Field Format



SCSI Sense Data Byte Two ("snsflgs") Specific Subfields:

Sense Key

The sense key provides generic categories in which events can be reported. The sense keys are described in Table C–12.

ILI

An incorrect length indicator (ILI) bit of one usually indicates that the requested logical block length did not match the logical block length of the data on the medium.

EOM

For sequential-access devices (that is, "devtype" is 1) an end-of-medium (EOM) bit set to one indicates that the unit is at or past the early-warning if the direction was forward or that the command could not be completed because beginning-of-partition was encountered if the direction was reverse.

FM

A filemark (FM) bit set to one indicates that the current command has read a filemark or setmark. The Additional Sense Code field (see "asc" field description) may be used to indicate whether a filemark or setmark was read. Note that the reporting of setmarks is optional.

info

This field contains bytes 3 through 6 (Information field) of the Sense Data returned in the response of a SCSI REQUEST SENSE command. The content of this field varies depending on the values contained in

the "devtype" and "cmdopcd" fields and the "bufmode," "uweuo," "msbd," and "fbw" subfields of the "sdqual" field as follows:

1. Regardless of the value of the "devtype" field and the "sdqual" subfields, if the "cmdopcd" is an 18 (COPY), 39 (COMPARE), or 3A (COPY AND VERIFY), this field contains the difference (residue) of the requested number of blocks minus the actual number of blocks copied or compared for the current segment descriptor.
2. Regardless of the value of the "sdqual" subfields, if "devtype" is 0 (Direct-Access Devices—such as magnetic disk) or 5 (CD-ROM Devices) and "cmdopcd" is not an 18 (COPY), 39 (COMPARE), or 3A (COPY AND VERIFY), this field contains the unsigned logical block address associated with the value contained in the Sense Key subfield of the "snsflgs" field (see Figure C–11).
3. Regardless of the value of "cmdopcd," if "devtype" is 1 (Sequential-Access Devices—such as magnetic tape) and "uweuo" is 1 and "bufmode" is either 1 or 2, this field contains:
   a. the total number of objects in the buffer if "msbd" and "fbw" are both 1.
   b. the number of bytes in the buffer, including filemarks and setmarks, if "msbd" is 1 and "fbw" is 0.

addsnsl

This field contains byte 7 (Additional Sense Length field) of the Sense Data returned in the response of a SCSI REQUEST SENSE command. This field contains the number of additional Sense Data bytes to follow.

If this value is less than 10, the content of some or all of the remaining event log fields (that is, cmdspec, asc, ascq, frucode, and keyspec) may be undefined. The "cmdspec" field is undefined unless this value is 4 or greater. The "asc" and "ascq" fields are undefined unless this value is 6 or greater. The "frucode" field is undefined unless this value is 7 or greater. The "keyspec" field is undefined unless this value is 10 or greater.

If this value is greater than 10, the device supplied the Additional Sense Bytes field, which begins at byte 12 of the Sense Data. The content of the Additional Sense Bytes field is not included in the event log.

cmdspec

If the value contained in the "addsnsl" field is 4 or greater, this field contains bytes 8 through 0B (Command-Specific Information field) of the Sense Data returned in the response of a SCSI REQUEST SENSE command. The content of this field varies depending on the value contained in the "cmdopcd" field as follows:

1. If the "cmdopcd" is an 18 (COPY), 39 (COMPARE), or 3A (COPY AND VERIFY), the low order byte of this field contains the starting byte number of an area relative to Sense Data byte 0 that contains (unchanged) the source logical unit's status byte and sense data and the next higher order byte contains the starting

byte number of an area relative to Sense Byte 0 that contains (unchanged) the destination logical unit's status byte and sense data. If the low order or next higher order byte of this field contains the value zero, no status byte or sense data was supplied for the corresponding (source or destination) logical unit. The content of the highest order two bytes of this field is undefined.

2. If the "cmdopcd" is a 7 (REASSIGN BLOCKS), this field contains the logical block address of the first defect descriptor not reassigned. If information about the first defect descriptor not reassigned is not available, or if all the defects have been reassigned, this field will contain the value FFFFFFFF.
3. If the "cmdopcd" is a 31 (SEARCH DATA EQUAL), 30 (SEARCH DATA HIGH), or 32 (SEARCH DATA LOW) and the Sense Key subfield of the "snsflgs" field (see Figure C–11) value is EQUAL, this field contains the record offset of the matching record.

asc
ascq

If the value contained in the "addsnsl" field is 6 or greater and the "dssd" subfield of the "sdqual" field is equal to 1, the "asc" and "ascq" fields contain the values supplied in the byte 0C (Additional Sense Code) and byte 0D (Additional Sense Code Qualifier) fields, respectively, of the Sense Data returned in the response of a SCSI REQUEST SENSE command issued to the target device. The Additional Sense Code (ASC) field and the Additional Sense Code Qualifier (ASCQ) field together describe the event being reported. The standard SCSI ASC/ASCQ codes are "devtype" dependent as shown in Table C–13 through Table C–16. Note that the SCSI specification defines ASCs within the range 80 through FF in combination with ASCQs within the range 00 through FF and ASCQs within the range 80 through FF regardless of ASC value as being vendor specific. Refer to documentation provided by the vendor of the target device for a description of an ASC/ASCQ value that falls within the defined vendor specific ranges.

If the value contained in the "addsnsl" field is 6 or greater and the "dssd" subfield of the "sdqual" field is equal to 0, the "asc" and "ascq" fields contain HS1CP controller vendor specific SCSI ASC/ASCQ codes generated by the HS1CP on behalf of the target device. See Table C–17 for the descriptions of the HS1CP controller vendor specific SCSI ASC/ASCQ codes.

frucode

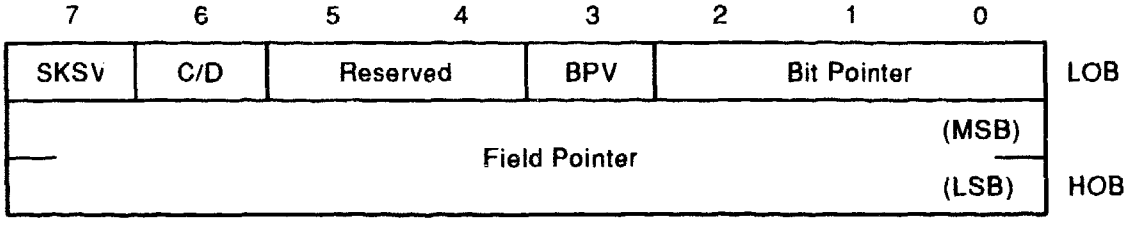
If the value contained in the "addsnsl" field is 7 or greater, this field contains byte 0E (Field Replaceable Unit field) of the Sense Data returned in the response of a SCSI REQUEST SENSE command. If this field is nonzero, the target device is identifying the "field replaceable unit" that has failed. See the service manual for the target device for complete details of the meaning of this value.

keyspec

If the value contained in the "addsnsl" field is 10 or greater, this field contains bytes 0F through 11 (Sense-Key Specific field) of the Sense Data returned in the response of a SCSI REQUEST SENSE command. The definition of this field is determined by the value of the Sense Key subfield of the "snsflgs" field. This field is reserved for Sense Key values other than ILLEGAL REQUEST, RECOVERED ERROR, HARDWARE ERROR, MEDIUM ERROR and NOT READY.

If the Sense Key value is ILLEGAL REQUEST, the format of this field is as shown in Figure C–12.

Figure C–12 SCSI Sense Data Byte 0F through 11 ("keyspec") Field—Field Pointer Bytes Format



SCSI Sense Data Byte 0F through 11 ("keyspec")—Field Pointer Bytes Specific Subfields:

Bit Pointer and BPV

A bit pointer valid (BPV) bit of zero indicates that the value in the Bit Pointer subfield is not valid. A BPV bit of one indicates that the Bit Pointer subfield specifies which bit of the byte designated by the Field Pointer field is in error. When a multiple-bit field is in error, the Bit Pointer subfield points to the most-significant (left-most) bit of the field.

C/D

A command data (C/D) bit of one indicates that the illegal parameter is in the command descriptor block. A C/D bit of zero indicates that the illegal parameter is in the data parameters sent by the initiator during the DATA OUT phase.

SKSV

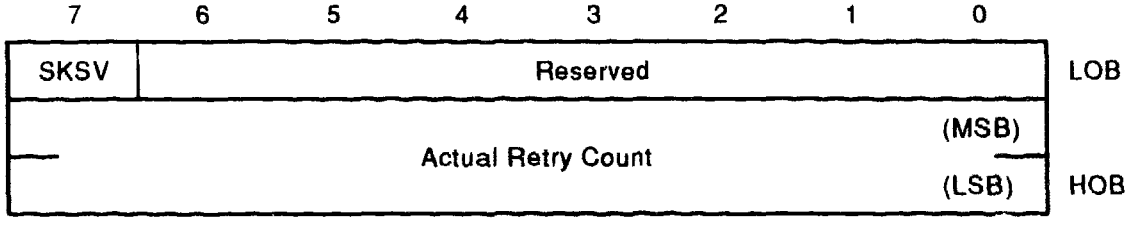
The content of the "keyspec" field is valid if and only if this bit is set to one.

Field Pointer

The Field Pointer subfield indicates which byte of the command descriptor block or of the parameter data was in error. When a multiple-byte field is in error, the pointer points to the most-significant (left-most) byte of the field.

If the Sense Key value is RECOVERED ERROR or HARDWARE ERROR or MEDIUM ERROR, the format of this field is as shown in Figure C–13.

Figure C–13 SCSI Sense Data Byte 0F through 11 ("keyspec") Field—Actual Retry Count Bytes Format



SCSI Sense Data Byte 0F through 11 ("keyspec")—Actual Retry Count Bytes Specific Subfields:

SKSV

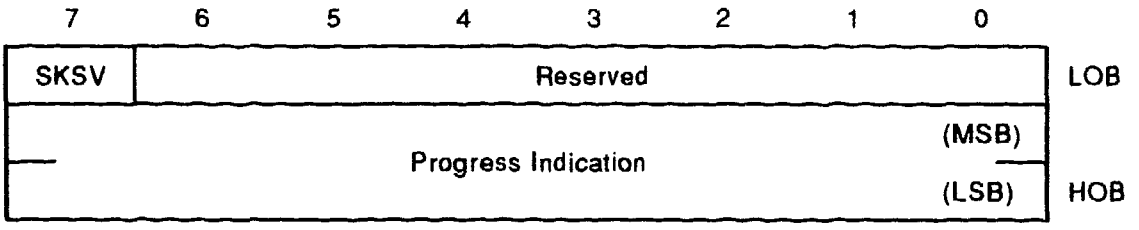
The content of the "keyspec" field is valid if and only if this bit is set to one.

Actual Retry Count

The actual retry count subfield contains the implementation-specific information on the actual number of retries of the recovery algorithm used in attempting to recover an error or exception condition.

If the Sense Key value is NOT READY and the last command issued to the device was a FORMAT UNIT, the format of this field is as shown in Figure C–14.

Figure C–14 SCSI Sense Data Byte 0F through 11 ("keyspec") Field—Progress Indication Bytes Format



SCSI Sense Data Byte 0F through 11 ("keyspec")—Progress Indication Bytes Specific Subfields:

SKSV

The content of the "keyspec" field is valid if and only if this bit is set to one.

Progress Indication

This subfield is a percent complete indication in which the returned value is the numerator that has 10000 as its denominator. The progress indication is based upon the total format operation including any certification or initialization operations.

C.3.2 Specific Event Log Formats

In addition to the common fields generated across certain event logs, there is specific information for each log, based on template type. The specific information is described in Sections C.3.2.1 through C.3.2.14.

C.3.2.1 Last Failure Event Log (Template 01)

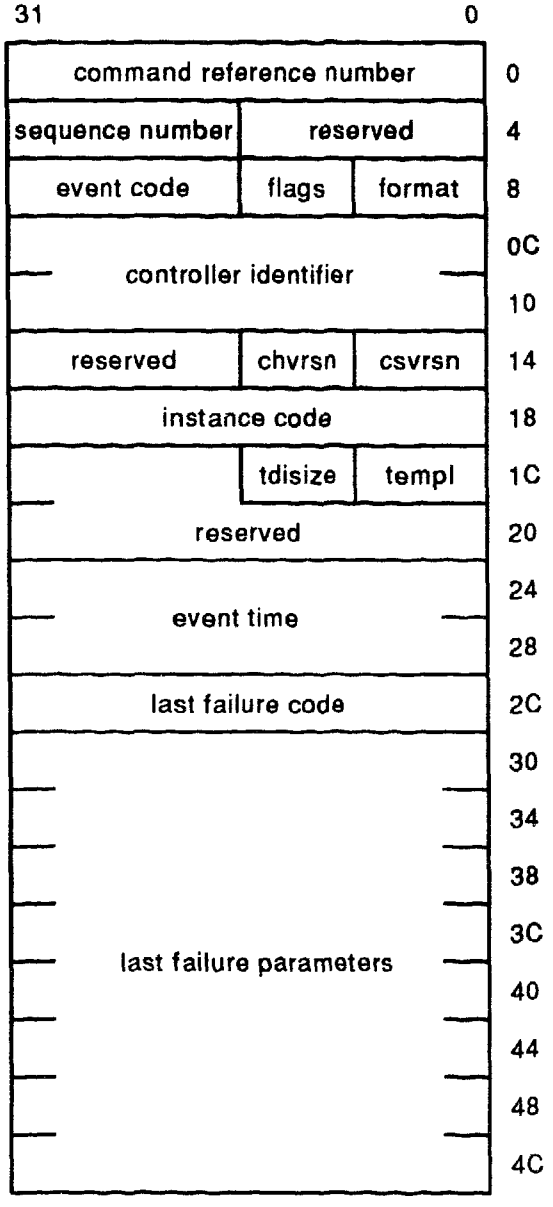
Unrecoverable conditions detected by either firmware or hardware and certain operator initiated conditions result in the termination of HS1CP controller operation.

In most cases, following such a termination the controller will attempt to restart (that is, restart) with hardware components and firmware data structures initialized to the states necessary to perform normal operations.

If the restart is successful and communications are re-established with the host systems and "Miscellaneous" error logging is enabled by one or more host systems, the HS1CP controller will send a Last Failure Event Log, which describes the condition that caused controller operation to terminate, to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Last Failure Event Log is reported via the T/MSCP Controller Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–15.

Figure C–15 Last Failure Event Log (Template 01) Format



Last Failure Event Log Format Specific Fields:

format

This field contains the value 00 (that is, T/MSCP Controller Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–18.

reserved (offset 16)

This field contains the value 0.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–18.

templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 01 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 24 for this event log.

reserved (offset 1E)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

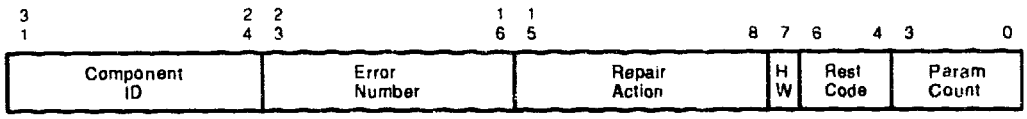
last failure code

A number that uniquely describes the unrecoverable condition being reported as shown in Table C–33 through Table C–49. The format of this field is shown in Figure C–16.

Note

Do not confuse this field with the "instance code" field. They are similar in format but convey different information.

Figure C–16 Last Failure Code Format



Last Failure Code Specific Subfields:

Parameter Count

The number of longwords of supplemental information provided in the "last failure parameters" field.

Restart Code

A number that describes the actions taken to restart the controller after the unrecoverable condition was detected as shown in Table C–51.

HW

Hardware/firmware flag. If this flag is equal to 1, the unrecoverable condition is due to a hardware detected fault. If this flag is equal to 0, the unrecoverable condition is due to a firmware detected inconsistency.

Repair Action

The recommended repair action code assigned to the condition. This value is used during Symptom Directed Diagnosis procedures to determine what notification/recovery action should be taken. See Section C.6 for more detail.

Error Number

A number, when combined with the value contained in the Component ID subfield, uniquely identifies the condition detected.

Component ID

A number that uniquely identifies the firmware component that reported the condition as shown in Table C–2.

last failure parameters

This field contains supplemental information specific to the failure being reported.

The content of the parameters supplied (if any) are described in the individual "last failure code" descriptions contained in Table C–33 through Table C–49.

C.3.2.2 Failover Event Log (Template 05)

The HS1CP controller Failover Control firmware component reports errors and other conditions encountered during redundant controller communications and failover operation via the Failover Event Log.

The Failover Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Failover Event Log is reported via the T/MSCP Controller Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–17.

Failover Event Log Format Specific Fields:

format

This field contains the value 00 (that is, T/MSCP Controller Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–19.

reserved (offset 16)

This field contains the value 0.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–19.

templ

See Implementation Dependent Information Area for the description of this field.

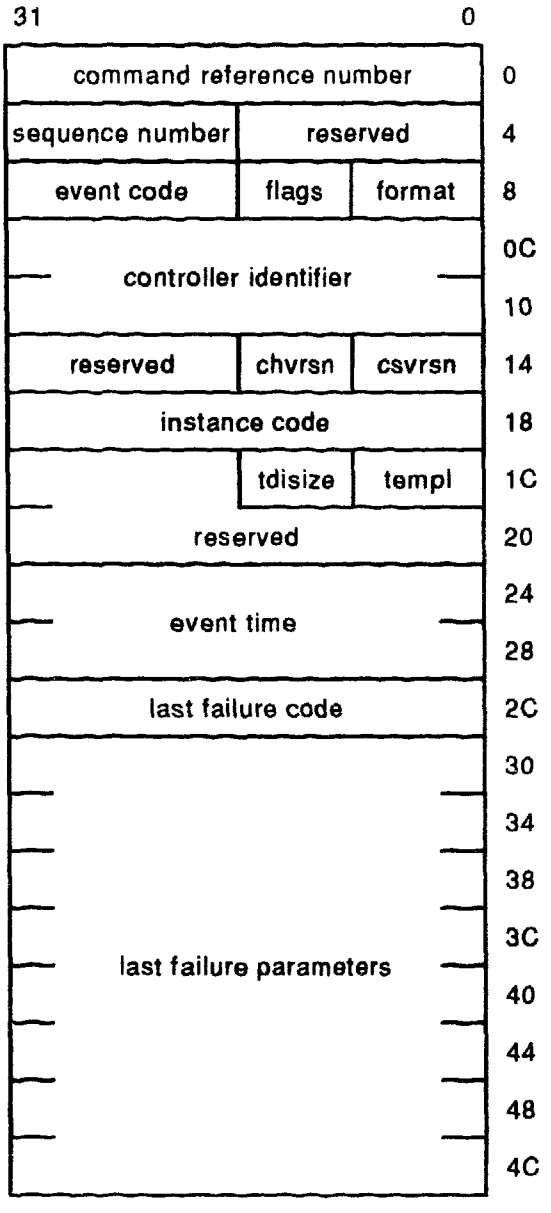
This field contains the value 05 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 24 for this event log.

Figure C–17 Failover Event Log (Template 05) Format



reserved (offset 1E)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

last failure code
last failure parameters

These fields contain the last failure information supplied in the last gasp message sent by the other HS1CP controller in a dual-redundant configuration as a normal part of terminating controller operation.

See Section C.3.2.1 for the description of the format of these fields.

Note that the content of certain of the fields described above may be undefined depending on the value supplied in the "instance code" field. See Table C–19 for more detail.

C.3.2.3 Firmware License Violation Event Log (Template 06)

The HS1CP controller Value Added firmware component reports firmware license violations via the Firmware License Violation Event Log.

The Firmware License Violation Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Firmware License Violation Event Log is reported via the T/MSCP Controller Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–18.

Firmware License Violation Event Log Format Specific Fields:

format

This field contains the value 00 (that is, T/MSCP Controller Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–20.

reserved (offset 16)

This field contains the value 0.

instance code

See Implementation Dependent Information Area for the description of this field.

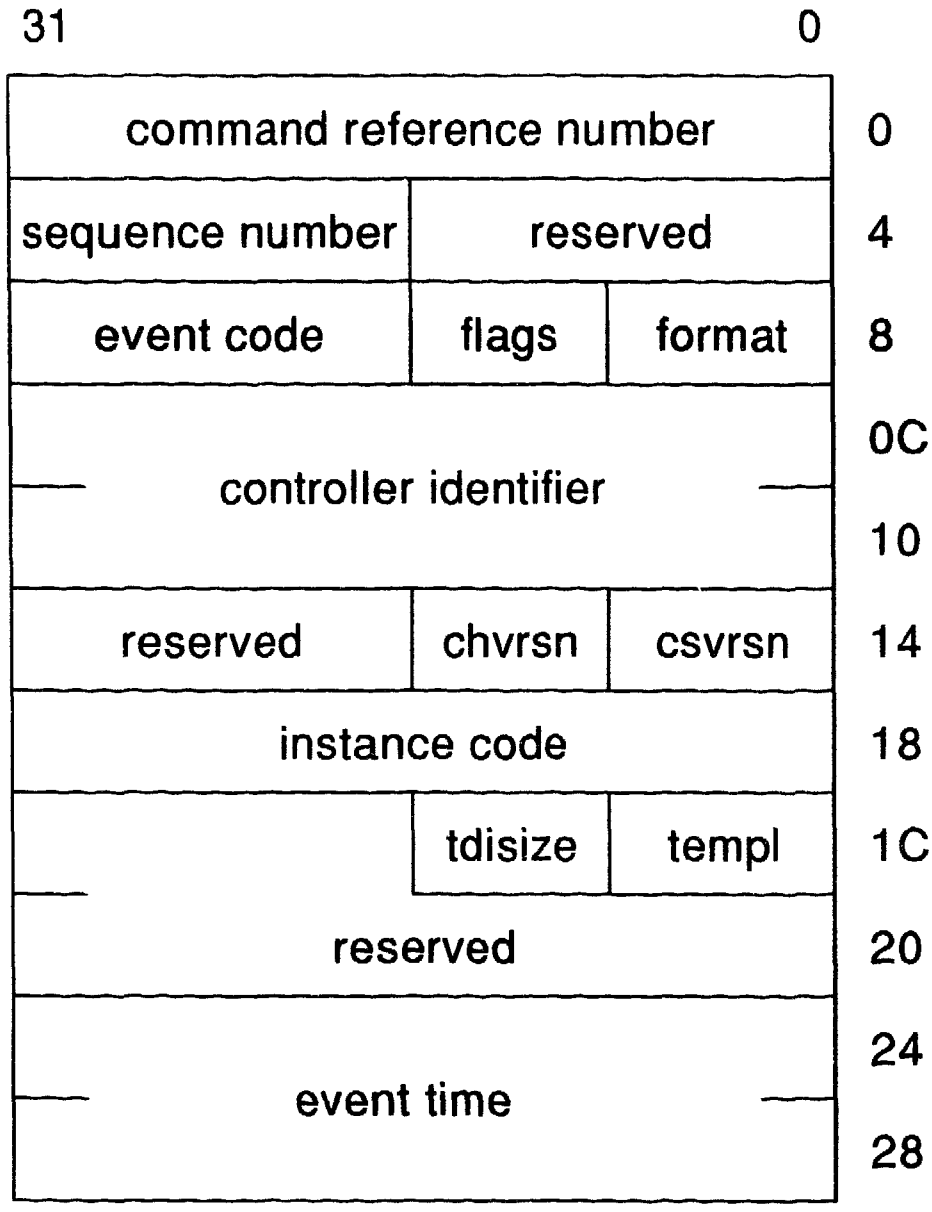
The values that can be reported in this field for this event log are shown in Table C–20.

templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 06 for this event log.

Figure C–18 Firmware License Violation Event Log (Template 06) Format



tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 00 for this event log.

reserved (offset 1E)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

C.3.2.4 Nonvolatile Parameter Memory Component Event Log (Template 11)

The HS1CP controller Executive firmware component reports errors detected while accessing a Nonvolatile Parameter Memory Component via the Nonvolatile Parameter Memory Component Event Log.

The Nonvolatile Parameter Memory Component Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Nonvolatile Parameter Memory Component Event Log is reported via the T/MSCP Memory Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–19.

Nonvolatile Parameter Memory Component Event Log Format Specific Fields:

format

This field contains the value 01 (that is, T/MSCP Memory Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–21.

memory address

The physical address of the beginning of the affected Nonvolatile Parameter Memory component area.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–21.

templ

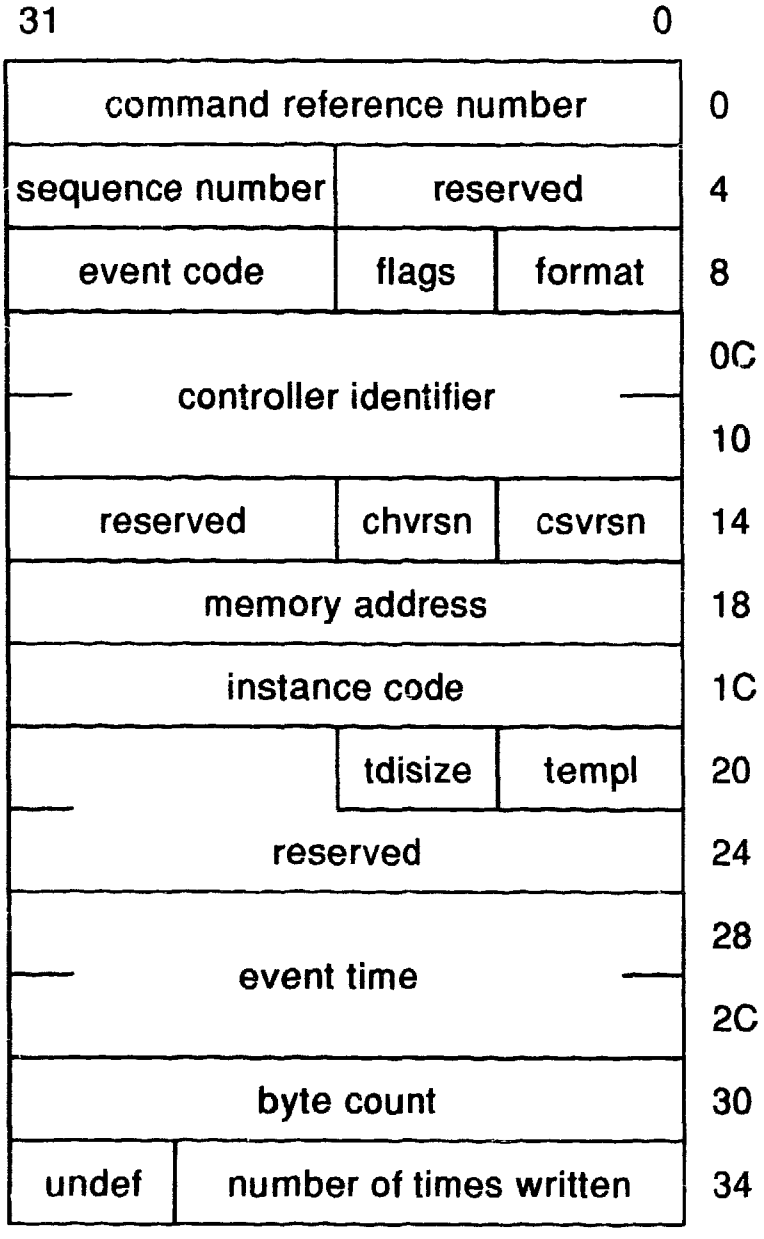
See Implementation Dependent Information Area for the description of this field.

This field contains the value 11 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

Figure C–19 Nonvolatile Parameter Memory Component Event Log (Template 11) Format



This field contains the value 08 for this event log.

reserved (offset 22)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

byte count

The number of bytes contained in the affected Nonvolatile Parameter Memory component area (that is, the area bounded by: "memory address" through "memory address" + "byte count" - 1).

number of times written

The number of times the affected Nonvolatile Parameter Memory component area has been written.

undef

This field is only present to provide longword alignment, its content is undefined.

C.3.2.5 Backup Battery Failure Event Log (Template 12)

The HS1CP controller Value Added Services firmware component reports backup battery failure conditions for the various hardware components that use a battery to maintain state during power-failures via the Backup Battery Failure Event Log.

The Backup Battery Failure Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Backup Battery Failure Event Log is reported via the T/MSCP Memory Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–20.

Backup Battery Failure Event Log Format Specific Fields:

format

This field contains the value 01 (that is, T/MSCP Memory Errors error log format code).

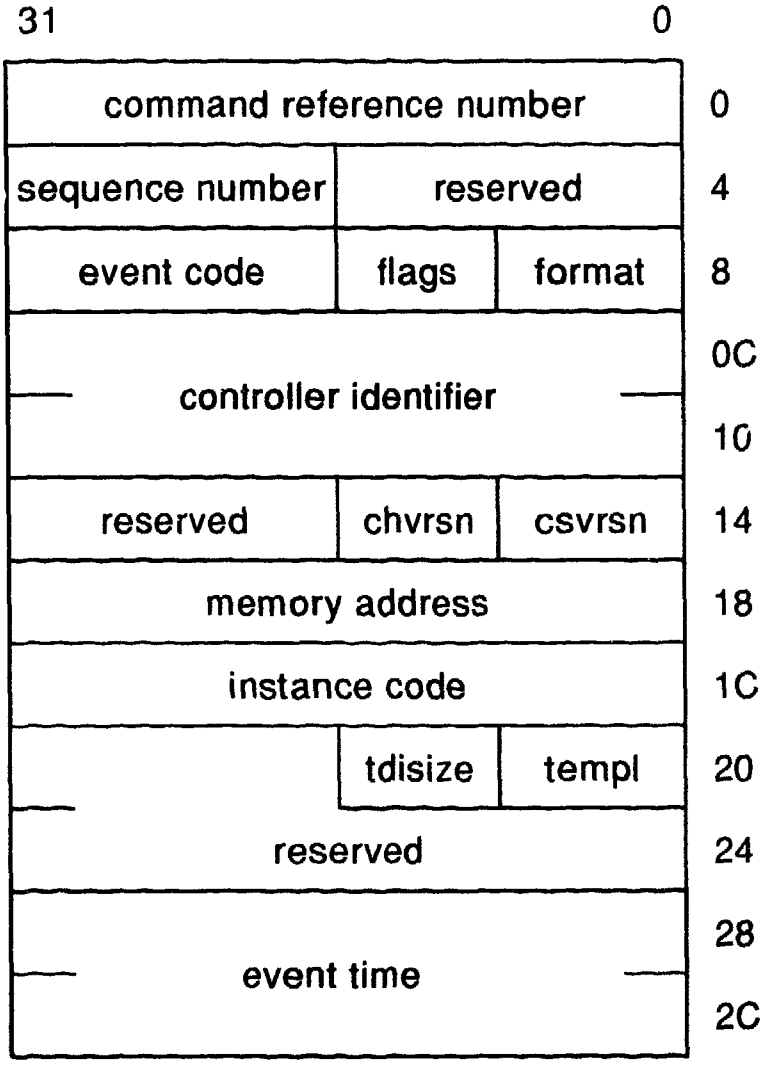
event code

The values that can be reported in this field for this event log are shown in Table C–22.

memory address

The content of this field depends on the value supplied in the "instance code" field. See Table C–22 for more detail.

Figure C–20 Backup Battery Failure Event Log (Template 12) Format



instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–22.

templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 12 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 00 for this event log.

reserved (offset 22)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

C.3.2.6 Subsystem Built-In Self Test Failure Event Log (Template 13)

The HS1CP controller Subsystem Built-In Self Tests firmware component reports errors detected during test execution via the Subsystem Built-In Self Test Failure Event Log.

The Subsystem Built-In Self Test Failure Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Subsystem Built-In Self Test Failure Event Log is reported via the T/MSCP Controller Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–21.

Subsystem Built-In Self Test Failure Event Log Format Specific Fields:

format

This field contains the value 00 (that is, T/MSCP Controller Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–23.

reserved (offset 16)

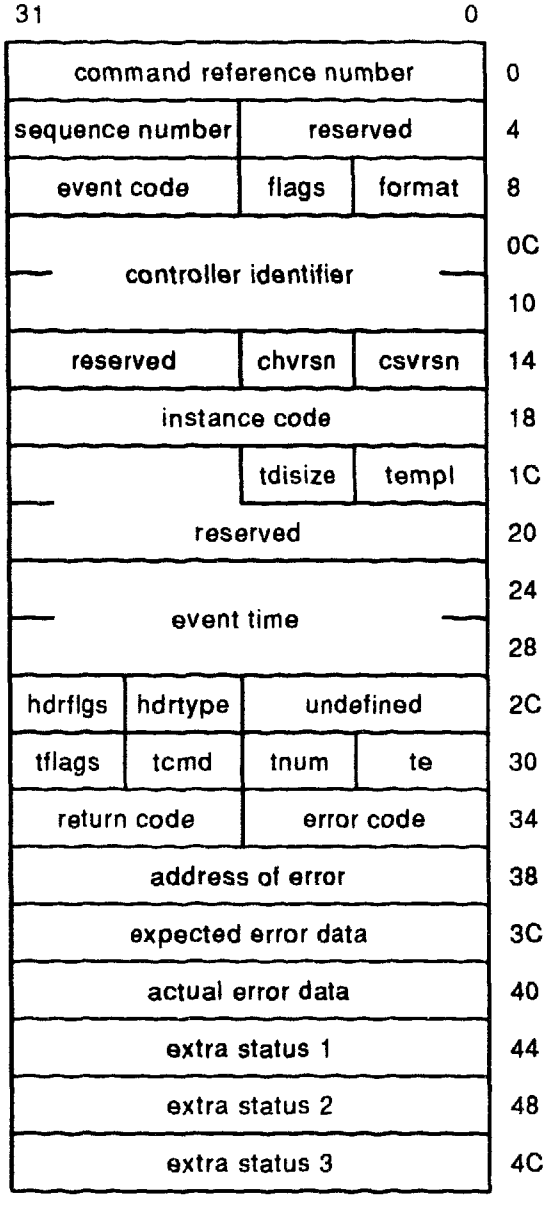
This field contains the value 0.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–23.

Figure C–21 Subsystem Built-In Self Test Failure Event Log (Template 13) Format



templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 13 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 24 for this event log.

reserved (offset 1E)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

undefined

This field is only present to provide longword alignment, its content is undefined.

hdrtype
hdrflgs
te
tnum
tcmd
tflags
error code
return code
address of error
expected error data
actual error data
extra status 1
extra status 2
extra status 3

The content of these fields varies depending on the HS1CP controller Subsystem Built-in Self Test that detected the error condition and the error condition that was detected.

C.3.2.7 Memory System Failure Event Log (Template 14)

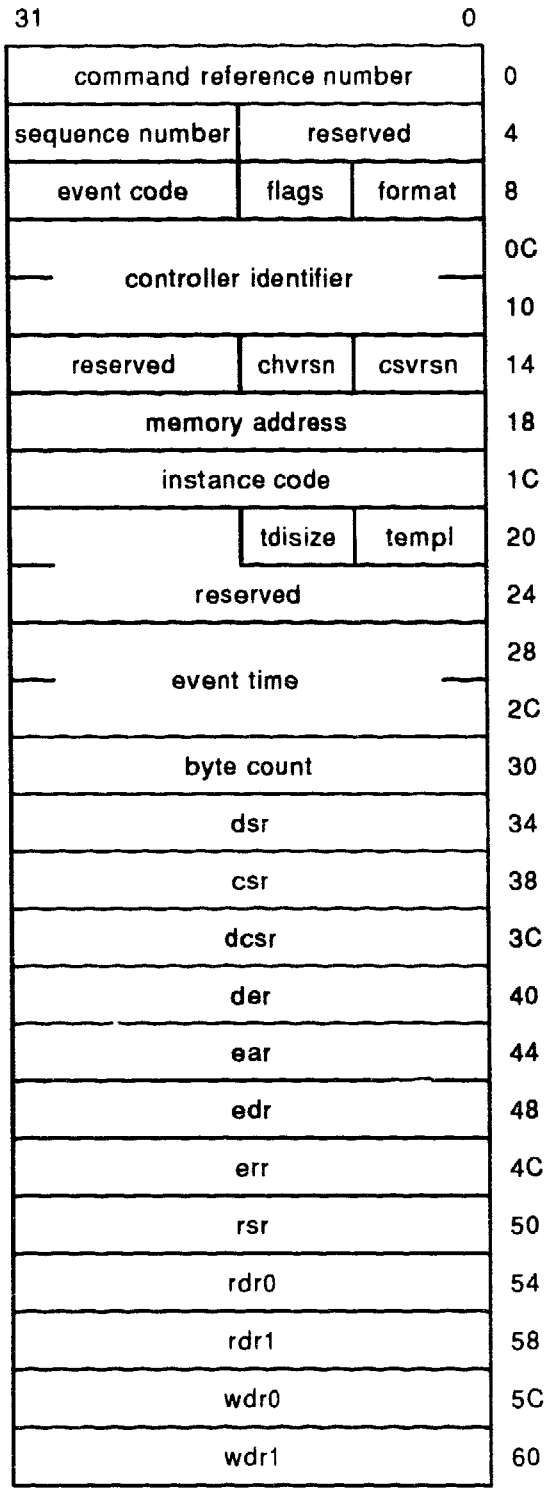
The HS1CP controller Executive firmware component and the Cache Manager, part of the Value Added firmware component, report the occurrence of memory errors via the Memory System Failure Event Log.

The Memory System Failure Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Memory System Failure Event Log is reported via the T/MSCP Memory Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–22.

Memory System Failure Event Log Format Specific Fields:

Figure C–22 Memory System Failure Event Log (Template 14) Format



format

This field contains the value 01 (that is, T/MSCP Memory Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–24.

memory address

The content of this field depends on the value supplied in the "instance code" field. See Table C–24 for more detail.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–24.

templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 14 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 34 for this event log.

reserved (offset 22)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

byte count

The number of bytes contained in the bad memory area (that is, the area bounded by: "memory address" through "memory address" + "byte count" - 1).

dsr
csr
dcsr
der
ear
edr
err
rsr

These fields contain the values contained in the registers of the DRAB that detected the memory failure.

rdr0
rdr1
wdr0
wdr1

These fields contain the values contained in the HS1CP controller's Read and Write Diagnostic registers.

Note that the content of certain of the fields described above may be undefined depending on the value supplied in the "instance code" field. See Table C–24 for more detail.

C.3.2.8 Internal Bus Port/Port Driver Event Log (Template 32)

The HS1CP controller Host Interconnect Services firmware component reports errors detected while performing work related to the Internal Bus Port/Port Driver (PPD) communication layer via the Internal Bus Port/Port Driver Event Log.

The Internal Bus Port/Port Driver Event Log is be sent to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Internal Bus Port/Port Driver Event Log is reported via the T/MSCP Controller Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–23.

Internal Bus Port/Port Driver Event Log Format Specific Fields:

format

This field contains the value 00 (that is, T/MSCP Controller Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–25.

reserved (offset 16)

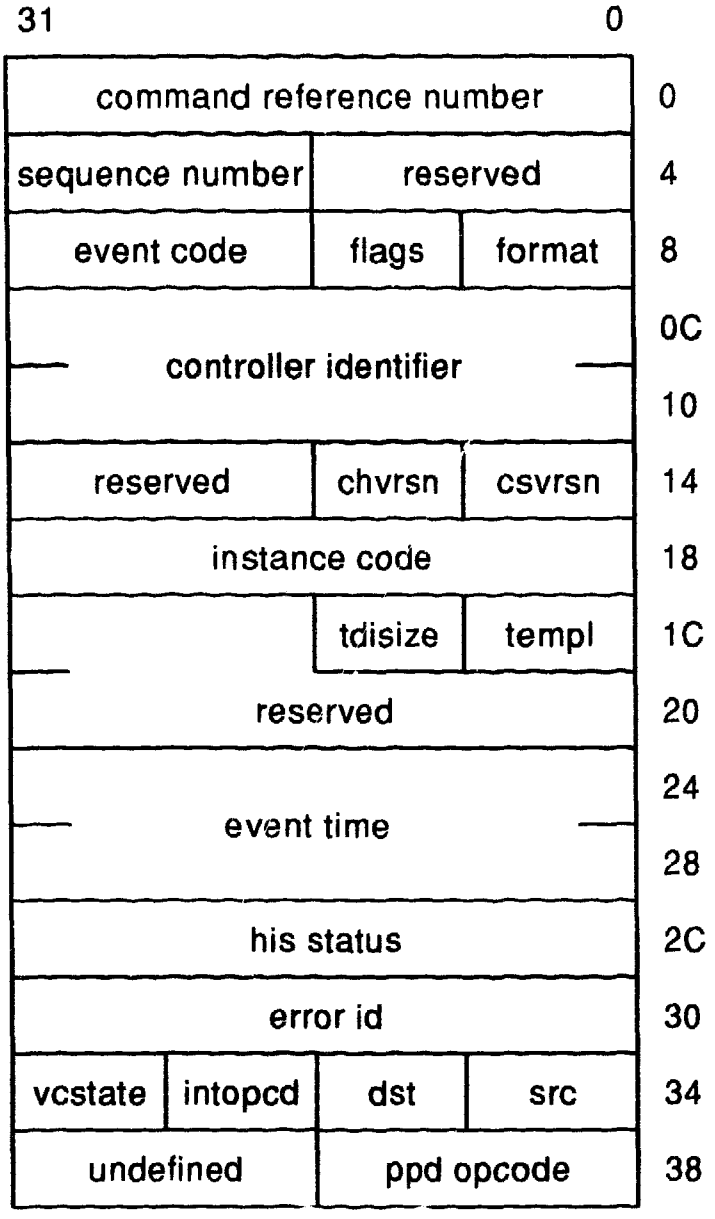
This field contains the value 0.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–25.

Figure C–23 Internal Bus Port/Port Driver Event Log (Template 32) Format



templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 32 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 10 for this event log.

reserved (offset 1E)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

his status
error id
src
dst
intopcd
vcstate
ppd opcode

See Section C.3.1.1 for the description of these fields.

undefined

This field is only present to provide longword alignment, its content is undefined.

Note that the content of certain of the fields described above may be undefined depending on the value supplied in the "instance code" field. See Table C–25 for more detail.

C.3.2.9 Internal Bus System Communication Services Event Log (Template 33)

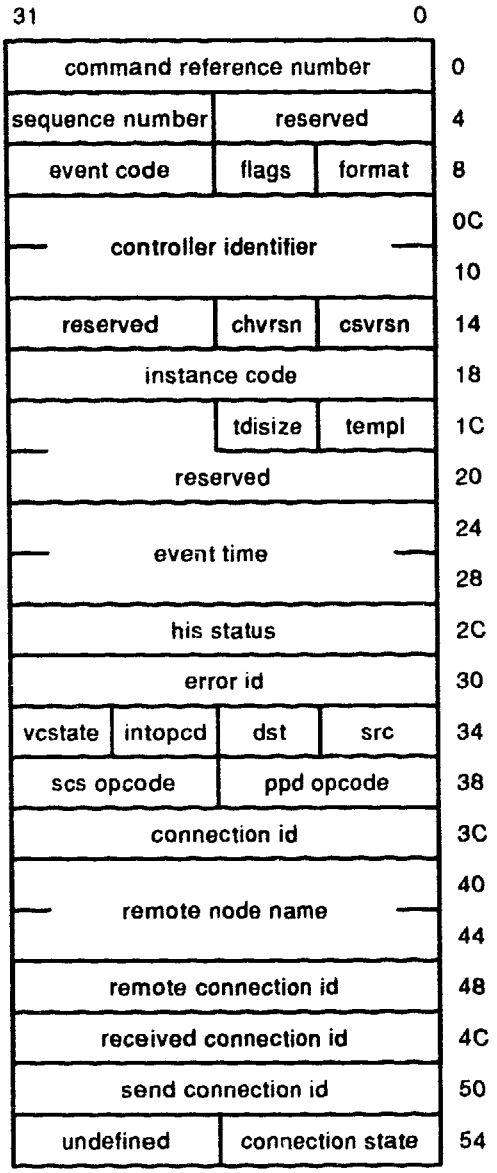
The HS1CP controller Host Interconnect Services firmware component reports errors detected while performing work related to the Internal Bus System Communication Services (SCS) communication layer via the Internal Bus System Communication Services Event Log.

The Internal Bus Communication Services Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Internal Bus System Communication Services Event Log is reported via the T/MSCP Controller Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–24.

Internal Bus System Communication Services Event Log Format Specific Fields:

Figure C–24 Internal Bus System Communication Services Event Log (Template 33) Format



format

This field contains the value 00 (that is, T/MSCP Controller Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–26.

reserved (offset 16)

This field contains the value 0.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–26.

templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 33 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 2C for this event log.

reserved (offset 1E)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

hia status
error id
src
dst
intopcd
vcstate
ppd opcode
scs opcode

See Section C.3.1.1 for the description of these fields.

connection id
remote node name

See Section C.3.1.2 for the description of these fields.

remote connection id

The remote connection identifier supplied by the host node.

received connection id

The connection identifier of the System Application (SYSAP) that is receiving the message contained in the Host Transaction Block.

send connection id

The connection identifier of the System Application (SYSAP) that is sending the message contained in the Host Transaction Block.

connection state

The connection state code as shown in Table C–8.

undefined

This field is present only to provide longword alignment; its content is undefined.

Note that the content of certain of the fields described above may be undefined depending on the value supplied in the "instance code" field. See Table C–26 for more detail.

C.3.2.10 Device Services Nontransfer Error Event Log (Template 41)

The HS1CP controller Device Services firmware component reports errors detected while performing nontransfer work related to disk, tape, or media loader device operations via the Device Services Nontransfer Event Log.

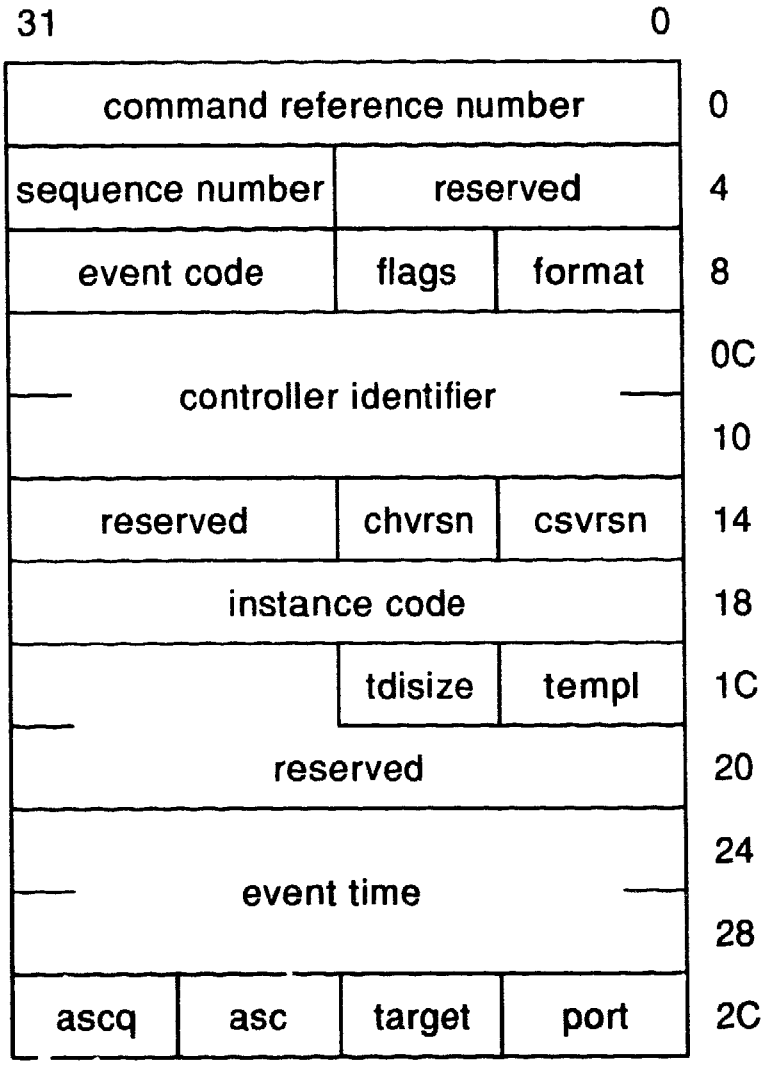
If the error is associated with a command issued by a host system, the Device Services Nontransfer Error Event Log will be sent to the host system that issued the command on the same connection upon which the command was received if "This Host" error logging is enabled on that connection, and to all host systems that have enabled "Other Host" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

If the error is associated with a command issued by an HS1CP controller firmware component, the Device Services Nontransfer Error Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection established with the HS1CP controller's Disk and/or Tape MSCP Server.

The Device Services Nontransfer Error Event Log is reported via the T/MSCP Controller Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–25.

Device Services Nontransfer Error Event Log Format Specific Fields:

Figure C–25 Device Services Nontransfer Error Event Log (Template 41) Format



format

This field contains the value 00 (that is, T/MSCP Controller Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–27.

reserved (offset 16)

This field contains the value 0.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–27.

templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 41 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 04 for this event log.

reserved (offset 1E)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

port

The SCSI bus number affected by the error being reported.

target

The SCSI target number on the "port" affected by the error being reported.

asc
ascq

The "asc" and "ascq" fields contain the values supplied in byte 0C (Additional Sense Code) and byte 0D (Additional Sense Code Qualifier) fields, respectively, of the Sense Data returned in the response of a SCSI REQUEST SENSE command issued to the target device. The description of the value supplied in the "instance code" field (see Table C–27) describes the Sense Key value supplied in the Sense Data returned.

Note that the content of certain of the fields described above may be undefined depending on the value supplied in the "instance code" field. See Table C–27 for more detail.

C.3.2.11 Disk Transfer Error Event Log (Template 51)

The HS1CP controller Device Services and Value Added Services firmware components report errors detected while performing work related to disk unit transfer operations via the Disk Transfer Error Event Log.

If the error is associated with a command issued by a host system, the Disk Transfer Error Event Log will be sent to the host system that issued the command on the same connection upon which the command was received if "This Host" error logging is enabled on that connection and to all host systems that have enabled "Other Host" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

If the error is associated with a command issued by an HS1CP controller firmware component, the Disk Transfer Error Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection established with the HS1CP controller's Disk MSCP Server.

The Disk Transfer Error Event Log is reported via the MSCP Disk Transfer Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–26.

Disk Transfer Error Event Log Format Specific Fields:

format

This field contains the value 02 (that is, MSCP Disk Transfer Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–28.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–28.

templ

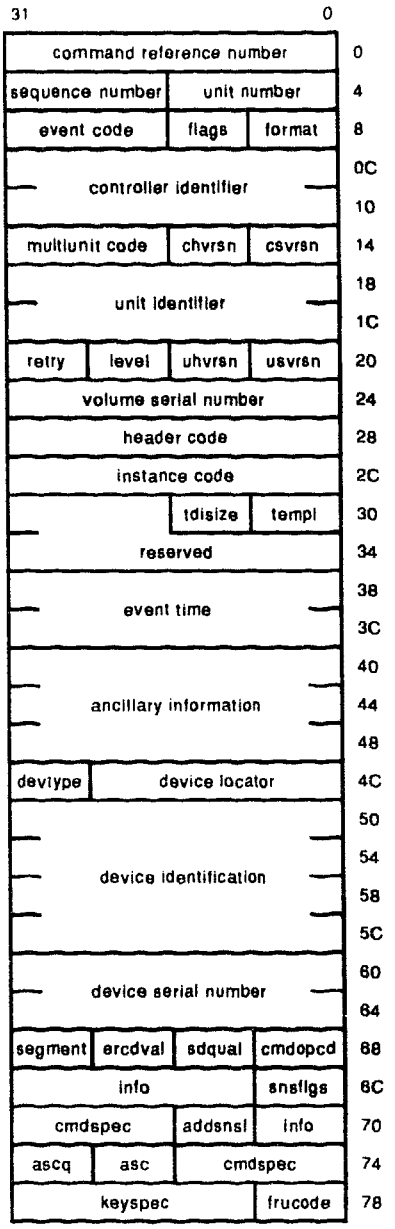
See Implementation Dependent Information Area for the description of this field.

This field contains the value 51 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

Figure C–26 Disk Transfer Error Event Log (Template 51) Format



This field contains the value 3C for this event log.

reserved (offset 32)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

ancillary information

The format of this field varies depending on whether the event being reported is associated with a command issued by a host system or one issued by an HS1CP controller firmware component.

If the event is associated with a command issued by a host system, this field is formatted as described in Section C.3.1.2.

If the event is associated with a command issued by an HS1CP controller firmware component, this field is formatted as described in Section C.3.1.3.

device locator
devtype
device identification
device serial number

See Section C.3.1.4 for the description of these fields.

cmdopcd
infoq
ercdval
segment
snsflgs
info
addsnsl
cmdspec
asc
ascq
frucode
keyspec

See Section C.3.1.5 for the description of these fields.

Note that the content of certain of the fields described above may be undefined depending on the value supplied in the "instance code" field. See Table C–28 for more detail.

C.3.2.12 Disk Bad Block Replacement Attempt Event Log (Template 57)

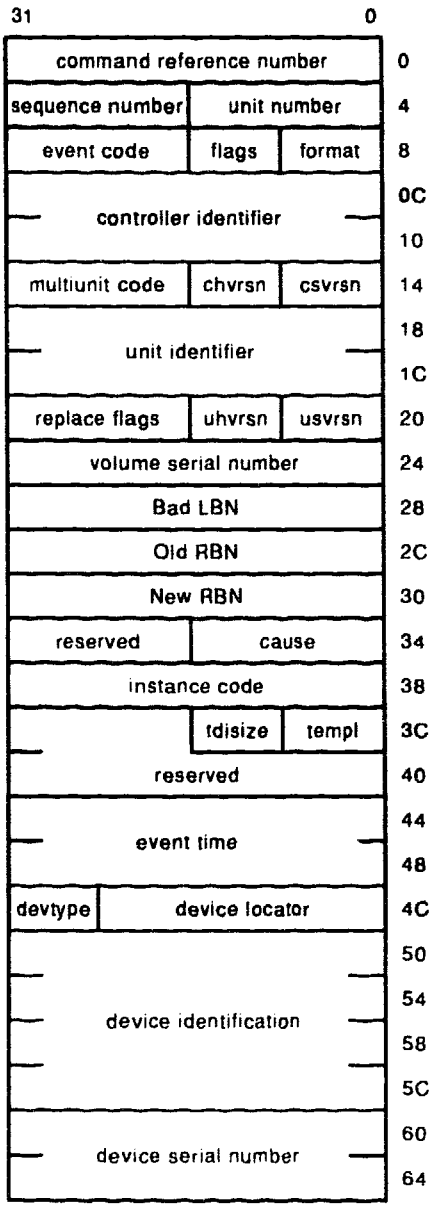
The HS1CP controller Value Added firmware component reports disk unit bad block replacement attempt results via the Disk Bad Block Replacement Attempt Event Log.

If the replacement is associated with a command issued by a host system, the Disk Bad Block Replacement Attempt Event Log will be sent to the host system that issued the command on the same connection upon which the command was received if "This Host" error logging is enabled on that connection, and to all host systems that have enabled "Other Host" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

If the replacement is associated with a command issued by an HS1CP controller firmware component, the Disk Bad Block Replacement Attempt Error Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection established with the HS1CP controller's Disk MSCP Server.

The Disk Bad Block Replacement Attempt Event Log is reported via the MSCP Bad Block Replacement Attempt error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–27.

Figure C–27 Disk Bad Block Replacement Attempt Event Log (Template 57) Format



Disk Bad Block Replacement Attempt Event Log Format Specific Fields:

format

This field contains the value 09 (that is, MSCP Bad Block Replacement Attempt error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–29.

reserved (offset 36)

This field contains the value 0.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–29.

templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 57 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 1C for this event log.

reserved (offset 3E)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

device locator
devtype
device identification
device serial number

See Section C.3.1.4 for the description of these fields.

Note that the content of certain of the fields described above may be undefined depending on the value supplied in the "instance code" field. See Table C–29 for more detail.

C.3.2.13 Tape Transfer Error Event Log (Template 61)

The HS1CP controller Device Services and Value Added Services firmware components report errors detected while performing work related to tape unit transfer operations via the Tape Transfer Error Event Log.

If the error is associated with a command issued by a host system, the Tape Transfer Error Event Log will be sent to the host system that issued the command on the same connection upon which the command was received if "This Host" error logging is enabled on that connection, and to all host systems that have enabled "Other Host" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

If the error is associated with a command issued by an HS1CP controller firmware component, the Tape Transfer Error Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection established with the HS1CP controller's Tape MSCP Server.

The Tape Transfer Error Event Log is reported via the TMSCP Tape Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–28.

Tape Transfer Error Event Log Format Specific Fields:

format

This field contains the value 05 (that is, TMSCP Tape Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–30.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–30.

templ

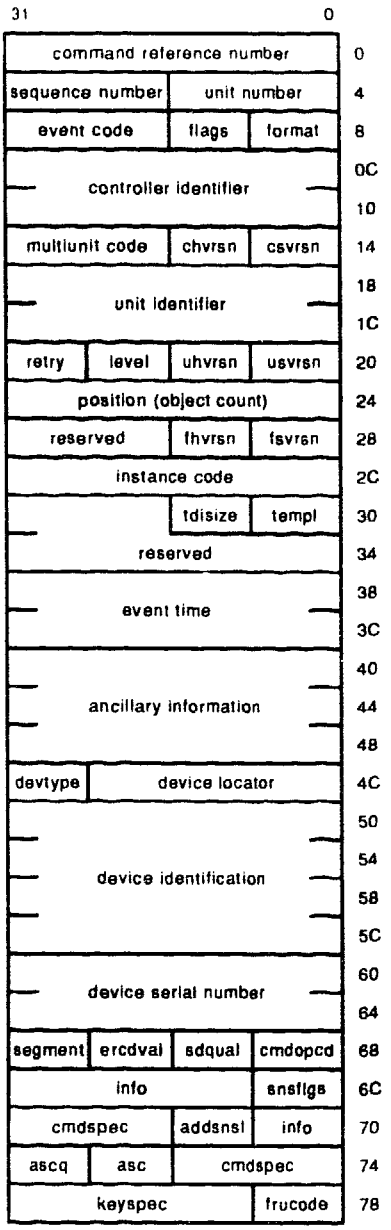
See Implementation Dependent Information Area for the description of this field.

This field contains the value 61 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

Figure C–28 Tape Transfer Error Event Log (Template 61) Format



This field contains the value 3C for this event log.

reserved (offset 32)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

ancillary information

The format of this field varies depending on whether the event being reported is associated with a command issued by a host system or one issued by an HS1CP controller firmware component.

If the event is associated with a command issued by a host system, this field is formatted as described in Section C.3.1.2.

If the event is associated with a command issued by an HS1CP controller firmware component, this field is considered "reserved" and contains the value 0.

device locator
devtype
device identification
device serial number

See Section C.3.1.4 for the description of these fields.

cmdopcd
infoq
ercdval
segment
snsflgs
info
addsnsl
cmdspec
asc
ascq
frucode
keyspec

See Section C.3.1.5 for the description of these fields.

Note that the content of certain of the fields described above may be undefined depending on the value supplied in the "instance code" field. See Table C–30 for more detail.

C.3.2.14 Media Loader Error Event Log (Template 71)

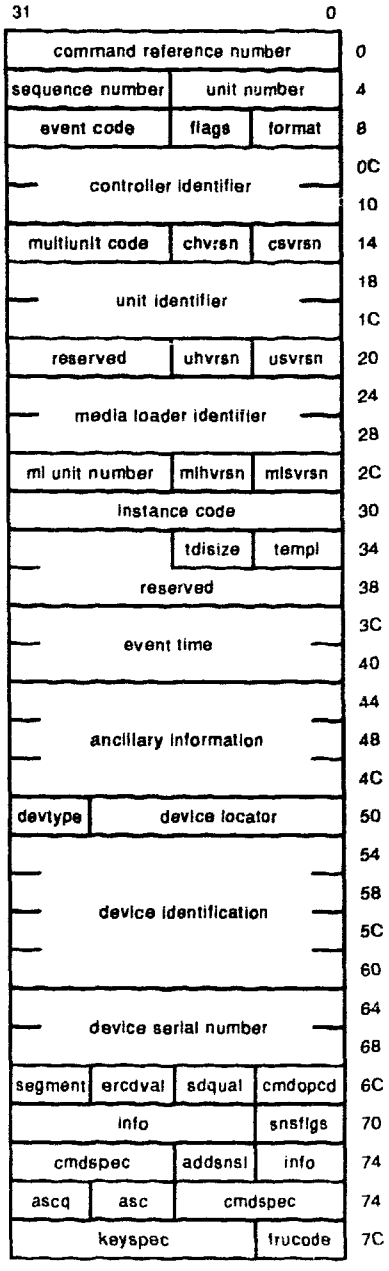
The HS1CP controller Device Services firmware component reports errors detected while performing work related to media loader operations via the Media Loader Error Event Log.

If the error is associated with a command issued by a host system, the Media Loader Error Event Log will be sent to the host system that issued the command on the same connection upon which the command was received if "This Host" error logging is enabled on that connection, and to all host systems that have enabled "Other Host" error logging on a connection or connections established with the HS1CP controller's Disk and/or Tape MSCP Server.

If the error is associated with a command issued by an HS1CP controller firmware component, the Media Loader Error Event Log will be sent to all host systems that have enabled "Miscellaneous" error logging on a connection established with the HS1CP controller's Tape MSCP Server.

The Media Loader Error Event Log is reported via the T/MSCP Media Loader Errors error log message format. The format of this event log, including the HS1CP controller specific fields, is shown in Figure C–29.

Figure C–29 Media Loader Error Event Log (Template 71) Format



Media Loader Error Event Log Format Specific Fields:

format

This field contains the value 0A (that is, T/MSCP Media Loader Errors error log format code).

event code

The values that can be reported in this field for this event log are shown in Table C–31.

instance code

See Implementation Dependent Information Area for the description of this field.

The values that can be reported in this field for this event log are shown in Table C–31.

templ

See Implementation Dependent Information Area for the description of this field.

This field contains the value 71 for this event log.

tdisize

See Implementation Dependent Information Area for the description of this field.

This field contains the value 3C for this event log.

reserved (offset 36)

This field contains the value 0.

event time

See Implementation Dependent Information Area for the description of this field.

ancillary information

The format of this field varies depending on whether the event being reported is associated with a command issued by a host system or one issued by an HS1CP controller firmware component.

If the event is associated with a command issued by a host system, this field is formatted as described in Section C.3.1.2.

If the event is associated with a command issued by an HS1CP controller firmware component, this field is considered "reserved" and contains the value 0.

device locator
devtype
device identification

device serial number

See Section C.3.1.4 for the description of these fields.

cmdopcd
infoq
ercdval
segment
snsflgs
info
addsnsl
cmdspec
asc
ascq
frucode
keyspec

See Section C.3.1.5 for the description of these fields.

C.3.2.15 Disk Copy Data Correlation Event Log

The HS1CP controller disk MSCP server firmware component reports errors detected while performing disk copy data commands via the disk copy data correlation event log.

The format of the Disk Copy Data Correlation Event Log is identical to the format of the MSCP Disk Copy Data Correlation error log message.

If a Controller Error (subcode "Local Connection Request Failed, Insufficient Resources to Request Local Connection") or a Controller Error (subcode "Remote Connection Request Failed, Insufficient Resources to Request Remote Connection") condition is detected, the HS1CP controller will store one of values shown in Table C–32 in the first longword of the event dependent information field of the MSCP Disk Copy Data Correlation error log message to identify the resource that is lacking.

C.4 Event Log Codes

Tables C–2 through C–51 list specific codes contained within the event log information.

Table C–2 Firmware Component Identifier Codes

Code	Description
01	Executive Services
02	Value Added Services
03	Device Services
04	Fault Manager

(continued on next page)

Table C–2 (Cont.) Firmware Component Identifier Codes

Code	Description
05	Common Library Routines
06	Dual Universal Asynchronous Receiver/Transmitter Services
07	Failover Control
08	Nonvolatile Parameter Memory Failover Control
20	Command Line Interpreter
40	Host Interconnect Services
42	Host Interconnect Port Services
60	Disk and Tape MSCP Server
61	Diagnostics and Utilities Protocol Server
62	System Communication Services Directory Service
80	Disk Inline Exerciser (DILX)
81	Tape Inline Exerciser (TILX)
82	Subsystem Built-In Self Tests (BIST)
83	Device Configuration Utilities (CONFIG/CFMENU)
84	Clone Utility(CLONE)

Table C–3 Host Interconnect Services Status Codes

Code	Description
00000000	Request succeeded.
00000001	The remote sent a message over a connection that has been invalidated.
00000002	The remote sent a message for which no receive credit is available.
00000003	Received a message from the remote while in an invalid or illegal connection state.
00000004	Pending work exists but connection state is invalid or illegal.
00000009	Request failed, no additional information available.
00000032	A PPD message was received from the remote but the Virtual Circuit is in an invalid or illegal state.
00000033	A PPD START was received from the remote but the Virtual Circuit state indicates that the Virtual Circuit is already OPEN.
00000034	A PPD NODE_STOP was received from the remote.
00000035	The "PPD START send without receiving a PPD START in response" limit has been reached; the remote node is acknowledging the packets but not responding to them.
00000036	The "PPD STACK send without receiving a PPD ACK in response" limit has been reached; the remote node is acknowledging the packets but not responding to them.
00000064	The "CI IDREQ send without receiving a CI ID in response" limit has been reached on both Path A and Path B; the remote node is acknowledging the packets but not responding to them.

(continued on next page)

Table C–3 (Cont.) Host Interconnect Services Status Codes

Code	Description
00000065	A CI ID or CI CNF packet (transmitted by the thread on behalf of Host Interconnect Services) could not be successfully transmitted.
00010009	VC closed due to CI ID request failure.
00020009	VC closed due to unexpected SCS state.
00030009	VC closed due to CI START failure.
00040009	VC closed due to CI STACK failure.
00050009	VC closed due to PPD ACK failure.
00060009	VC closed due to PPD NODE_STOP or PPD START message received.
00070009	VC closed due to NAK ADP retry CI ID transmit failure.
00080009	VC closed due to NAK ADP retry transmit failure.
00090009	VC closed due to NOR DDL retry transmit failure on Path A.
000A0009	VC closed due to NOR DDL retry transmit failure on Path B.
000B0009	VC closed due to NOR ADP retry CI ID transmit failure.
000C0009	VC closed due to NOR ADP retry transmit failure.
000D0009	VC closed due to NAK DDL retry transmit failure on Path A.
000E0009	VC closed due to NAK DDL retry transmit failure on Path B.
000F0009	VC closed due to arbitration timeout on Path A.
00100009	VC closed due to arbitration timeout on Path B.
00110009	VC closed due to Path A off.
00120009	VC closed due to Path B off.
00130009	VC closed due to dual receive.
00140009	VC closed due to invalid receive data structure state.
00150009	VC closed due to no path.
00160009	VC closed due to message transmit closed.
00170009	VC closed due to data transmit closed.
00180009	VC closed due to message scan.
00190009	VC closed due to data scan.
001A0009	VC closed due to data timeout.
001B0009	VC closed due to unrecognized packet.
001C0009	VC closed due to data transmit failure.
001D0009	VC closed due to CI ID complete failure.
001E0009	VC closed due to lost command.
001F0009	Not implemented in CI environment.
00200009	VC closed due to sequence number mismatch.

Table C–4 CI Message Operation Codes

Code	Description
00	Reserved
01	DG
02	MSG
03	CNF
04	MCNF
05	IDREQ
06	RST
07	STRT
08	DATREQ0
09	DATREQ1
0A	DATREQ2
0B	ID
0C	PSREQ
0D	LB
0E	MDATREQ
0F	RETPS
10	SNTDAT
11	RETDAT
12	SNTMDAT
13	RETMDAT

Table C–5 CI Virtual Circuit State Codes

Code	Description
01	VC_CLOSED
02	START_SENT
03	START_REC
04	VC_OPEN
05	VC_CLOSING

Table C–6 Port/Port Driver Message Operation Codes

Code	Description
0000	START
0001	STACK
0002	ACK
0003	SCS_DG
0004	SCS_MSG

(continued on next page)

Table C–6 (Cont.) Port/Port Driver Message Operation Codes

Code	Description
0005	ERROR_LOG
0006	NODE_STOP

Table C–7 System Communication Services Message Operation Codes

Code	Description
0000	CONNECT_REQ
0001	CONNECT_RSP
0002	ACCEPT_REQ
0003	ACCEPT_RSP
0004	REJECT_REQ
0005	REJECT_RSP
0006	DISCONNECT_REQ
0007	DISCONNECT_RSP
0008	CREDIT_REQ
0009	CREDIT_RSP
000A	APPL_MSG
000B	APPL_DG

Table C–8 CI Connection State Codes

Code	Description
0000	CLOSED
0001	LISTENING
0002	CONNECT_SENT
0003	CONNECT_ACK
0004	CONNECT_REC
0005	ACCEPT_SENT
0006	REJECT_SENT
0007	OPEN
0008	DISCONNECT_SENT
0009	DISCONNECT_REC
000A	DISCONNECT_ACK
000B	DISCONNECT_MATCH

Table C–9 Supported SCSI Device Type Codes

Code	Description
00	Direct-Access Devices (such as magnetic disk).

(continued on next page)

Table C–9 (Cont.) Supported SCSI Device Type Codes

Code	Description
01	Sequential-Access Devices (such as magnetic tape).
05	CD-ROM Devices.
08	Medium Changer Devices (such as jukeboxes).

Table C–10 SCSI Command Operation Codes

Code	Supported Device Types (See Table C–9)	Description
00	00, 01, 05, 08	TEST UNIT READY
01	01	REWIND
01	00, 05, 08	REZERO UNIT
03	00, 01, 05, 08	REQUEST SENSE
04	00	FORMAT UNIT
05	01	READ BLOCK LIMITS
07	08	INITIALIZE ELEMENT STATUS
07	00	REASSIGN BLOCKS
08	00, 01, 05	READ (6 byte)
0A	00, 01	WRITE (6 byte)
0B	00, 05	SEEK (6 byte)
0F	01	READ REVERSE
10	01	WRITE FILEMARKS
11	01	SPACE
12	00, 01, 05, 08	INQUIRY
13	01	TAPE VERIFY
14	01	RECOVER BUFFERED DATA
15	00, 01, 05, 08	MODE SELECT (6 byte)
16	00, 01, 05, 08	RESERVE UNIT
17	00, 01, 05, 08	RELEASE UNIT
18	00, 01, 05	COPY
19	01	ERASE
1A	00, 01, 05, 08	MODE SENSE (6 byte)
1B	00, 05	START STOP UNIT
1B	01	LOAD UNLOAD
1C	00, 01, 05, 08	RECEIVE DIAGNOSTIC RESULTS
1D	00, 01, 05, 08	SEND DIAGNOSTIC
1E	00, 01, 05, 08	PREVENT-ALLOW MEDIUM REMOVAL
25	00, 05	READ CAPACITY
28	00, 05	READ (10 byte)

(continued on next page)

Table C–10 (Cont.) SCSI Command Operation Codes

Code	Supported Device Types (See Table C–9)	Description
2A	00	WRITE (10 byte)
2B	08	POSITION TO ELEMENT
2B	01	LOCATE
2B	00, 05	SEEK (10 byte)
2E	00	WRITE AND VERIFY (10 byte)
2F	00, 05	VERIFY (10 byte)
30	00, 05	SEARCH DATA HIGH (10 byte)
31	00, 05	SEARCH DATA EQUAL (10 byte)
32	00, 05	SEARCH DATA LOW (10 byte)
33	00, 05	SET LIMITS (10 byte)
34	01	READ POSITION
34	00, 05	PRE-FETCH
35	00, 05	SYNCHRONIZE CACHE
36	00, 05	LOCK-UNLOCK CACHE
37	00	READ DEFECT DATA (10 byte)
39	00, 01, 05	COMPARE
3A	00, 01, 05	COPY AND VERIFY
3B	00, 01, 05, 08	WRITE BUFFER
3C	00, 01, 05, 08	READ BUFFER
3E	00, 05	READ LONG
3F	00	WRITE LONG
40	00, 01, 05, 08	CHANGE DEFINITION
41	00	WRITE SAME
42	05	READ SUB-CHANNEL
43	05	READ TOC (table of contents)
44	05	READ HEADER
45	05	PLAY AUDIO (10 byte)
47	05	PLAY AUDIO MSF
48	05	PLAY AUDIO TRACK/INDEX
49	05	PLAY TRACK RELATIVE (10 byte)
4B	05	PAUSE/RESUME
4C	00, 01, 05, 08	LOG SELECT
4D	00, 01, 05, 08	LOG SENSE
55	00, 01, 05, 08	MODE SELECT (10 byte)
5A	00, 01, 05, 08	MODE SENSE (10 byte)
A5	05	PLAY AUDIO (12 byte)
A5	08	MOVE MEDIUM

(continued on next page)

Table C–10 (Cont.) SCSI Command Operation Codes

Code	Supported Device Types (See Table C–9)	Description
A6	08	EXCHANGE MEDIUM
A8	05	READ (12 byte)
A9	05	PLAY TRACK RELATIVE (12 byte)
AF	05	VERIFY (12 byte)
B0	05	SEARCH DATA HIGH (12 byte)
B1	05	SEARCH DATA EQUAL (12 byte)
B2	05	SEARCH DATA LOW (12 byte)
B3	05	SET LIMITS (12 byte)
B5	08	REQUEST VOLUME ELEMENT ADDRESS
B6	08	SEND VOLUME TAG
B8	08	READ ELEMENT STATUS

Table C–11 SCSI Buffered Modes Codes

Code	Description
0	The target shall not report GOOD status on write commands until the data blocks are actually written on the medium.
1	The target may report GOOD status on write commands as soon as as all the data specified in the write command has been transferred to the target's buffer. One or more blocks may be buffered prior to writing the blocks to the medium.
2	The target may report GOOD status on write commands as soon as: (1) All the data specified in the write command has been successfully transferred to the target's buffer, and (2) All buffered data from different initiators has been successfully written to the medium.
3	Reserved for future use.
4	Reserved for future use.
5	Reserved for future use.
6	Reserved for future use.
7	Reserved for future use.

Table C–12 SCSI Sense Key Codes

Code	Description
0	NO SENSE. Indicates that there is no specific sense key information to be reported for the designated logical unit. This would be the case for a successful command or a command that received CHECK CONDITION or COMMAND TERMINATED status because one of the FM, EOM, or ILI bits is set to one in the "snsflgs" field.
1	RECOVERED ERROR. Indicates that the last command completed successfully with some recovery action performed by the target. Details may be determinable by examining the "info" field.

(continued on next page)

Table C–12 (Cont.) SCSI Sense Key Codes

Code	Description
2	NOT READY. Indicates that the logical unit addressed cannot be accessed. Operator intervention may be required to correct this condition.
3	MEDIUM ERROR. Indicates that the command terminated with a nonrecovered error condition that was probably caused by a flaw in the medium or an error in the recorded data. This sense key may also be returned if the target is unable to distinguish between a flaw in the medium and a specific hardware failure (sense key 4).
4	HARDWARE ERROR. Indicates that the target detected a nonrecoverable hardware failure (for example, controller failure, device failure, parity error, and so forth) while performing the command or during a self test.
5	ILLEGAL REQUEST. Indicates that there was an illegal parameter in the command descriptor block or in the additional parameters supplied as data for some commands (FORMAT UNIT, SEARCH DATA, and so forth). If the target detects an invalid parameter in the command descriptor block, then it shall terminate the command without altering the medium. If the target detects an invalid parameter in the additional parameters supplied as data, then the target may have already altered the medium. This sense key may also indicate that an invalid IDENTIFY message was received.
6	UNIT ATTENTION. Indicates that the removable medium may have been changed or the target has been reset.
7	DATA PROTECT. Indicates that a command that reads or writes the medium was attempted on a block that is protected from this operation. The read or write operation is not performed.
8	BLANK CHECK. Indicates that a write-once device or a sequential-access device encountered blank medium or format-defined end-of-data indication while reading or a write-once device encountered a nonblank medium while writing.
9	Vendor Specific. This sense key is available for reporting vendor specific conditions.
A	COPY ABORTED. Indicates a COPY, COMPARE, or COPY AND VERIFY command was aborted due to an error condition on the source device, the destination device, or both.
B	ABORTED COMMAND. Indicates that the target aborted the command. The initiator may be able to recover by trying the command again.
C	EQUAL. Indicates a SEARCH DATA command has satisfied an equal comparison.
D	VOLUME OVERFLOW. Indicates that a buffered peripheral device has reached the end-of-partition and data may remain in the buffer that has not been written to the medium. A RECOVER BUFFERED DATA commands may be issued to read the unwritten data from the buffer.
E	MISCOMPARE. Indicates that the source data did not match the data read from the medium.
F	RESERVED.

Table C–13 SCSI ASC/ASCQ Codes for Direct-Access Devices (Such as Magnetic Disk)

ASC Code	ASCQ Code	Description
00	00	No additional sense information.

(continued on next page)

Table C–13 (Cont.) SCSI ASC/ASCQ Codes for Direct-Access Devices (Such as Magnetic Disk)

ASC Code	ASCQ Code	Description
00	06	I/O process terminated.
01	00	No index/sector signal.
02	00	No seek complete.
03	00	Peripheral device write fault.
04	00	Logical unit not ready, cause not reportable.
04	01	Logical unit is in process of becoming ready.
04	02	Logical unit not ready, initializing command required.
04	03	Logical unit not ready, manual intervention required.
04	04	Logical unit not ready, format in progress.
06	00	No reference position found.
07	00	Multiple peripheral devices selected.
08	00	Logical unit communication failure.
08	01	Logical unit communication time-out.
08	02	Logical unit communication parity error.
09	00	Track following error.
0A	00	Error log overflow.
0C	01	Write error recovered with auto reallocation.
0C	02	Write error - auto reallocation failed.
10	00	Id crc or ecc error.
11	00	Unrecovered read error.
11	01	Read retries exhausted.
11	02	Error too long to correct.
11	03	Multiple read errors.
11	04	Unrecovered read error - auto reallocate failed.
11	0A	Miscorrected error.
11	0B	Unrecovered read error - recommend reassignment.
11	0C	Unrecovered read error - recommend rewrite the data.
12	00	Address mark not found for id field.
13	00	Address mark not found for data field.
14	00	Recorded entity not found.
14	01	Record not found.
15	00	Random positioning error.
15	01	Mechanical positioning error.
15	02	Positioning error detected by read of medium.
16	00	Data synchronization mark error.
17	00	Recovered data with no error correction applied.

(continued on next page)

Table C–13 (Cont.) SCSI ASC/ASCQ Codes for Direct-Access Devices (Such as Magnetic Disk)

ASC Code	ASCQ Code	Description
17	01	Recovered data with retries.
17	02	Recovered data with positive head offset.
17	03	Recovered data with negative head offset.
17	05	Recovered data using previous sector id.
17	06	Recovered data without ecc - data auto-reallocated.
17	07	Recovered data without ecc - recommend reassignment.
17	08	Recovered data without ecc - recommend rewrite.
18	00	Recovered data with error correction applied.
18	01	Recovered data with error correction & retries applied.
18	02	Recovered data - data auto-reallocated.
18	05	Recovered data - recommend reassignment.
18	06	Recovered data - recommend rewrite.
19	00	Defect list error.
19	01	Defect list not available.
19	02	Defect list error in primary list.
19	03	Defect list error in grown list.
1A	00	Parameter list length error.
1B	00	Synchronous data transfer error.
1C	00	Defect list not found.
1C	01	Primary defect list not found.
1C	02	Grown defect list not found.
1D	00	Miscompare during verify operation.
1E	00	Recovered id with ecc correction.
20	00	Invalid command operation code.
21	00	Logical block address out of range.
22	00	Illegal function (should use 20 00, 24 00, or 26 00).
24	00	Invalid field in cdb.
25	00	Logical unit not supported.
26	00	Invalid field in parameter list.
26	01	Parameter not supported.
26	02	Parameter value invalid.
26	03	Threshold parameters not supported.
27	00	Write protected.
28	00	Not ready to ready transition, medium may have changed.
29	00	Power on, reset, or bus device reset occurred.
29	01	Power on occurred.

(continued on next page)

Table C–13 (Cont.) SCSI ASC/ASCQ Codes for Direct-Access Devices (Such as Magnetic Disk)

ASC Code	ASCQ Code	Description
29	02	SCSI bus reset occurred.
29	03	Bus device reset occurred.
2A	00	Parameters changed.
2A	01	Mode parameters changed.
2A	02	Log parameters changed.
2B	00	Copy cannot execute since host cannot disconnect.
2C	00	Command sequence error.
2F	00	Commands cleared by another initiator.
30	00	Incompatible medium installed.
30	01	Cannot read medium - unknown format.
30	02	Cannot read medium - incompatible format.
30	03	Cleaning cartridge installed.
31	00	Medium format corrupted.
31	01	Format command failed.
32	00	No defect spare location available.
32	01	Defect list update failure.
37	00	Rounded parameter.
39	00	Saving parameters not supported.
3A	00	Medium not present.
3D	00	Invalid bits in identify message.
3E	00	Logical unit has not self-configured yet.
3F	00	Target operating conditions have changed.
3F	01	Microcode has been changed.
3F	02	Changed operating definition.
3F	03	Inquiry data has changed.
40	*nn*	Diagnostic failure detected on component *nn*; where *nn* identifies a specific target device component (*nn* range 80 through FF). Refer to documentation provided by the vendor of the target device for a description of the component identified by *nn*.
40	00	Ram failure (should use 40 80 through 40 FF).
41	00	Data path failure (should use 40 80 through 40 FF).
42	00	Power-on or self-test failure (should use 40 80 through 40 FF).
43	00	Message error.
44	00	Internal target failure.
45	00	Select or reselect failure.
46	00	Unsuccessful soft reset.
47	00	SCSI parity error.

(continued on next page)

Table C–13 (Cont.) SCSI ASC/ASCQ Codes for Direct-Access Devices (Such as Magnetic Disk)

ASC Code	ASCQ Code	Description
48	00	Initiator detected error message received.
49	00	Invalid message error.
4A	00	Command phase error.
4B	00	Data phase error.
4C	00	Logical unit failed self-configuration.
4E	00	Overlapped commands attempted.
53	00	Media load or eject failed.
53	02	Medium removal prevented.
5A	00	Operator request or state change input (unspecified).
5A	01	Operator medium removal request.
5A	02	Operator selected write protect.
5A	03	Operator selected write permit.
5B	00	Log exception.
5B	01	Threshold condition met.
5B	02	Log counter at maximum.
5B	03	Log list codes exhausted.
5C	00	Rpl status change.
5C	01	Spindles synchronized.
5C	02	Spindles not synchronized.

Table C–14 SCSI ASC/ASCQ Codes for Sequential-Access Devices (Such as Magnetic Tape)

ASC Code	ASCQ Code	Description
00	00	No additional sense information.
00	01	Filemark detected.
00	02	End-of-partition/medium detected.
00	03	Setmark detected.
00	04	Beginning-of-partition/medium detected.
00	05	End-of-data detected.
00	06	I/O process terminated.
03	00	Peripheral device write fault.
03	01	No write current.
03	02	Excessive write errors.
04	00	Logical unit not ready, cause not reportable.
04	01	Logical unit is in process of becoming ready.

(continued on next page)

Table C–14 (Cont.) SCSI ASC/ASCQ Codes for Sequential-Access Devices (Such as Magnetic Tape)

ASC Code	ASCQ Code	Description
04	02	Logical unit not ready, initializing command required.
04	03	Logical unit not ready, manual intervention required.
04	04	Logical unit not ready, format in progress.
07	00	Multiple peripheral devices selected.
08	00	Logical unit communication failure.
08	01	Logical unit communication time-out.
08	02	Logical unit communication parity error.
09	00	Track following error.
0A	00	Error log overflow.
0C	00	Write error.
11	00	Unrecovered read error.
11	01	Read retries exhausted.
11	02	Error too long to correct.
11	03	Multiple read errors.
11	08	Incomplete block read.
11	09	No gap found.
11	0A	Miscorrected error.
14	00	Recorded entity not found.
14	01	Record not found.
14	02	Filemark or setmark not found.
14	03	End-of-data not found.
14	04	Block sequence error.
15	00	Random positioning error.
15	01	Mechanical positioning error.
15	02	Positioning error detected by read of medium.
17	00	Recovered data with no error correction applied.
17	01	Recovered data with retries.
17	02	Recovered data with positive head offset.
17	03	Recovered data with negative head offset.
18	00	Recovered data with error correction applied.
1A	00	Parameter list length error.
1B	00	Synchronous data transfer error.
20	00	Invalid command operation code.
21	00	Logical block address out of range.
24	00	Invalid field in cdb.
25	00	Logical unit not supported.

(continued on next page)

Table C–14 (Cont.) SCSI ASC/ASCQ Codes for Sequential-Access Devices (Such as Magnetic Tape)

ASC Code	ASCQ Code	Description
26	00	Invalid field in parameter list.
26	01	Parameter not supported.
26	02	Parameter value invalid.
26	03	Threshold parameters not supported.
27	00	Write protected.
28	00	Not ready to ready transition, medium may have changed.
29	00	Power on, reset, or bus device reset occurred.
29	01	Power on occurred.
29	02	SCSI bus reset occurred.
29	03	Bus device reset occurred.
2A	00	Parameters changed.
2A	01	Mode parameters changed.
2A	02	Log parameters changed.
2B	00	Copy cannot execute since host cannot disconnect.
2C	00	Command sequence error.
2D	00	Overwrite error on update in place.
2F	00	Commands cleared by another initiator.
30	00	Incompatible medium installed.
30	01	Cannot read medium - unknown format.
30	02	Cannot read medium - incompatible format.
30	03	Cleaning cartridge installed.
31	00	Medium format corrupted.
33	00	Tape length error.
37	00	Rounded parameter.
39	00	Saving parameters not supported.
3A	00	Medium not present.
3B	00	Sequential positioning error.
3B	01	Tape position error at beginning-of-medium.
3B	02	Tape position error at end-of-medium.
3B	08	Reposition error.
3D	00	Invalid bits in identify message.
3E	00	Logical unit has not self-configured yet.
3F	00	Target operating conditions have changed.
3F	01	Microcode has been changed.
3F	02	Changed operating definition.
3F	03	Inquiry data has changed.

(continued on next page)

Table C–14 (Cont.) SCSI ASC/ASCQ Codes for Sequential-Access Devices (Such as Magnetic Tape)

ASC Code	ASCQ Code	Description
40	*nn*	Diagnostic failure detected on component *nn*; where *nn* identifies a specific target device component (*nn* range 80 through FF). Refer to documentation provided by the vendor of the target device for a description of the component identified by *nn*.
43	00	Message error.
44	00	Internal target failure.
45	00	Select or reselect failure.
46	00	Unsuccessful soft reset.
47	00	SCSI parity error.
48	00	Initiator detected error message received.
49	00	Invalid message error.
4A	00	Command phase error.
4B	00	Data phase error.
4C	00	Logical unit failed self-configuration.
4E	00	Overlapped commands attempted.
50	00	Write append error.
50	01	Write append position error.
50	02	Position error related to timing.
51	00	Erase failure.
52	00	Cartridge fault.
53	00	Media load or eject failed.
53	01	Unload tape failure.
53	02	Medium removal prevented.
5A	00	Operator request or state change input (unspecified).
5A	01	Operator medium removal request.
5A	02	Operator selected write protect.
5A	03	Operator selected write permit.
5B	00	Log exception.
5B	01	Threshold condition met.
5B	02	Log counter at maximum.
5B	03	Log list codes exhausted.

Table C–15 SCSI ASC/ASCQ Codes for CD-ROM Devices

ASC Code	ASCQ Code	Description
00	00	No additional sense information.
00	06	I/O process terminated.

(continued on next page)

Table C–15 (Cont.) SCSI ASC/ASCQ Codes for CD-ROM Devices

ASC Code	ASCQ Code	Description
00	11	Audio play operation in progress.
00	12	Audio play operation paused.
00	13	Audio play operation successfully completed.
00	14	Audio play operation stopped due to error.
00	15	No current audio status to return.
02	00	No seek complete.
04	00	Logical unit not ready, cause not reportable.
04	01	Logical unit is in process of becoming ready.
04	02	Logical unit not ready, initializing command required.
04	03	Logical unit not ready, manual intervention required.
06	00	No reference position found.
07	00	Multiple peripheral devices selected.
08	00	Logical unit communication failure.
08	01	Logical unit communication time-out.
08	02	Logical unit communication parity error.
09	00	Track following error.
09	01	Tracking servo failure.
09	02	Focus servo failure.
09	03	Spindle servo failure.
0A	00	Error log overflow.
11	00	Unrecovered read error.
11	05	L-ec uncorrectable error.
11	06	Circ unrecovered error.
14	00	Recorded entity not found.
14	01	Record not found.
15	00	Random positioning error.
15	01	Mechanical positioning error.
15	02	Positioning error detected by read of medium.
17	00	Recovered data with no error correction applied.
17	01	Recovered data with retries.
17	02	Recovered data with positive head offset.
17	03	Recovered data with negative head offset.
17	04	Recovered data with retries and/or circ applied.
17	05	Recovered data using previous sector id.
18	00	Recovered data with error correction applied.
18	01	Recovered data with error correction & retries applied.
18	02	Recovered data - data auto-reallocated.

(continued on next page)

Table C–15 (Cont.) SCSI ASC/ASCQ Codes for CD-ROM Devices

ASC Code	ASCQ Code	Description
18	03	Recovered data with circ.
18	04	Recovered data with lec.
18	05	Recovered data - recommend reassignment.
18	06	Recovered data - recommend rewrite.
1A	00	Parameter list length error.
1B	00	Synchronous data transfer error.
20	00	Invalid command operation code.
21	00	Logical block address out of range.
24	00	Invalid field in cdb.
25	00	Logical unit not supported.
26	00	Invalid field in parameter list.
26	01	Parameter not supported.
26	02	Parameter value invalid.
26	03	Threshold parameters not supported.
28	00	Not ready to ready transition, medium may have changed.
29	00	Power on, reset, or bus device reset occurred.
29	01	Power on occurred.
29	02	SCSI bus reset occurred.
29	03	Bus device reset occurred.
2A	00	Parameters changed.
2A	01	Mode parameters changed.
2A	02	Log parameters changed.
2B	00	Copy cannot execute since host cannot disconnect.
2C	00	Command sequence error.
2F	00	Commands cleared by another initiator.
30	00	Incompatible medium installed.
30	01	Cannot read medium - unknown format.
30	02	Cannot read medium - incompatible format.
37	00	Rounded parameter.
39	00	Saving parameters not supported.
3A	00	Medium not present.
3D	00	Invalid bits in identify message.
3E	00	Logical unit has not self-configured yet.
3F	00	Target operating conditions have changed.
3F	01	Microcode has been changed.
3F	02	Changed operating definition.
3F	03	Inquiry data has changed.

(continued on next page)

Table C–15 (Cont.) SCSI ASC/ASCQ Codes for CD-ROM Devices

ASC Code	ASCQ Code	Description
40	*nn*	Diagnostic failure detected on component *nn*; where *nn* identifies a specific target device component (*nn* range 80 through FF). Refer to documentation provided by the vendor of the target device for a description of the component identified by *nn*.
43	00	Message error.
44	00	Internal target failure.
45	00	Select or reselect failure.
46	00	Unsuccessful soft reset.
47	00	SCSI parity error.
48	00	Initiator detected error message received.
49	00	Invalid message error.
4A	00	Command phase error.
4B	00	Data phase error.
4C	00	Logical unit failed self-configuration.
4E	00	Overlapped commands attempted.
53	00	Media load or eject failed.
53	02	Medium removal prevented.
57	00	Unable to recover table-of-contents.
5A	00	Operator request or state change input (unspecified).
5A	01	Operator medium removal request.
5B	00	Log exception.
5B	01	Threshold condition met.
5B	02	Log counter at maximum.
5B	03	Log list codes exhausted.
63	00	End of user area encountered on this track.
64	00	Illegal mode for this track.

Table C–16 SCSI ASC/ASCQ Codes for Medium Changer Devices (Such as Jukeboxes)

ASC Code	ASCQ Code	Description
00	00	No additional sense information.
00	06	I/O process terminated.
02	00	No seek complete.
04	00	Logical unit not ready, cause not reportable.
04	01	Logical unit is in process of becoming ready.
04	02	Logical unit not ready, initializing command required.
04	03	Logical unit not ready, manual intervention required.

(continued on next page)

Table C–16 (Cont.) SCSI ASC/ASCQ Codes for Medium Changer Devices (Such as Jukeboxes)

ASC Code	ASCQ Code	Description
06	00	No reference position found.
07	00	Multiple peripheral devices selected.
08	00	Logical unit communication failure.
08	01	Logical unit communication time-out.
08	02	Logical unit communication parity error.
0A	00	Error log overflow.
15	00	Random positioning error.
15	01	Mechanical positioning error.
1A	00	Parameter list length error.
1B	00	Synchronous data transfer error.
20	00	Invalid command operation code.
21	00	Logical block address out of range.
21	01	Invalid element address.
24	00	Invalid field in cdb.
25	00	Logical unit not supported.
26	00	Invalid field in parameter list.
26	01	Parameter not supported.
26	02	Parameter value invalid.
26	03	Threshold parameters not supported.
28	00	Not ready to ready transition, medium may have changed.
28	01	Import or export element accessed.
29	00	Power on, reset, or bus device reset occurred.
29	01	Power on occurred.
29	02	SCSI bus reset occurred.
29	03	Bus device reset occurred.
2A	00	Parameters changed.
2A	01	Mode parameters changed.
2A	02	Log parameters changed.
2C	00	Command sequence error.
2F	00	Commands cleared by another initiator.
30	00	Incompatible medium installed.
37	00	Rounded parameter.
39	00	Saving parameters not supported.
3A	00	Medium not present.
3B	0D	Medium destination element full.
3B	0E	Medium source element empty.

(continued on next page)

Table C–16 (Cont.) SCSI ASC/ASCQ Codes for Medium Changer Devices (Such as Jukeboxes)

ASC Code	ASCQ Code	Description
3D	00	Invalid bits in identify message.
3E	00	Logical unit has not self-configured yet.
3F	00	Target operating conditions have changed.
3F	01	Microcode has been changed.
3F	02	Changed operating definition.
3F	03	Inquiry data has changed.
40	*nn*	Diagnostic failure detected on component *nn*; where *nn* identifies a specific target device component (*nn* range 80 through FF). Refer to documentation provided by the vendor of the target device for a description of the component identified by *nn*.
43	00	Message error.
44	00	Internal target failure.
45	00	Select or reselect failure.
46	00	Unsuccessful soft reset.
47	00	SCSI parity error.
48	00	Initiator detected error message received.
49	00	Invalid message error.
4A	00	Command phase error.
4B	00	Data phase error.
4C	00	Logical unit failed self-configuration.
4E	00	Overlapped commands attempted.
53	00	Media load or eject failed.
53	02	Medium removal prevented.
5A	00	Operator request or state change input (unspecified).
5A	01	Operator medium removal request.
5B	00	Log exception.
5B	01	Threshold condition met.
5B	02	Log counter at maximum.
5B	03	Log list codes exhausted.

Table C–17 HS1CP Controller Vendor Specific SCSI ASC/ASCQ Codes

ASC Code	ASCQ Code	Description
3F	85	Test Unit Ready or Read Capacity Command failed.
3F	87	Drive failed by a Host Mode Select command.
3F	88	Drive failed due to a deferred error reported by drive.
3F	90	Unrecovered Read/Write error.

(continued on next page)

Table C–17 (Cont.) HS1CP Controller Vendor Specific SCSI ASC/ASCQ Codes

ASC Code	ASCQ Code	Description
3F	C0	No response from one or more drives.
3F	C2	NV memory and drive metadata indicate conflicting drive configurations.
3F	D2	Synchronous Transfer Value differences between drives.
82	01	No command control structures available.
84	04	Command failed - SCSI ID verification failed.
85	05	Data returned from drive is invalid.
89	00	Request Sense command to drive failed.
8A	00	Illegal command for pass through mode.
8C	04	Data transfer request error.
8F	00	Premature completion of a drive command.
93	00	Drive returned vendor unique sense data.
A0	00	Last failure event report.
A0	01	Nonvolatile parameter memory component event report.
A0	02	Backup battery failure event report.
A0	03	Subsystem built-in self test failure event report.
A0	04	Memory system failure event report.
A0	05	Failover event report.
A0	06	Firmware license violation event report.
A0	07	RAID membership event report.
A1	00	Shelf OK is not properly asserted.
A1	01	Unable to clear SWAP interrupt, interrupt disabled.
A1	02	Swap interrupt re-enabled.
A1	03	Asynchronous SWAP detected.
A1	04	Controller shelf OK is not properly asserted.
B0	00	Command timeout.
B0	01	Watchdog timer timeout.
D0	01	Disconnect timeout.
D0	02	Chip command timeout.
D0	03	Byte transfer timeout.
D1	00	Bus errors.
D1	02	Unexpected bus phase.
D1	03	Disconnect expected.
D1	04	ID Message not sent.
D1	05	Synchronous negotiation error.
D1	07	Unexpected disconnect.
D1	08	Unexpected message.

(continued on next page)

Table C–17 (Cont.) HS1CP Controller Vendor Specific SCSI ASC/ASCQ Codes

ASC Code	ASCQ Code	Description
D1	09	Unexpected Tag message.
D1	0A	Channel busy.
D1	0B	Device initialization failure, device sense data available.
D2	00	Miscellaneous SCSI driver error.
D2	03	Device services had to reset the bus.
D3	00	Drive SCSI chip reported gross error.
D4	00	Non-SCSI bus parity error.
D5	02	Message Reject received on a valid message.
D7	00	Source driver programming error.
E0	03	Fault Manager detected an unknown error code.
E0	06	Maximum number of errors for this I/O exceeded.
E0	07	Drive reported recovered error without transferring all data.

Table C–18 Last Failure Event Log (Template 01) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01010302	03EA	An unrecoverable hardware detected fault occurred.
0102030A	040A	An unrecoverable firmware inconsistency was detected.

Table C–19 Failover Event Log (Template 05) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
07030B0A	022A	Failover Control detected a receive packet sequence number mismatch. The controllers are out of synchronization with each other and are unable to communicate. Note that in this instance, the "last failure code" and "last failure parameters" fields are undefined.
07040B0A	022A	Failover Control detected a transmit packet sequence number mismatch. The controllers are out of synchronization with each other and are unable to communicate. Note that in this instance, the "last failure code" and "last failure parameters" fields are undefined.
07050064	022A	Failover Control received a Last Gasp message from the other controller. The other controller is expected to restart itself within a given time period. If it does not, it will be held reset with the "Kill" line.

(continued on next page)

Table C–19 (Cont.) Failover Event Log (Template 05) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
07060C01	022A	Failover Control detected that both controllers are acting as SCSI ID 6. Because ids are determined by hardware, it is unknown which controller is the real SCSI ID 6. Note that in this instance, the "last failure code" and "last failure parameters" fields are undefined.
07070C01	022A	Failover Control detected that both controllers are acting as SCSI ID 7. Because ids are determined by hardware, it is unknown which controller is the real SCSI ID 7. Note that in this instance, the "last failure code" and "last failure parameters" fields are undefined.
07080B0A	022A	Failover Control was unable to send keep alive communication to the other controller. It is assumed that the other controller is hung or not started. Note that in this instance, the "last failure code" and "last failure parameters" fields are undefined.

Table C–20 Firmware License Violation Event Log (Template 06) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
0217000A	0016	RAID support is enabled but not licensed on this controller. Any use of this feature requires licensing. Continued use does not comply with the terms and conditions of licensing this product.
0218000A	0016	Write-Back Cache support is enabled but not licensed on this controller. Any use of this feature requires licensing. Continued use does not comply with the terms and conditions of licensing this product.
0227000A	0016	mirroring support is enabled but not licensed on this controller. Any use of this feature requires licensing. Continued use does not comply with the terms and conditions of licensing this product.

Table C–21 Nonvolatile Parameter Memory Component Event Log (Template 11) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01032002	012A	Nonvolatile parameter memory component EDC check failed; content of the component reset to default settings.

Table C–22 Backup Battery Failure Event Log (Template 12) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
02032001	012A	Journal SRAM backup battery failure; detected during system restart. The "memory address" field contains the starting physical address of the Journal SRAM.
02042001	012A	Journal SRAM backup battery failure; detected during periodic check. The "memory address" field contains the starting physical address of the Journal SRAM.
02052301	012A	A processor interrupt was generated by the CACHE Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging). The "memory address" field contains the starting physical address of the CACHEA0 memory.
02062301	012A	The CACHE backup battery has failed testing performed by the Cache Diagnostics during system startup. The "memory address" field contains the starting physical address of the CACHEA0 memory.

Table C–23 Subsystem Built-In Self Test Failure Event Log (Template 13) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
82012002	020A	An unrecoverable error was detected during execution of the Device Port Subsystem Built-In Self Test. One or more of the device ports on the controller module has failed; some/all of the attached storage is no longer accessible via this controller.
82042002	020A	A spurious interrupt was detected during the execution of a Subsystem Built-In Self Test.
82052002	020A	An unrecoverable error was detected during execution of the HOST PORT Subsystem Test. The system will not be able to communicate with the host.
82062002	020A	An unrecoverable error was detected during execution of the UART/DUART Subsystem Test. This will cause the console to be unusable. This will cause failover communications to fail.
82072002	020A	An unrecoverable error was detected during execution of the FX Subsystem Test.
82082002	020A	An unrecoverable error was detected during execution of the nbuss init Test.

Table C–24 Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
0121370A	012A	Memory System Error Analysis is indicated in the information preserved during a previous last failure but no error conditions are indicated in the available DRAB registers. The Master DRAB registers content is supplied.
0122330A	012A	An error condition detected by one of the CACHE DRABs that requires supplemental information has been reported in another event report. This event report contains the Master DRAB and Diagnostic registers content associated with that initial event report.
01292602	012A	The Master DRAB detected a Cache Timeout condition during an access attempt to a CACHEA0, CACHEA1, CACHEB0, or CACHEB1 DRAB registers region or memory region, the addressed region failed to respond.
012A3702	012A	The CACHEA0 DRAB unexpectedly reported a Cache Timeout condition.
012B3702	012A	The CACHEA1 DRAB unexpectedly reported a Cache Timeout condition.
012C3702	012A	The CACHEB0 DRAB unexpectedly reported a Cache Timeout condition.
012D3702	012A	The CACHEB1 DRAB unexpectedly reported a Cache Timeout condition.
012E2702	012A	The Master DRAB detected an Nbus Transfer Error Acknowledge (TEA) condition—that is, the TEA signal was asserted by an Nbus device during an I960 to Nbus device transaction.
012F3702	012A	The Master DRAB unexpectedly reported an Nbus Transfer Error Acknowledge condition.
01303702	012A	The CACHEA0 DRAB unexpectedly reported an Nbus Transfer Error Acknowledge condition.
01313702	012A	The CACHEA1 DRAB unexpectedly reported an Nbus Transfer Error Acknowledge condition.
01323702	012A	The CACHEB0 DRAB unexpectedly reported an Nbus Transfer Error Acknowledge condition.
01333702	012A	The CACHEB1 DRAB unexpectedly reported an Nbus Transfer Error Acknowledge condition.
01342802	012A	A Multiple Bit ECC error was detected during a memory refresh attempt by the Master DRAB.
01352902	012A	A Multiple Bit ECC error was detected during a memory refresh attempt by the CACHEA0 DRAB.
01362902	012A	A Multiple Bit ECC error was detected during a memory refresh attempt by the CACHEA1 DRAB.
01372A02	012A	A Multiple Bit ECC error was detected during a memory refresh attempt by the CACHEB0 DRAB.
01382A02	012A	A Multiple Bit ECC error was detected during a memory refresh attempt by the CACHEB1 DRAB.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01392802	012A	The Master DRAB detected a Multiple Bit ECC error during an FX attempt to read buffer memory.
013A2802	012A	The Master DRAB detected a Multiple Bit ECC error during a Host Port attempt to read buffer memory.
013B2802	012A	The Master DRAB detected a Multiple Bit ECC error during a Device Port attempt to read buffer memory.
013C2802	012A	The Master DRAB detected a Multiple Bit ECC error during an I960 attempt to read buffer memory.
013D2902	012A	The CACHEA0 DRAB detected a Multiple Bit ECC error during an FX attempt to read CACHEA0 memory.
013E2902	012A	The CACHEA0 DRAB detected a Multiple Bit ECC error during a Host Port attempt to read CACHEA0 memory.
013F2902	012A	The CACHEA0 DRAB detected a Multiple Bit ECC error during a Device Port attempt to read CACHEA0 memory.
01402902	012A	The CACHEA0 DRAB detected a Multiple Bit ECC error during an I960 attempt to read CACHEA0 memory.
01412902	012A	The CACHEA1 DRAB detected a Multiple Bit ECC error during an FX attempt to read CACHEA1 memory.
01422902	012A	The CACHEA1 DRAB detected a Multiple Bit ECC error during a Host Port attempt to read CACHEA1 memory.
01432902	012A	The CACHEA1 DRAB detected a Multiple Bit ECC error during a Device Port attempt to read CACHEA1 memory.
01442902	012A	The CACHEA1 DRAB detected a Multiple Bit ECC error during an I960 attempt to read CACHEA1 memory.
01452A02	012A	The CACHEB0 DRAB detected a Multiple Bit ECC error during an FX attempt to read CACHEB0 memory.
01462A02	012A	The CACHEB0 DRAB detected a Multiple Bit ECC error during a Host Port attempt to read CACHEB0 memory.
01472A02	012A	The CACHEB0 DRAB detected a Multiple Bit ECC error during a Device Port attempt to read CACHEB0 memory.
01482A02	012A	The CACHEB0 DRAB detected a Multiple Bit ECC error during an I960 attempt to read CACHEB0 memory.
01492A02	012A	The CACHEB1 DRAB detected a Multiple Bit ECC error during an FX attempt to read CACHEB1 memory.
014A2A02	012A	The CACHEB1 DRAB detected a Multiple Bit ECC error during a Host Port attempt to read CACHEB1 memory.
014B2A02	012A	The CACHEB1 DRAB detected a Multiple Bit ECC error during a Device Port attempt to read CACHEB1 memory.
014C2A02	012A	The CACHEB1 DRAB detected a Multiple Bit ECC error during an I960 attempt to read CACHEB1 memory.
014D3702	012A	The Master DRAB unexpectedly reported a Multiple Bit ECC error.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
014E3702	012A	The CACHEA0 DRAB unexpectedly reported a Multiple Bit ECC error.
014F3702	012A	The CACHEA1 DRAB unexpectedly reported a Multiple Bit ECC error.
01503702	012A	The CACHEB0 DRAB unexpectedly reported a Multiple Bit ECC error.
01513702	012A	The CACHEB1 DRAB unexpectedly reported a Multiple Bit ECC error.
01522B02	012A	The Master DRAB detected an Ibus to Nbus Timeout condition during an I960 to Nbus device transaction, the Nbus device failed to respond.
01533702	012A	The CACHEA0 DRAB unexpectedly reported an Ibus to Nbus Timeout condition.
01543702	012A	The CACHEA1 DRAB unexpectedly reported an Ibus to Nbus Timeout condition.
01553702	012A	The CACHEB0 DRAB unexpectedly reported an Ibus to Nbus Timeout condition.
01563702	012A	The CACHEB1 DRAB unexpectedly reported an Ibus to Nbus Timeout condition.
01572C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during an FX attempt to write buffer memory.
01582C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during an FX attempt to write a byte to buffer memory.
01592C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during an FX attempt to read buffer memory.
015A2C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write buffer memory.
015B2C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write a byte to buffer memory.
015C2C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to read buffer memory.
015D2C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write buffer memory.
015E2C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write a byte to buffer memory.
015F2C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to read buffer memory.
01602C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write buffer memory.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01612C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write a byte to buffer memory.
01622C02	012A	The Master DRAB detected a Nonexistent Memory Error condition during an I960 attempt to read buffer memory.
01632D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during an FX attempt to write CACHEA0 memory.
01642D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during an FX attempt to write a byte to CACHEA0 memory.
01652D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during an FX attempt to read CACHEA0 memory.
01662D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write CACHEA0 memory.
01672D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write a byte to CACHEA0 memory.
01682D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to read CACHEA0 memory.
01692D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write CACHEA0 memory.
016A2D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write a byte to CACHEA0 memory.
016B2D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to read CACHEA0 memory.
016C2D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write CACHEA0 memory.
016D2D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write a byte to CACHEA0 memory.
016E2D02	012A	The CACHEA0 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to read CACHEA0 memory.
016F2D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during an FX attempt to write CACHEA1 memory.
01702D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during an FX attempt to write a byte to CACHEA1 memory.
01712D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during an FX attempt to read CACHEA1 memory.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01722D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write CACHEA1 memory.
01732D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write a byte to CACHEA1 memory.
01742D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to read CACHEA1 memory.
01752D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write CACHEA1 memory.
01762D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write a byte to CACHEA1 memory.
01772D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to read CACHEA1 memory.
01782D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write CACHEA1 memory.
01792D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write a byte to CACHEA1 memory.
017A2D02	012A	The CACHEA1 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to read CACHEA1 memory.
017B2E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during an FX attempt to write CACHEB0 memory.
017C2E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during an FX attempt to write a byte to CACHEB0 memory.
017D2E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during an FX attempt to read CACHEB0 memory.
017E2E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write CACHEB0 memory.
017F2E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write a byte to CACHEB0 memory.
01802E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to read CACHEB0 memory.
01812E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write CACHEB0 memory.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01822E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write a byte to CACHEB0 memory.
01832E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to read CACHEB0 memory.
01842E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write CACHEB0 memory.
01852E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write a byte to CACHEB0 memory.
01862E02	012A	The CACHEB0 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to read CACHEB0 memory.
01872E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during an FX attempt to write CACHEB1 memory.
01882E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during an FX attempt to write a byte to CACHEB1 memory.
01892E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during an FX attempt to read CACHEB1 memory.
018A2E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write CACHEB1 memory.
018B2E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to write a byte to CACHEB1 memory.
018C2E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during a Host Port attempt to read CACHEB1 memory.
018D2E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write CACHEB1 memory.
018E2E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to write a byte to CACHEB1 memory.
018F2E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during a Device Port attempt to read CACHEB1 memory.
01902E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write CACHEB1 memory.
01912E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to write a byte to CACHEB1 memory.
01922E02	012A	The CACHEB1 DRAB detected a Nonexistent Memory Error condition during an I960 attempt to read CACHEB1 memory.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01933702	012A	The Master DRAB unexpectedly reported a Nonexistent Memory Error condition.
01943702	012A	The CACHEA0 DRAB unexpectedly reported a Nonexistent Memory Error condition.
01953702	012A	The CACHEA1 DRAB unexpectedly reported a Nonexistent Memory Error condition.
01963702	012A	The CACHEB0 DRAB unexpectedly reported a Nonexistent Memory Error condition.
01973702	012A	The CACHEB1 DRAB unexpectedly reported a Nonexistent Memory Error condition.
01982F02	012A	An Address Parity error was detected during a memory refresh attempt by the Master DRAB.
01993002	012A	An Address Parity error was detected during a memory refresh attempt by the CACHEA0 DRAB.
019A3002	012A	An Address Parity error was detected during a memory refresh attempt by the CACHEA1 DRAB.
019B3102	012A	An Address Parity error was detected during a memory refresh attempt by the CACHEB0 DRAB.
019C3102	012A	An Address Parity error was detected during a memory refresh attempt by the CACHEB1 DRAB.
019D2F02	012A	The Master DRAB detected an Address Parity error during an FX attempt to read buffer memory.
019E2F02	012A	The Master DRAB detected an Address Parity error during a Host Port attempt to read buffer memory.
019F2F02	012A	The Master DRAB detected an Address Parity error during a Device Port attempt to read buffer memory.
01A02F02	012A	The Master DRAB detected an Address Parity error during an I960 attempt to read buffer memory.
01A13002	012A	The CACHEA0 DRAB detected an Address Parity error during an FX attempt to read CACHEA0 memory.
01A23002	012A	The CACHEA0 DRAB detected an Address Parity error during a Host Port attempt to read CACHEA0 memory.
01A33002	012A	The CACHEA0 DRAB detected an Address Parity error during a Device Port attempt to read CACHEA0 memory.
01A43002	012A	The CACHEA0 DRAB detected an Address Parity error during an I960 attempt to read CACHEA0 memory.
01A53002	012A	The CACHEA1 DRAB detected an Address Parity error during an FX attempt to read CACHEA1 memory.
01A63002	012A	The CACHEA1 DRAB detected an Address Parity error during a Host Port attempt to read CACHEA1 memory.
01A73002	012A	The CACHEA1 DRAB detected an Address Parity error during a Device Port attempt to read CACHEA1 memory.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01A83002	012A	The CACHEA1 DRAB detected an Address Parity error during an I960 attempt to read CACHEA1 memory.
01A93102	012A	The CACHEB0 DRAB detected an Address Parity error during an FX attempt to read CACHEB0 memory.
01AA3102	012A	The CACHEB0 DRAB detected an Address Parity error during a Host Port attempt to read CACHEB0 memory.
01AB3102	012A	The CACHEB0 DRAB detected an Address Parity error during a Device Port attempt to read CACHEB0 memory.
01AC3102	012A	The CACHEB0 DRAB detected an Address Parity error during an I960 attempt to read CACHEB0 memory.
01AD3102	012A	The CACHEB1 DRAB detected an Address Parity error during an FX attempt to read CACHEB1 memory.
01AE3102	012A	The CACHEB1 DRAB detected an Address Parity error during a Host Port attempt to read CACHEB1 memory.
01AF3102	012A	The CACHEB1 DRAB detected an Address Parity error during a Device Port attempt to read CACHEB1 memory.
01B03102	012A	The CACHEB1 DRAB detected an Address Parity error during an I960 attempt to read CACHEB1 memory.
01B13702	012A	The Master DRAB unexpectedly reported an Address Parity error.
01B23702	012A	The CACHEA0 DRAB unexpectedly reported an Address Parity error.
01B33702	012A	The CACHEA1 DRAB unexpectedly reported an Address Parity error.
01B43702	012A	The CACHEB0 DRAB unexpectedly reported an Address Parity error.
01B53702	012A	The CACHEB1 DRAB unexpectedly reported an Address Parity error.
01B63202	012A	The Master DRAB detected an Ibus Parity Error during an I960 ID/Cache access attempt.
01B73202	012A	The Master DRAB detected an Ibus Parity Error during an I960 buffer memory access attempt.
01B83202	012A	The Master DRAB detected an Ibus Parity Error during an I960 buffer memory access attempt with a simultaneous but unrelated CACHExn memory access.
01B93202	012A	The Master DRAB detected an Ibus Parity Error during an I960 CACHEA memory access with a simultaneous but unrelated buffer memory access.
01BA3202	012A	The Master DRAB detected an Ibus Parity Error during an I960 CACHEB memory access with a simultaneous but unrelated buffer memory access.
01BB3202	012A	The Master DRAB detected an Ibus Parity Error during an I960 CACHEA Control and Status Register access with a simultaneous but unrelated buffer memory access.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01BC3202	012A	The Master DRAB detected an Ibus Parity Error during an I960 CACHEB Control and Status Register access with a simultaneous but unrelated buffer memory access.
01BD3702	012A	The Master DRAB unexpectedly reported an Ibus Parity error.
01BE3702	012A	The CACHEA0 DRAB unexpectedly reported an Ibus Parity error.
01BF3702	012A	The CACHEA1 DRAB unexpectedly reported an Ibus Parity error.
01C03702	012A	The CACHEB0 DRAB unexpectedly reported an Ibus Parity error.
01C13702	012A	The CACHEB1 DRAB unexpectedly reported an Ibus Parity error.
01C22F02	012A	The Master DRAB detected a Write Data Parity error during an FX attempt to write buffer memory.
01C32F02	012A	The Master DRAB detected a Write Data Parity error during an FX attempt to write a byte to buffer memory.
01C42F02	012A	The Master DRAB detected a Write Data Parity error during a Host Port attempt to write buffer memory.
01C52F02	012A	The Master DRAB detected a Write Data Parity error during a Host Port attempt to write a byte to buffer memory.
01C62F02	012A	The Master DRAB detected a Write Data Parity error during a Device Port attempt to write buffer memory.
01C72F02	012A	The Master DRAB detected a Write Data Parity error during a Device Port attempt to write a byte to buffer memory.
01C82F02	012A	The Master DRAB detected a Write Data Parity error during an I960 attempt to write buffer memory.
01C92F02	012A	The Master DRAB detected a Write Data Parity error during an I960 attempt to write a byte to buffer memory.
01CA3002	012A	The CACHEA0 DRAB detected a Write Data Parity error during an FX attempt to write CACHEA0 memory.
01CB3002	012A	The CACHEA0 DRAB detected a Write Data Parity error during an FX attempt to write a byte to CACHEA0 memory.
01CC3002	012A	The CACHEA0 DRAB detected a Write Data Parity error during a Host Port attempt to write CACHEA0 memory.
01CD3002	012A	The CACHEA0 DRAB detected a Write Data Parity error during a Host Port attempt to write a byte to CACHEA0 memory.
01CE3002	012A	The CACHEA0 DRAB detected a Write Data Parity error during a Device Port attempt to write CACHEA0 memory.
01CF3002	012A	The CACHEA0 DRAB detected a Write Data Parity error during a Device Port attempt to write a byte to CACHEA0 memory.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01D03002	012A	The CACHEA0 DRAB detected a Write Data Parity error during an I960 attempt to write CACHEA0 memory.
01D13002	012A	The CACHEA0 DRAB detected a Write Data Parity error during an I960 attempt to write a byte to CACHEA0 memory.
01D23002	012A	The CACHEA1 DRAB detected a Write Data Parity error during an FX attempt to write CACHEA1 memory.
01D33002	012A	The CACHEA1 DRAB detected a Write Data Parity error during an FX attempt to write a byte to CACHEA1 memory.
01D43002	012A	The CACHEA1 DRAB detected a Write Data Parity error during a Host Port attempt to write CACHEA1 memory.
01D53002	012A	The CACHEA1 DRAB detected a Write Data Parity error during a Host Port attempt to write a byte to CACHEA1 memory.
01D63002	012A	The CACHEA1 DRAB detected a Write Data Parity error during a Device Port attempt to write CACHEA1 memory.
01D73002	012A	The CACHEA1 DRAB detected a Write Data Parity error during a Device Port attempt to write a byte to CACHEA1 memory.
01D83002	012A	The CACHEA1 DRAB detected a Write Data Parity error during an I960 attempt to write CACHEA1 memory.
01D93002	012A	The CACHEA1 DRAB detected a Write Data Parity error during an I960 attempt to write a byte to CACHEA1 memory.
01DA3102	012A	The CACHEB0 DRAB detected a Write Data Parity error during an FX attempt to write CACHEB0 memory.
01DB3102	012A	The CACHEB0 DRAB detected a Write Data Parity error during an FX attempt to write a byte to CACHEB0 memory.
01DC3102	012A	The CACHEB0 DRAB detected a Write Data Parity error during a Host Port attempt to write CACHEB0 memory.
01DD3102	012A	The CACHEB0 DRAB detected a Write Data Parity error during a Host Port attempt to write a byte to CACHEB0 memory.
01DE3102	012A	The CACHEB0 DRAB detected a Write Data Parity error during a Device Port attempt to write CACHEB0 memory.
01DF3102	012A	The CACHEB0 DRAB detected a Write Data Parity error during a Device Port attempt to write a byte to CACHEB0 memory.
01E03102	012A	The CACHEB0 DRAB detected a Write Data Parity error during an I960 attempt to write CACHEB0 memory.
01E13102	012A	The CACHEB0 DRAB detected a Write Data Parity error during an I960 attempt to write a byte to CACHEB0 memory.
01E23102	012A	The CACHEB1 DRAB detected a Write Data Parity error during an FX attempt to write CACHEB1 memory.
01E33102	012A	The CACHEB1 DRAB detected a Write Data Parity error during an FX attempt to write a byte to CACHEB1 memory.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
01E43102	012A	The CACHEB1 DRAB detected a Write Data Parity error during a Host Port attempt to write CACHEB1 memory.
01E53102	012A	The CACHEB1 DRAB detected a Write Data Parity error during a Host Port attempt to write a byte to CACHEB1 memory.
01E63102	012A	The CACHEB1 DRAB detected a Write Data Parity error during a Device Port attempt to write CACHEB1 memory.
01E73102	012A	The CACHEB1 DRAB detected a Write Data Parity error during a Device Port attempt to write a byte to CACHEB1 memory.
01E83102	012A	The CACHEB1 DRAB detected a Write Data Parity error during an I960 attempt to write CACHEB1 memory.
01E93102	012A	The CACHEB1 DRAB detected a Write Data Parity error during an I960 attempt to write a byte to CACHEB1 memory.
01EA3702	012A	The Master DRAB unexpectedly reported a Write Data Parity error.
01EB3702	012A	The CACHEA0 DRAB unexpectedly reported a Write Data Parity error.
01EC3702	012A	The CACHEA1 DRAB unexpectedly reported a Write Data Parity error.
01ED3702	012A	The CACHEB0 DRAB unexpectedly reported a Write Data Parity error.
01EE3702	012A	The CACHEB1 DRAB unexpectedly reported a Write Data Parity error.
02072201	012A	The CACHE Dynamic Ram controller and Arbitration engine 0 (DRAB0) failed testing performed by the Cache Diagnostics. The "memory address" field contains the starting physical address of the CACHEA0 memory.
02082201	012A	The CACHE Dynamic Ram controller and Arbitration engine 1 (DRAB1) failed testing performed by the Cache Diagnostics. The "memory address" field contains the starting physical address of the CACHEA1 memory.
020B2201	012A	Failed read test of a write-back metadata page residing in cache. Dirty write-back cached data exists and cannot be flushed to media. The dirty data is lost. The "memory address" field contains the starting physical address of the CACHEA0 memory.
020C2201	012A	Cache Diagnostics have declared the cache bad during testing. The "memory address" field contains the starting physical address of the CACHEA0 memory.
020D2401	012A	The wrong write cache module is configured. The serial numbers do not match. Either the existing or the expected cache contains dirty write-back cached data. Note that in this instance the "memory address," "byte count," DRAB register, and Diagnostic register fields are undefined.

(continued on next page)

Table C–24 (Cont.) Memory System Failure Event Log (Template 14) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
020E2401	012A	The write cache module is missing. A cache is expected to be configured and contains dirty write-back cached data. Note that in this instance the "memory address," "byte count," DRAB register, and Diagnostic register fields are undefined.
020F2401	012A	The write cache modules are not configured properly for a dual-redundant configuration. One of the write cache modules is not present to perform cache failover of dirty write-back cached data. Note that in this instance, the "memory address," "byte count," DRAB register, and Diagnostic register fields are undefined.
02102401	012A	The write cache modules are not configured properly for a dual-redundant configuration. One of the cache modules is not the same size to perform cache failover of dirty write-back cached data. Note that in this instance, the "memory address," "byte count," DRAB register, and Diagnostic register fields are undefined.
02192401	012A	The cache modules are not configured properly for a dual-redundant configuration. One of the cache modules is not the same type (that is, both write cache) to perform cache failover of dirty write-back cached data. Note that in this instance the "memory address," "byte count," DRAB register, and Diagnostic register fields are undefined.
021D0064	012A	Unable to lock the other controller's cache in a write-cache failover attempt. Either a latent error could not be cleared on the cache or the other controller did not release its cache. Note that in this instance, the "memory address," "byte count," DRAB register, and Diagnostic register fields are undefined.

Table C–25 Internal Bus Port/Port Driver Event Log (Template 32) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
4003640A	006A	Internal Bus Port detected a Dual Receive condition that resulted in the closure of the Virtual Circuit. This error condition will be eliminated in a future CI interface chip.
4004020A	006A	Host Interconnect Services detected protocol error upon validating a received packet.
4007640A	006A	Internal Bus Port detected error upon attempting to transmit a packet. This resulted in the closure of the Virtual Circuit.
400F640A	006A	Host Interconnect Services detected packet sequence number mismatch.
403D020A	006A	Received packet with an unrecognized PPD opcode. Note that the content of the "vcstate" field is undefined in this instance.
40440064	006A	Received a PPD NODE_STOP and closed virtual circuit.

Table C–26 Internal Bus System Communication Services Event Log (Template 33) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
4015020A	006A	Remote SYSAP sent an SCS APPL_MSG but no receive credit was available.
4029010A	006A	Illegal connection state. Not in CONNECT_REC connection state when an SCS ACCEPT_REQ is pending.
402A010A	006A	Illegal connection state. Not in CONNECT_REC connection state when an SCS REJECT_REQ is pending.
402B010A	006A	Illegal connection state. Not in CLOSED connection state when an SCS CONNECT_REQ is pending.
402C010A	006A	Illegal connection state. Not in OPEN or DISCONNECT_REC connection state when an SCS DISCONNECT_REQ is pending.
4051020A	006A	Received SCS CONNECT_RSP when not in CONNECT_SENT connection state.
4052020A	006A	Received SCS CONNECT_RSP when the connection is no longer valid.
4053020A	006A	Received SCS ACCEPT_REQ when not in CONNECT_ACK connection state.
4054020A	006A	Received SCS ACCEPT_RSP when not in the ACCEPT_SENT connection state.
4055020A	006A	Received SCS REJECT_REQ when not in the CONNECT_ACK connection state.
4056020A	006A	Received SCS REJECT_RSP when not in the REJECT_SENT connection state.
4057020A	006A	Received SCS DISCONNECT_REQ when not in the OPEN, DISCONNECT_SENT or DISCONNECT_ACK connection state.
4058020A	006A	Received SCS DISCONNECT_RSP when not in the DISCONNECT_SENT or DISCONNECT_MATCH connection state.
4059020A	006A	Received SCS CREDIT_REQ when in the DISCONNECT_REC or DISCONNECT_MATCH connection state.
405A020A	006A	Received SCS APPL_MSG when in the DISCONNECT_SENT or DISCONNECT_ACK connection state.
405B020A	006A	Received SCS ACCEPT_REQ on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
405C020A	006A	Received SCS ACCEPT_RSP on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
405D020A	006A	Received SCS REJECT_REQ on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.

(continued on next page)

Table C–26 (Cont.) Internal Bus System Communication Services Event Log (Template 33) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
405E020A	006A	Received SCS REJECT_RSP on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
405F020A	006A	Received SCS DISCONNECT_REQ on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
4060020A	006A	Received SCS DISCONNECT_RSP on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
4061020A	006A	Received SCS CREDIT_REQ on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
4062020A	006A	Received SCS CREDIT_RSP on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
4063020A	006A	Received SCS APPL_MSG on a connection that is no longer valid. Note that in this instance, if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
4064020A	006A	Received an unrecognized SCS message. Note that in this instance if the "connection id" field is zero, the content of the "vcstate," "remote node name," "remote connection id" and "connection state" fields are undefined.
4065020A	006A	Received SCS CONNECT_RSP with an unrecognized status. Connection is broken by Host Interconnect Services.
4066020A	006A	Received SCS REJECT_REQ with an invalid reason.
4067020A	006A	Received SCS APPL_MSG with no receive credit available.

Table C–27 Device Services Nontransfer Error Event Log (Template 41) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
021B0064	0014	Disk Bad Block Replacement attempt completed for a read of controller metadata from a location outside the user data area of the disk. Note that due to the way Bad Block Replacement is performed on SCSI disk drives, information on the actual replacement blocks is not available to the controller and is therefore not included in the event report.
021A0064	0014	Disk Bad Block Replacement attempt completed for a write of controller metadata to a location outside the user data area of the disk. Note that due to the way Bad Block Replacement is performed on SCSI disk drives, information on the actual replacement blocks is not available to the controller and is therefore not included in the event report.
03010101	006A	No command control structures available for disk operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03820101	006A	No command control structures available for tape operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03B40101	006A	No command control structures available for media loader operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03C80101	006A	No command control structures available for operation to a device that is unkown to the controller. Note that in this instance the "asc" and "ascq" fields are undefined.
03022002	002A	SCSI interface chip command timeout during disk operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03832002	002A	SCSI interface chip command timeout during tape operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03B52002	002A	SCSI interface chip command timeout during media loader operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03C92002	002A	SCSI interface chip command timeout during operation to a device that is unknown to the controller. Note that in this instance the "asc" and "ascq" fields are undefined.
03034002	016A	Byte transfer timeout during disk operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03844002	016A	Byte transfer timeout during tape operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03B64002	016A	Byte transfer timeout during media loader operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03CA4002	016A	Byte transfer timeout during operation to a device that is unknown to the controller. Note that in this instance, the "asc" and "ascq" fields are undefined.
03044402	01AA	SCSI bus errors during disk operation. Note that in this instance the "asc" and "ascq" fields are undefined.

(continued on next page)

Table C–27 (Cont.) Device Services Nontransfer Error Event Log (Template 41) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
03854402	01AA	SCSI bus errors during tape operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03B74402	01AA	SCSI bus errors during media loader operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03D24402	01AA	SCSI bus errors during device operation. The device type is unknown to the controller. Note that in this instance, the "asc" and "ascq" fields are undefined.
03052002	002A	Device port SCSI chip reported gross error during disk operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03862002	002A	Device port SCSI chip reported gross error during tape operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03B82002	002A	Device port SCSI chip reported gross error during media loader operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03CD2002	002A	Device port SCSI chip reported gross error during operation to a device that is unknown to the controller. Note that in this instance the "asc" and "ascq" fields are undefined.
03062002	008A	Non-SCSI bus parity error during disk operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03872002	008A	Non-SCSI bus parity error during tape operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03B920(2	008A	Non-SCSI bus parity error during media loader operation. Note that in this instance the "asc" and "ascq" fields are undefined.
03CE2002	008A	Non-SCSI bus parity error during operation to a device that is unknown to the controller. Note that in this instance, the "asc" and "ascq" fields are undefined.
03070101	01CA	Source driver programming error encountered during disk operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03880101	01CA	Source driver programming error encountered during tape operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03BA0101	01CA	Source driver programming error encountered during media loader operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03CF0101	01CA	Source driver programming error encountered during operation to a device that is unknown to the controller. Note that in this instance the "asc" and "ascq" fields are undefined.
03080101	01EA	Miscellaneous SCSI Port Driver coding error detected during disk operation. Note that in this instance, the "asc" and "ascq" fields are undefined.

(continued on next page)

Table C–27 (Cont.) Device Services Nontransfer Error Event Log (Template 41) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
03890101	01EA	Miscellaneous SCSI Port Driver coding error encountered during tape operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03BB0101	01EA	Miscellaneous SCSI Port Driver coding error detected during media loader operation. Note that in this instance, the "asc" and "ascq" fields are undefined.
03CB0101	01EA	Miscellaneous SCSI Port Driver coding error detected during operation to a device that is unkown to the controller. Note that in this instance the "asc" and "ascq" fields are undefined.
03270101	01EA	A disk related error code was reported that was unknown to the Fault Management firmware. Note that in this instance, the "asc" and "ascq" fields are undefined.
038A0101	01EA	A tape related error code was reported that was unknown to the Fault Management firmware. Note that in this instance, the "asc" and "ascq" fields are undefined.
03BC0101	01EA	A media loader related error code was reported that was unknown to the Fault Management firmware. Note that in this instance the "asc" and "ascq" fields are undefined.
03CC0101	01EA	A error code was reported that was unknown to the Fault Management firmware. Note that in this instance, the "asc" and "ascq" fields are undefined.
03D04002	01AA	A failure occurred while attempting a SCSI Test Unit Ready or Read Capacity command to a device. The device type is unknown to the controller. Note that in this instance, the "asc" and "ascq" fields are undefined.
03D14002	006A	The identification of a device does not match the configuration information. The actual device type is unknown to the controller. Note that in this instance, the "asc" and "ascq" fields are undefined.
03F40064	016A	Device services had to reset the port to clear a bad condition. Note that in this instance, the "target," "asc," and "ascq" fields are undefined.
03F00402	00EB	The shelf indicated by the "port" field is reporting a problem. This could mean one or both of the following: • If the shelf is using dual power supplies, one power supply has failed. • One of the shelf cooling fans has failed. Note that in this instance the "target," "asc," and "ascq" fields are undefined.

(continued on next page)

Table C–27 (Cont.) Device Services Nontransfer Error Event Log (Template 41) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
03F10502	00EB	The SWAP interrupt from the shelf indicated by the "port" field can not be cleared. All SWAP interrupts from all ports will be disabled until corrective action is taken. When SWAP interrupts are disabled, both controller front panel button presses and removal/insertion of devices are not detected by the controller. Note that in this instance, the "target," "asc," and "ascq" fields are undefined.
03F20064	00EB	The SWAP interrupts have been cleared and re-enabled for all shelves. Note that in this instance, the "port," "target," "asc," and "ascq" fields are undefined.
03F30064	00EB	An asynchronous SWAP interrupt was detected by the controller for the shelf indicated by the "port" field. Possible reasons for this occurance include: • Device insertion/removal • Shelf power failure • SWAP interrupts re-enabled. Note that in this instance, the "target," "asc," and "ascq" fields are undefined.
03D3450A	00EB	During device initialization, the device reported the SCSI Sense Key NO SENSE. This indicates that there is no specific sense key information to be reported for the designated logical unit. This would be the case for a successful command or a command that received CHECK CONDITION or COMMAND TERMINATED status because one of the FM, EOM, or ILI bits is set to one in the sense data flags field.
03D4450A	00EB	During device initialization, the device reported the SCSI Sense Key RECOVERED ERROR. This indicates the last command completed successfully with some recovery action performed by the target.
03D5450A	00EB	During device initialization, the device reported the SCSI Sense Key NOT READY. This indicates that the logical unit addressed cannot be accessed. Operator intervention may be required to correct this condition.
03D6450A	00EB	During device initialization, the device reported the SCSI Sense Key MEDIUM ERROR. This indicates that the command terminated with a nonrecovered error condition that was probably caused by a flaw in the medium or an error in the recorded data. This sense key may also be returned if the target is unable to distinguish between a flaw in the medium and a specific hardware failure (HARDWARE ERROR sense key).
03D7450A	00EB	During device initialization, the device reported the SCSI Sense Key HARDWARE ERROR. This indicates that the target detected a nonrecoverable hardware failure (for example, controller failure, device failure, parity error, and so forth) while performing the command or during a self test.

(continued on next page)

Table C–27 (Cont.) Device Services Nontransfer Error Event Log (Template 41) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
03D8450A	00EB	During device initialization, the device reported the SCSI Sense Key ILLEGAL REQUEST. Indicates that there was an illegal parameter in the command descriptor block or in the additional parameters supplied as data for some commands (FORMAT UNIT, SEARCH DATA, and so forth). If the target detects an invalid parameter in the command descriptor block, then it shall terminate the command without altering the medium. If the target detects an invalid parameter in the additional parameters supplied as data, then the target may have already altered the medium. This sense key may also indicate that an invalid IDENTIFY message was received.
03D9450A	00EB	During device initialization, the device reported the SCSI Sense Key UNIT ATTENTION. This indicates that the removable medium may have been changed or the target has been reset.
03DA450A	00EB	During device initialization, the device reported the SCSI Sense Key DATA PROTECT. This indicates that a command that reads or writes the medium was attempted on a block that is protected from this operation. The read or write operation is not performed.
03DB450A	00EB	During device initialization, the device reported the SCSI Sense Key BLANK CHECK. This indicates that a write-once device or a sequential-access device encountered blank medium or format-defined end-of-data indication while reading or a write-once device encountered a medium that was not blank while writing.
03DC450A	00EB	During device initialization, the device reported a SCSI Vendor Specific Sense Key. This sense key is available for reporting vendor specific conditions.
03DD450A	00EB	During device initialization, the device reported the SCSI Sense Key COPY ABORTED. This indicates a COPY, COMPARE, or COPY AND VERIFY command was aborted due to an error condition on the source device, the destination device, or both.
03DE450A	00EB	During device initialization, the device reported the SCSI Sense Key ABORTED COMMAND. This indicates the target aborted the command. The initiator may be able to recover by trying the command again.
03DF450A	00EB	During device initialization, the device reported the SCSI Sense Key EQUAL. This indicates a SEARCH DATA command has satisfied an equal comparison.
03E0450A	00EB	During device initialization, the device reported the SCSI Sense Key VOLUME OVERFLOW. This indicates a buffered peripheral device has reached the end-of-partition and data may remain in the buffer that has not been written to the medium. A RECOVER BUFFERED DATA commands may be issued to read the unwritten data from the buffer.
03E1450A	00EB	During device initialization, the device reported the SCSI Sense Key MISCOMPARE. This indicates the source data did not match the data read from the medium.

(continued on next page)

Table C–27 (Cont.) Device Services Nontransfer Error Event Log (Template 41) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
03E2450A	00EB	During device initialization, the device reported a reserved SCSI Sense Key.
03F50402	016A	The controller shelf is reporting a problem. This could mean one or both of the following: • If the shelf is using dual power supplies, one power supply has failed. • One of the shelf cooling fans has failed. Note that in this instance, the "target," "asc," and "ascq" fields are undefined.

Table C–28 Disk Transfer Error Event Log (Template 51) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
02090064	0007	A data compare error was detected during the execution of a compare modified READ or WRITE command. Note that in this instance the SCSI Device Sense Data fields, "cmdopcd" through "keyspec," are undefined.
021E0064	0016	The device specified in the "device locator" field has been added to the RAIDset unit specified in the "unit identifier" field. The RAIDset is now in Reconstructing state. Note that in this instance the "header code," "ancillary information," and "cmdopcd" fields are undefined.
021F0064	0016	The device specified in the "device locator" field has been removed from the RAIDset unit specified in the "unit identifier" field. The removed device is now in the failedset. The RAIDset is now in Reduced state. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
02200064	0016	The device specified in the "device locator" field failed to be added to the RAIDset unit specified in the "unit identifier" field. The device will remain in the spareset. Note that in this instance the "header code," "ancillary information," and "cmdopcd" fields are undefined.
02210064	0016	The device specified in the "device locator" field failed to be added to the RAIDset unit specified in the "unit identifier" field. The failed device has been moved to the failedset. Note that in this instance the "header code," "ancillary information," and "cmdopcd" fields are undefined.

(continued on next page)

Table C–28 (Cont.) Disk Transfer Error Event Log (Template 51) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
02220064	0016	The RAIDset unit specified in the "unit identifier" field has transitioned from Reconstructing state to Normal state. Note that in this instance the "header code," "ancillary information," and "cmdopcd" fields are undefined. Note also that information supplied in the "device locator," "devtype," "device identification," and "device serial number" fields is for the first device in the RAID set.
02280064	0016	The device specified in the "device locator" field has been added to the mirrorset unit specified in the "unit identifier" field. The new mirrorset member is now in Copying state. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
02230064	0016	The RAIDset unit specified in the "unit identifier" field has gone inoperative. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined. Note also that information supplied in the "device locator," "devtype," "device identification," and "device serial number" fields is for the first device in the RAID set.
02240064	0016	The RAIDset unit specified in the "unit identifier" field has transitioned from Normal state to Reconstructing state. Note that in this instance the "header code," "ancillary information," and "cmdopcd" fields are undefined. Note also that information supplied in the "device locator," "devtype," "device identification," and "device serial number" fields is for the first device in the RAID set.
02280064	0016	The device specified in the "device locator" field has added the mirrorset unit specified in the "unit identifier" field. The new mirrorset member is now in Copying state. Note that in thie instance the "he r code," "ancillary information," and "cmdopcd" fields are u. .efined.
02290064	0016	The device specified in the "device locator" field has been removed from the mirrorset unit specified in the "unit identifier" field. The removed device is now in the failedset. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
022A0064	0016	The device specified in the "device locator" field failed to be added to the mirrorset unit specified in the "unit identifier" field. The device will remain in the spareset. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
022B0064	0016	The device specified in the "device locator" field failed to be added to the mirrorset unit specified in the "unit identifier" field. The failed device has been moved to the failedset. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
022C0064	0016	The device specified in the "device locator" has transitioned from Copying or Normalizing state to Normal state. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.

(continued on next page)

Table C–28 (Cont.) Disk Transfer Error Event Log (Template 51) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
022D0064	0016	The mirrorset unit specified in the "unit identifier" field has gone inoperative. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined. Note also that information supplied in the "device locator," "devtype," "device identification," and "device serial number" fields is for the first device in the mirrorset.
022E0064	0016	The device specified in the "device locator" field has been converted to a mirrorset unit. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
022F0064	0016	The mirrored device specified in the "device locator" field has been converted to a single device unit Note that in this instance the "header code," "ancillary information," and "cmdopcd" fields are undefined.
02300064	0016	The device specified in the "device locator" field has been reduced from the mirrorset unit specified in the "unit identifier" field. The nominal number of members in the mirrorset has been decreased by one. The reduced device is now available for use. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
02310064	0016	The mirrorset unit specified in the "unit identifier" field has had its nominal membership changed. The new nominal number of members for the mirrorset is specified in the "port" field. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined. Note also that information supplied in the "device locator," "devtype," "device identification," and "device serial number" fields is for the first device in the mirrorset.
023B0064	0016	The device specified in the "device locator" field had a read error which has been repaired with data from another mirrorset member. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
023C0064	0016	The device specified in the "device locator" field had a read error. Attempts to repair the error with data from another mirrorset member failed due to lack of alternate error-free data source. Note that in this instance the "header code," "ancillary information," and "cmdopcd" fields are undefined.
023D0064	0016	The device specified in the "device locator" field had a read error. Attempts to repair the error with data from another mirrorset member failed due to a write error on the original device. The original device will be removed from the mirrorset. Note that in this instance, the "header code," "ancillary information," and "cmdopcd" fields are undefined.
03094002	000B	An unrecoverable disk drive error was encountered while performing work related to disk unit operations.
0328450A	000B	The disk device reported standard SCSI Sense Data.
030C4002	014B	A Drive failed because a Test Unit Ready command or a Read Capacity command failed.

(continued on next page)

Table C–28 (Cont.) Disk Transfer Error Event Log (Template 51) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
030D000A	0103	Drive was failed by a Mode Select command received from the host.
030E4002	00EB	Drive failed due to a deferred error reported by drive.
030F4002	00E8	Unrecovered Read or Write error.
03104002	002B	No response from one or more drives.
0311430A	012B	Nonvolatile memory and drive metadata indicate conflicting drive configurations.
0312430A	012B	The Synchronous Transfer Value differs between drives in the same storageset.
03134002	012B	Maximum number of errors for this data transfer operation exceeded.
03144002	00CB	Drive reported recovered error without transferring all data.
03154002	00E8	Data returned from drive is invalid.
03164002	012B	Request Sense command to drive failed.
03170064	0016	Illegal command for pass through mode.
03180064	0016	Data transfer request error.
03194002	012B	Premature completion of a drive command.
031A4002	002B	Command timeout.
031B0101	002B	Watchdog timer timeout.
031C4002	002B	Disconnect timeout.
031D4002	012B	Unexpected bus phase.
031E4002	012B	Disconnect expected.
031F4002	012B	ID Message not sent by drive.
03204002	012B	Synchronous negotiation error.
03214002	012B	The drive unexpectedly disconnected from the SCSI bus.
03224002	012B	Unexpected message.
03234002	012B	Unexpected Tag message.
03244002	012B	Channel busy.
03254002	012B	Message Reject received on a valid message.
0326450A	00EB	The disk device reported Vendor Unique SCSI Sense Data.

Table C–29 Disk Bad Block Replacement Attempt Event Log (Template 57) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
02110064	0014	Disk Bad Block Replacement attempt completed for a read within the user data area of the disk. Note that due to the way Bad Block Replacement is performed on SCSI disk drives, information on the actual replacement blocks is not available to the controller and is therefore not included in the "Old RBN" and "New RBN" fields. The content of those fields is undefined.
02020064	0014	Disk Bad Block Replacement attempt completed for a write within the user data area of the disk. Note that due to the way Bad Block Replacement is performed on SCSI disk drives, information on the actual replacement blocks is not available to the controller and is therefore not included in the "Old RBN" and "New RBN" fields. The content of those fields is undefined.

Table C–30 Tape Transfer Error Event Log (Template 61) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
020A0064	0007	A data compare error was detected during the execution of a compare modified READ or WRITE command. Note that in this instance the SCSI Device Sense Data fields, "cmdopcd" through "keyspec," are undefined.
03644002	000B	An unrecoverable tape drive error was encountered while performing work related to tape unit operations.
038B450A	000B	The tape device reported standard SCSI Sense Data.
03674002	014B	A Drive failed because a Test Unit Ready command or a Read Capacity command failed.
0368000A	0103	Drive was failed by a Mode Select command received from the host.
03694002	00EB	Drive failed due to a deferred error reported by drive.
036A4002	00E8	Unrecovered Read or Write error.
036B4002	002B	No response from one or more drives.
036C430A	012B	Nonvolatile memory and drive metadata indicate conflicting drive configurations.
036D430A	012B	The Synchronous Transfer Value differs between drives in the same storageset.
036E4002	012B	Maximum number of errors for this data transfer operation exceeded.
036F4002	00CB	Drive reported recovered error without transferring all data.
03704002	00E8	Data returned from drive is invalid.
03714002	012B	Request Sense command to drive failed.
03720064	0016	Illegal command for pass through mode.

(continued on next page)

Table C–30 (Cont.) Tape Transfer Error Event Log (Template 61) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
03730064	0016	Data transfer request error.
03744002	012B	Premature completion of a drive command.
03754002	002B	Command timeout.
03760101	002B	Watchdog timer timeout.
03774002	002B	Disconnect timeout.
03784002	012B	Unexpected bus phase.
03794002	012B	Disconnect expected.
037A4002	012B	ID Message not sent by drive.
037B4002	012B	Synchronous negotiation error.
037C4002	012B	The drive unexpectedly disconnected from the SCSI bus.
037D4002	012B	Unexpected message.
037E4002	012B	Unexpected Tag message.
037F4002	012B	Channel busy.
03804002	012B	Message Reject received on a valid message.
0381450A	00EB	The tape device reported Vendor Unique SCSI Sense Data.

Table C–31 Media Loader Error Event Log (Template 71) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
03964002	0097	An unrecoverable media loader error was encountered while performing work related to media loader operations.
03BD450A	0097	The media changer device reported standard SCSI Sense Data.
03994002	0097	A Drive failed because a Test Unit Ready command or a Read Capacity command failed.
039A000A	0077	Drive was failed by a Mode Select command received from the host.
039B4002	0097	Drive failed due to a deferred error reported by drive.
039C4002	0097	Unrecovered Read or Write error.
039D4002	0037	No response from one or more drives.
039E430A	0097	Nonvolatile memory and drive metadata indicate conflicting drive configurations.
039F430A	0097	The Synchronous Transfer Value differs between drives in the same storageset.
03A04002	0097	Maximum number of errors for this data transfer operat on exceeded.
03A14002	0097	Drive reported recovered error without transferring all data.

(continued on next page)

Table C–31 (Cont.) Media Loader Erro Event Log (Template 71) Instance/MSCP Event Codes

Instance Code	MSCP Event Code	Description
03A24002	0097	Data returned from drive is invalid.
03A34002	0097	Request Sense command to drive failed.
03A40064	0016	Illegal command for pass through mode.
03A50064	0016	Data transfer request error.
03A64002	0097	Premature completion of a drive command.
03A74002	0037	Command timeout.
03A80101	0037	Watchdog timer timeout.
03A94002	0037	Disconnect timeout.
03AA4002	0097	Unexpected bus phase.
03AB4002	0097	Disconnect expected.
03AC4002	0097	ID Message not sent by drive.
03AD4002	0097	Synchronous negotiation error.
03AE4002	0097	The drive unexpectedly disconnected from the SCSI bus.
03AF4002	0097	Unexpected message.
03B04002	0097	Unexpected Tag message.
03B14002	0097	Channel busy.
03B24002	0097	Message Reject received on a valid message.
03B3450A	0097	The media changer device reported Vendor Unique SCSI Sense Data.

Table C–32 Disk Copy Data Correlation Event Log "Event Dependent Information" Values

Value	Description
00000001	Unable to allocate a sufficient number of DCD Context Blocks to support this host.
00000002	Unable to find an inactive Unit Path Block.
00000003	Unable to find an inactive Source Unit Block.
00000004	Insufficient resources returned by HIS$CONNECT.

Table C–33 Executive Services Last Failure Codes

Code	Description
01000100	Memory allocation failure during executive initialization.
01010100	An interrupt without any handler was triggered.
01020100	Entry on timer que was not of type AQ or BQ.
01030100	Memory allocation for a facility lock failed.

(continued on next page)

Table C–33 (Cont.) Executive Services Last Failure Codes

Code	Description
01040100	Memory initialization called with invalid memory type.
01050104	The I960 reported a fault. • Last Failure Parameter[0] contains the PC value. • Last Failure Parameter[1] contains the AC value. • Last Failure Parameter[2] contains the fault type and subtype values. • Last Failure Parameter[3] contains the address of the faulting instruction.
01070100	Timer chip setup failed.
01082004	The core diagnostics reported a fault. • Last Failure Parameter[0] contains the error code value (same as blinking OCP LEDs error code). • Last Failure Parameter[1] contains the address of the fault. • Last Failure Parameter[2] contains the actual data value. • Last Failure Parameter[3] contains the expected data value.
01090105	An NMI occurred during EXEC$BUGCHECK processing. • Last Failure Parameter[0] contains the executive flags value. • Last Failure Parameter[1] contains the RIP from the NMI stack. • Last Failure Parameter[2] contains the read diagnostic register 0 value. • Last Failure Parameter[3] contains the Master DRAB CSR value. • Last Failure Parameter[4] contains the SIP last failure code value.
010A2080	A single bit or a multi-single bit ECC error was detected. Due to a hardware bug, this could result in data corruption. If this event occurs frequently, the controller module should be replaced.
010B2380	A processor interrupt was generated by the CACHEA Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).
010C2380	A processor interrupt was generated by the CACHEB Dynamic Ram controller and Arbitration engine (DRAB) with an indication that the CACHE backup battery has failed or is low (needs charging).
010D0110	The System Information structure within the System Information Page has been reset to default settings. The only known cause for this event is an I960 processor hang caused by an unimplemented memory region reference. When such a hang occurs, controller modules equipped with inactivity watchdog timer circuitry will spontaneously reboot after the watchdog timer expires (within seconds of the hang). Controller modules not so equipped will just hang as indicated by the green LED on the OCP remaining in a steady state.

(continued on next page)

Table C–33 (Cont.) Executive Services Last Failure Codes

Code	Description
010E0110	All structures contained in the System Information Page (SIP) and the Last Failure entries have been reset to their default settings. This is a normal occurance for the first boot following manufacture of the controller module and during the transition from one firmware version to another if and only if the format of the SIP is different between the two versions. If this event is reported at any other time, follow the recommended repair action associated with this Last Failure code.
010F0110	All structures contained in the System Information Page and the Last Failure entries have been reset to their default settings as the result of certain controller manufacturing configuration activities. If this event is reported at any other time, follow the recommended repair action associated with this Last Failure code.
018000A0	A powerfail interrupt occured.
018600A0	A processor interrupt was generated with an indication that the other controller in a dual controller configuration asserted the KILL line to disable this controller.
018700A0	A processor interrupt was generated with an indication that the (//) RESET button on the controller module was depressed.
018800A0	A processor interrupt was generated with an indication that the program card was removed.
018900A0	A processor interrupt was generated with an indication that the controller inactivity watch dog timer expired.
018B2580	A NMI interrupt was generated with an indication that a memory system problem occurred.
018C2580	A DRAB_INT interrupt was generated with an indication that a memory system problem occurred.

Table C–34 Value-Added Services Last Failure Codes

Code	Description
02000100	Initialization code was unable to allocate enough memory to set up the receive data descriptors.
02010100	Initialization code was unable to allocate enough memory to set up the send data descriptors.
02040100	Unable to allocate memory necessary for data buffers.
02050100	Unable to allocate memory for the Free Buffer Array.
02080100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when populating the disk read DWD stack.
02090100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when populating the disk write DWD stack.
020A0100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when populating the tape read DWD stack.
020B0100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when populating the tape write DWD stack.
020C0100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when populating the miscellaneous DWD stack.

(continued on next page)

Table C–34 (Cont.) Value-Added Services Last Failure Codes

Code	Description
02100100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when creating the device services state table.
02170100	Unable to allocate memory for the Free Node Array.
02180100	Unable to allocate memory for the Free Buffer Descriptor Array.
021D0100	Unable to allocate memory for the Free Buffer Array.
021E0100	Unable to allocate memory for the Free Strip Node Array.
021F0100	Unable to allocate memory for WARPs and RMDs.
02210100	Invalid parameters in CACHE$OFFER_META call.
02220100	No buffer found for CACHE$MARK_META_DIRTY call.
02270104	A callback from DS on a transfer request has returned a bad or illegal DWD status. • Last Failure Parameter[0] contains the DWD Status. • Last Failure Parameter[1] contains the DWD address. • Last Failure Parameter[2] contains the PUB Address. • Last Failure Parameter[3] contains the Device Port.
022E0102	An invalid mapping type was specified for a logical unit. • Last Failure Parameter[0] contains the USB address. • Last Failure Parameter[1] contains the Unit Mapping Type.
02360101	Unrecognized state supplied to FOC$SEND callback routine va_dap_snd_cmd_complete. Last Failure Parameter[0] contains the unrecognized value.
02370102	Unsupported return from HIS$GET_CONN_INFO routine • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
02392084	A processor interrupt was generated by the controller's XOR engine (FX), with no bits set in the CSR to indicate a reason for the interrupt. • Last Failure Parameter[0] contains the FX Control and Status Register (CSR). • Last Failure Parameter[1] contains the FX DMA Indirect List Pointer register (DILP). • Last Failure Parameter[2] contains the FX DMA Page Address register (DADDR). • Last Failure Parameter[3] contains the FX DMA Command and control register (DCMD).

(continued on next page)

Table C–34 (Cont.) Value-Added Services Last Failure Codes

Code	Description
023A2084	A processor interrupt was generated by the controller's XOR engine (FX), indicating an unrecoverable error condition. • Last Failure Parameter[0] contains the FX Control and Status Register (CSR). • Last Failure Parameter[1] contains the FX DMA Indirect List Pointer register (DILP). • Last Failure Parameter[2] contains the FX DMA Page Address register (DADDR). • Last Failure Parameter[3] contains the FX DMA Command and control register (DCMD).
02440100	The logical unit mapping type was detected invalid in va_set_disk_ geometry( )
02530102	An invalid status was returned from CACHE$LOOKUP_LOCK( ). • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
02560102	An invalid status was returned from CACHE$LOOKUP_LOCK( ). • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
02570102	An invalid status was returned from VA$XFER( ) during a operation. • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
025A0102	An invalid status was returned from CACHE$LOOKUP_LOCK( ). • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
025B0102	An invalid mapping type was specified for a logical unit. • Last Failure Parameter[0] contains the USB address. • Last Failure Parameter[1] contains the Unit Mapping Type.
025C0102	An invalid mapping type was specified for a logical unit. • Last Failure Parameter[0] contains the USB address. • Last Failure Parameter[1] contains the Unit Mapping Type.

(continued on next page)

Table C–34 (Cont.) Value-Added Services Last Failure Codes

Code	Description
02620102	An invalid status was returned from CACHE$LOOKUP_LOCK( ). • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
02690102	An invalid status was returned from CACHE$OFFER_WRITE_DATA( ). • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
02730100	A request was made to write a device metadata block with an invalid block type.
02790102	An invalid status was returned from VA$XFER( ) in a complex read operation. • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
027B0102	An invalid status was returned from VA$XFER( ) in a complex ACCESS operation. • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
027D0100	Unable to allocate memory for a Failover Control Block.
027E0100	Unable to allocate memory for a Failover Control Block.
027F0100	Unable to allocate memory for a Failover Control Block.
02800100	Unable to allocate memory for a Failover Control Block.
02820100	Unable to allocate memory for the Dirty Count Array.
02830100	Unable to allocate memory for the Cache Buffer Index Array.
02840100	Unable to allocate memory for the XNode Array.
02850100	Cache was declared bad by the Cache Diagnostics after first Meg was tested. Cannot recover and use local memory because cannot get those initial buffers back.
02860100	Unable to a'locate memory for the Fault Management Event Information Packet used by the Cache Manager in generating error logs to the host.
02880100	Invalid FOC Message in cmfoc_snd_cmd.
02890100	Invalid FOC Message in cmfoc_rcv_cmd.
028A0100	Invalid return status from DIAG$CACHE_MEMORY_TEST.
028B0100	Invalid return status from DIAG$CACHE_MEMORY_TEST.
028C0100	Invalid error status given to cache_fail.
028E0100	Invalid DCA state detected in init_crashover.
028F0100	Invalid status returned from CACHE$CHECK_METADATA.
02900100	Unable to allocate memory for the First Cache Buffer Index Array.

(continued on next page)

Table C–34 (Cont.) Value-Added Services Last Failure Codes

Code	Description
02910100	Invalid metadata combination detected in build_raid_node.
02920100	Unable to handle that many bad dirty pages (exceeded MAX_BAD_ DIRTY). Cache memory is bad.
02930100	There was no free or freeable buffer to convert bad metadata or to borrow a buffer during failover of bad dirty.
02940100	A free Device Correlation Array entry could not be found during write-back cache failover.
02950100	Invalid DCA state detected in start_crashover.
02960100	Invalid DCA state detected in start_failover.
02970100	Invalid DCA state detected in init_failover.
02990100	A free RAID Correlation Array entry could not be found during write-back cache failover.
029A0100	Invalid cache buffer metadata detected while scanning the Buffer Metadata Array. Found a page containing dirty data but the corresponding Device Correlation Array entry does exist.
029D0100	Invalid metadata combination detected in build_bad_raid_node.
029F0100	The Cache Manager software has insufficient resources to handle a buffer request pending.
02A00100	VA change state is trying to change device affinity and the cache has data for this device.
02A10100	Pubs not one when transportable
02A20100	Pubs not one when transportable
02A30100	No available data buffers. If the cache module exists then this is true after testing the whole cache. Otherwise there were no buffers allocated from BUFFER memory on the controller module.
02A40100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when allocating VAXDs.
02A50100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when allocating DILPs.
02A60100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when allocating Change State Work Items.
02A70100	A call to EXEC$ALLOCATE_MEM_ZEROED failed to return memory when allocating VA Request Items.
02A80000	Controller is being rebooted as a result of a CLI CLEAR INVALID_CACHE command being issued on the other controller.
02A90100	Too many pending FOC$SEND requests by the Cache Manager. Code is not designed to handle more than one FOC$SEND to be pending because there's no reason to expect more than one pending.
)2AA0100	An invalid call was made to CACHE$DEALLOCATE_CLD. Either that device had dirty data or it was bound to a RAIDset.
02AB0100	An invalid call was made to CACHE$DEALLOCATE_SLD. A RAIDset member either had dirty data or write-back already turned on.
02AC0100	An invalid call was made to CACHE$DEALLOCATE_SLD. The RAIDset still has data (strip nodes).

(continued on next page)

Table C–34 (Cont.) Value-Added Services Last Failure Codes

Code	Description
02AD0180	The FX detected a compare error for data that was identical. This error has always previously occurred due to a hardware problem.
02AE0100	The mirrorset member count and individual member states are inconsistent. Discovered during a mirrorset write or erase.
02AF0102	An invalid status was returned from VA$XFER( ) in a write operation. • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
02B00102	An invalid status was returned from VA$XFER( ) in an erase operation. • Last Failure Parameter[0] contains the DD address. • Last Failure Parameter[1] contains the invalid status.
02B10100	A mirrorset read operation was received and the round robin selection algorithm found no normal members in the mirrorset. Internal inconsistency.
02BC0100	A mirrorset read operation was received and the preferred member selection algorithm found no normal members in the mirrorset. Internal inconsistency.
02BD0100	A mirrorset metadata online operation found no normal members in the mirrorset. Internal inconsistency.

Table C–35 Device Services Last Failure Codes

Code	Description
03020101	Invalid SCSI direct-access device opcode in misc command DWD. Last Failure Parameter[0] contains the SCSI command opcode.
03030101	Invalid SCSI sequential-access device opcode in misc cmd DWD. Last Failure Parameter[0] contains the SCSI command opcode.
03040101	Invalid SCSI CD-ROM device opcode in misc command DWD. Last Failure Parameter[0] contains the SCSI command opcode.
03050101	Invalid SCSI medium changer device opcode in misc command DWD. Last Failure Parameter[0] contains the SCSI command opcode.
03060101	Invalid SCSI device type in PUB. Last Failure Parameter[0] contains the SCSI device type.
03070101	Invalid CDB Group Code detected during create of misc cmd DWD Last Failure Parameter[0] contains the SCSI command opcode.
03080101	Invalid SCSI OPTICAL MEMORY device opcode in misc command DWD. Last Failure Parameter[0] contains the SCSI command opcode.
030A0100	Error DWD not found in port in_proc_q.

(continued on next page)

Table C–35 (Cont.) Device Services Last Failure Codes

Code	Description
030B0188	A dip error was detected when pcb_busy was set. • Last Failure Parameter[0] contains the PCB port_ptr value. • Last Failure Parameter[1] contains the new info NULL - SSTAT0 - DSTAT - ISTAT. • Last Failure Parameter[2] contains the PCB copy of the device port DBC register. • Last Failure Parameter[3] contains the PCB copy of the device port DNAD register. • Last Failure Parameter[4] contains the PCB copy of the device portDSP register. • Last Failure Parameter[5] contains the PCB copy of the device portDSPS register. • Last Failure Parameter[6] contains the PCB copies of the device port SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the device port LCRC/RESERVED/ISTAT/DFIFO registers.
03150100	More DBDs than allowded for in mask
031E0100	Cannot find in_error dwd on in-process queue
031F0100	either DWD_ptr is null or bad value in dsps
03280100	SCSI CDB contains an invalid group code for a transfer command
03290100	The required error information packet (EIP) or device work descriptor (DWD) were not supplied to the Device Services error logging code.
032A0100	HIS$GET_CONN_INFO( ) returned an unexpected completion code.
032B0100	A Device Work Discriptor (DWD) was supplied with a NULL Physical Unit Block (PUB) pointer.
03320101	An invalid code was passed to the error recovery thread in the error_stat field of the PCB. Last Failure Parameter[0] contains the PCB error_stat code.

(continued on next page)

Table C–35 (Cont.) Device Services Last Failure Codes

Code	Description
03330188	A parity error was detected by a 710 while sending data out onto the SCSI bus. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.
03350188	The TEA (bus fault) signal was asserted into a 710. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.

(continued on next page)

Table C–35 (Cont.) Device Services Last Failure Codes

Code	Description
03360188	A 710's host bus watchdog timer expired. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.
03370108	A 710 detected an illegal script instruction. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.

(continued on next page)

Table C–35 (Cont.) Device Services Last Failure Codes

Code	Description
03380188	A 710's DSTAT register contains multiple asserted bits, or an invalidly asserted bit, or both. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.
03390108	An unknown interrupt code was found in a 710's DSPS register. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.
033C0101	An invalid code was seen by the error recovery thread in the er_funct_step field of the PCB. Last Failure Parameter[0] contains the PCB er_funct_step code.

(continued on next page)

Table C–35 (Cont.) Device Services Last Failure Codes

Code	Description
033E0108	An attempt was made to restart a 710 at the SDP DBD. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.
033F0108	An EDC error was detected on a read of a soft-sectored device - path not yet implemented. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.
03410101	Invalid SCSI device type in PUB. Last Failure Parameter[0] contains the PUB SCSI device type.

(continued on next page)

Table C–35 (Cont.) Device Services Last Failure Codes

Code	Description
03420188	A UDC interrupt could not be associated with either a DWD or the noncallable scripts. • Last Failure Parameter[0] contains the PCB reg710_ptr value. • Last Failure Parameter[1] contains the PCB copy of the 710 TEMP register. • Last Failure Parameter[2] contains the PCB copy of the 710 DBC register. • Last Failure Parameter[3] contains the PCB copy of the 710 DNAD register. • Last Failure Parameter[4] contains the PCB copy of the 710 DSP register. • Last Failure Parameter[5] contains the PCB copy of the 710 DSPS register. • Last Failure Parameter[6] contains the PCB copies of the 710 SSTAT2/SSTAT1/SSTAT0/DSTAT registers. • Last Failure Parameter[7] contains the PCB copies of the 710 LCRC/RESERVED/ISTAT/DFIFO registers.
03470100	Insufficient memory available for target block allocation.
03480100	Insufficient memory available for 710 info block allocation.
03490100	Insufficient memory available for autoconfig buffer allocation.
034A0100	Insufficient memory available for PUB allocation.
034B0100	Insufficient memory available for DS init buffer allocation.
034C0100	Insufficient memory available for static structure allocation.
034D0100	DS init DWDs exhausted.
034E2080	Diagnostics report all NCR710s are broken.
03500100	Insufficient memory available for command disk allocation.
03510100	Insufficient resourses available for command disk data region.
03520100	A failure resulted when an attempt was made to allocate a DWD for use by DS CDI.

Table C–36 Fault Manager Last Failure Codes

Code	Description
04010101	The requestor id component of the instance code passed to FM$REPORT_EVENT is larger than the maximum allowed for this environment. Last Failure Parameter[0] contains the instance code value.

(continued on next page)

Table C–36 (Cont.) Fault Manager Last Failure Codes

Code	Description
04020102	The requestor's error table index passed to FM$REPORT_EVENT is larger than the maximum allowed for this requestor. • Last Failure Parameter[0] contains the instance code value. • Last Failure Parameter[1] contains the requester error table index value.
04030102	The USB index supplied in the EIP is larger than the maximum number of USBs. • Last Failure Parameter[0] contains the instance code value. • Last Failure Parameter[1] contains the USB index value.
04040103	The event log format found in V_fm_template_table is not supported by the Fault Manager. The bad format was discovered while trying to fill in a supplied eip. • Last Failure Parameter[0] contains the instance code value. • Last Failure Parameter[1] contains the format code value. • Last Failure Parameter[2] contains the requester error table index value.
04050100	The Fault Manager could not allocate memory for his Event Information Packet (EIP) buffers.
04060100	The Fault Manager could not allocate a Datagram HTB in his initialization routine.
04070103	There is more EIP information than will fit into a datagram. The requestor specific size is probably too large. • Last Failure Parameter[0] contains the instance code value. • Last Failure Parameter[1] contains the format code value. • Last Failure Parameter[2] contains the requester error table index value.
04080102	The event log format found in the already built eip is not supported by the Fault Manager. The bad format was discovered while trying to copy the eip information into a datagram HTB. • Last Failure Parameter[0] contains the format code value. • Last Failure Parameter[1] contains the instance code value.
04090100	The caller of FM$CANCEL_EVENT_NOTIFICATION passed an address of an event notification routine which does not match the address of any routines for which event notification is enabled.
040D0100	FM$ENABLE_EVENT_NOTIFICATION was called to enable eip notification but the specified routine was already enabled to receive eip notification.

(continued on next page)

Table C–36 (Cont.) Fault Manager Last Failure Codes

Code	Description
040F0102	The eip->generic.mscp1.flgs field of the EIP passed to FM$REPORT_EVENT contains an invalid flag. • Last Failure Parameter[0] contains the instance code value. • Last Failure Parameter[1] contains the value supplied in the eip->generic.mscp1.flgs field.
04100101	Unexpected template type found during fmu_display_errlog processing. Last Failure Parameter[0] contains the unexpected template value.
04110101	Unexpected instance code found during fmu_memerr_report processing. Last Failure Parameter[0] contains the unexpected instance code value.
04120101	CLIB$SDD_FAO call failed. Last Failure Parameter[0] contains the failure status code value.

Table C–37 Common Library Routines Last Failure Codes

Code	Description
05010100	In recursive_noconflict could not get enough memory for scanning the keyword tables for configuration name conflicts.
05030100	In clib$scan_for_noconflict could not get enough memory for scanning the keyword tables for configuration name conflicts.

Table C–38 Dual Universal Asynchronous Receiver/Transmitter Services Last Failure Codes

Code	Description
06010100	The DUART was unable to allocate enough memory to establish a connection to the CLI.
06020100	A port other than terminal port A was referred to by a set terminal characteristics command. This is illegal.
06030100	A DUP question or default question message type was passed to the DUART driver, but the pointer to the input area to receive the response to the question was NULL.
06040100	Attempted to detach unattached maintenance terminal.
06050100	Attempted output to unattached maintenance terminal.
06060100	Attempted input from output only maintenance terminal service.

Table C–39 Failover Control Last Failure Codes

Code	Description
07010100	All available slots in the FOC notify table are filled
07020100	FOC$CANCEL_NOTIFY( ) was called to disable notification for a rtn that did not have notification enabled.
07030100	Unable to start the Failover Control Timer before main loop.

(continued on next page)

Table C–39 (Cont.) Failover Control Last Failure Codes

Code	Description
07040100	Unable to restart the Failover Control Timer.
07050100	Unable to allocate flush buffer.
07060100	Unable to allocate active receive fcb.
07070100	The other controller killed this, but could not assert the kill line because nindy on or in debug. So it killed this now.

Table C–40 Nonvolatile Parameter Memory Failover Control Last Failure Codes

Code	Description
08010101	A remote state change was received from the FOC thread that NVFOC does not recognize. Last Failure Parameter[0] contains the unrecognized state value.
08020100	No memory could be allocated for a NVFOC information packet.
08030101	Work received on the S_nvfoc_bque did not have a NVFOC work id. Last Failure Parameter[0] contains the id type value that was received on the NVFOC work queue.
08040101	Unknown work value received by the S_nvfoc_bque. Last Failure Parameter[0] contains the unknown work value
08050100	An unlock was received and the controller was not locked by the other controller.
08060100	A really write command was received when the NV memory was not locked.
08070100	A write to NV memory was received while not locked.
08080000	The other controller requested this controller to restart.
08090010	The other controller requested this controller to shutdown.
080A0000	The other controller requested this controller to selftest.
080B0100	Could not get enough memory to build a FCB to send to the remote routines on the other controller.
080C0100	Could not get enough memory for FCBs to receive information from the other controller.
080D0100	Could not get enough memory to build a FCB to reply to a request from the other controller.
080E0101	An out-of-range receiver ID was received by the NVFOC communication utility (master send to slave send ACK). Last Failure Parameter[0] contains the bad id value.
080F0101	An out-of-range receiver ID was received by the NVFOC communication utility (received by master). Last Failure Parameter[0] contains the bad id value.
08100101	A call to NVFOC$TRANSACTION had a from field (id) that was out of range for the NVFOC communication utility. Last Failure Parameter[0] contains the bad id value.
08110101	NVFOC tried to defer more than one FOC send. Last Failure Parameter[0] contains the master ID of the connection that had the multiple delays.

(continued on next page)

Table C–40 (Cont.) Nonvolatile Parameter Memory Failover Control Last Failure Codes

Code	Description
08120100	Unable to lock other controller's NVmemory despite the fact that the running and handshake_complete flags are set.
08130100	Could not allocate memory to build a callback context block on an unlock NVmemory call.
08140100	Could not allocate memory to build a workblock to queue to the NVFOC thread.
08150100	A lock was requested by the other controller but the memory is already locked by the other controller.
08160100	A request to clear the remote configuration was received but the memory was not locked.
08170100	A request to read the next configuration was received but the memory was not locked.
08180100	Could not get enough memory for FLS FCBs to receive information from the other controller.
08190100	An unlock command was received when the NV memory was not locked.
081A0100	Unable to allocate memory for remote work
081B0101	Bad remote work received on remote work queue Last Failure Parameter[0] contains the id type value that was received on the NVFOC remote work queue.
081C0101	Bad member management work received Last Failure Parameter[0] contains the bad member management value that was detected.

Table C–41 Command Line Interpreter Last Failure Codes

Code	Description
20010100	The action for work on the CLI queue should be CLI_CONNECT, CLI_COMMAND_IN or CLI_PROMPT. If it is not one of these three, this bugcheck will result.
20020100	The FAO returned a nonsuccessful response. This will only happen if a bad format is detected or the formatted string overflows the output buffer.
20030100	The type of work received on the CLI work queue was not of type CLI.
20070100	A work item of an unknown type was placed on the CLI's DUP Virtual Terminal thread's work queue by the CLI.
20080000	This controller requested this controller to restart.
20090010	This controller requested this controller to shutdown.
200A0000	This controller requested this controller to selftest.
200B0100	Could not get enough memory for FCBs to receive information from the other controller.
200C0100	After a CLI command the NV memory was still locked. The CLI should *always* unlock NV memory when the command is complete (if it had an error or not).
200D0101	After many calls to DS$PORT_BLOCKED, we never got a FALSE status back (which signals that nothing is blocked). Last Failure Parameter[0] contains the port number (1 - n) that we were waiting on to be unblocked.

(continued on next page)

Table C–41 (Cont.) Command Line Interpreter Last Failure Codes

Code	Description
200E0101	While traversing the structure of a unit, a config_info node was discovered with an unrecognized structure type. Last Failure Parameter[0] contains the structure type number that was unrecognized.
200F0101	A config_info node was discovered with an unrecognized structure type. Last Failure Parameter[0] contains the structure type number that was unrecognized.
20100101	A config_node of type VA_MA_DEVICE had an unrecognized SCSI device type. Last Failure Parameter[0] contains the SCSI device type number that was unrecognized.
20110100	An attempt to allocate memory so the CLI prompt messages could be deleted failed.
20120101	While traversing the structure of a unit, a config_info node was discovered with an unrecognized structure type. Last Failure Parameter[0] contains the structure type number that was unrecognized.
20130101	While traversing the structure of a unit, the device was of a unrecognized type. Last Failure Parameter[0] contains the SCSI device type that was unrecognized.

Table C–42 Host Interconnect Services Last Failure Codes

Code	Description
40000101	An unrecognized CI opcode was received by HIS. These packets are packets with CI opcodes recognized by the port but not by HIS. Last Failure Parameter[0] contains the CI opcode value.
40150100	LOCAL VC Timer in unexpected state.
40280100	Failed to allocate Buffer Name Table.
40290100	Failed to allocate ACB.
402A0100	Failed to allocate ID member template.
402B0100	Failed to allocate DG HTBs.
402C0100	Failed to allocate message HTBs.
402D0101	S_max_node greater than MAX_VC_ENTRIES. Last Failure Parameter[0] contains the S_ci_max_nodes value.
402E0101	S_max_node not set to valid value (8, 16, 32, 64, 128, 256). Last Failure Parameter[0] contains the S_ci_max_nodes value.
402F0100	Failure to allocate a HIS EIP structure
40300100	Failure in memory allocation
40510100	htb_id type not DG, when attempting to deallocate DG HTB.
40520100	htb_id type not RCV_SND, when attempting to dealloc recv queue HTB.
40530100	htb_id type not RCV_SND, when attempting to dealloc SCS queue HTB.
40560100	Failed to find a vc entry for ccb during his_close_connection routine.
407B0100	SCS command timeout unexpectedly inactive during SCS Accept Request.
407C0100	SCS command timeout unexpectedly inactive during SCS Reject Request.
408E0100	Message receive queue count disagrees with # HTBs on the queue.

(continued on next page)

Table C–42 (Cont.) Host Interconnect Services Last Failure Codes

Code	Description
408F0100	Unrecognized HTB id type.
40900100	htb_id type not DG, when attempting to xmit DG HTB.
40930100	Message receive queue count disagrees with # HTBs on the queue.
40950100	Create xfer request with 0-byte count
40960100	Create xfer request with 0-byte count
40970100	Create xfer request with 0-byte count
40980100	Create xfer request with 0-byte count
409D0100	Illegal return value from HIS$MAP.
40B40101	Invalid value in max_nodes field of se_params structure. Last Failure Parameter[0] contains the max_nodes field value.

Table C–43 Host Interconnect Port Services Last Failure Codes

Code	Description
42000100	Cmpl_main routine found invalid port xmt status.
42020100	Cannot start timer.
42030100	Cannot restart work timer.
42060100	HP_INIT could not allocate initial buffers.
420B0100	HP_INIT could not allocate initial bufs for Path A dl_ctl table.
42120100	Host Port polling HTB failed to complete in time.
42332080	Receive_main found destination address in the rcv packet does not match node address.
42340100	HP could not allocate buffers for I/O rundown in VC Close.
42382080	Ci_isr found that the yaci hardware had invalid xmt status on Path A, no bits set.
42392080	Ci_isr found that the yaci hardware had invalid xmt status on Path B, no bits set.
423A2080	CI_ISR found the abort bit set with out any valid reason. Path A.
423B2080	CI_ISR found transmit parity error without abort bit set. Path A.
423C2080	CI_ISR found buffer underflow without abort bit set. Path A.
423D2080	CI_ISR found the abort bit set with out any valid reason. Path B.
423E2080	CI_ISR found transmit parity error without abort bit set. Path B.
423F2080	CI_ISR found buffer underflow without abort bit set. Path B.
42442080	Ci_isr found that yaci hardware had a parity error.
42452080	Ci_isr found that yaci hardware had a bus timeout error.
42472080	Ci_isr found Data parity on Transmit Path A.
42482080	Ci_isr found Data parity on Transmit Path B.
424B0001	Ci_isr found Host Reset on Path A. Last Failure Parameter[0] contains the node number of the resetting node.

(continued on next page)

Table C–43 (Cont.) Host Interconnect Port Services Last Failure Codes

Code	Description
424C0001	Ci_isr found Host Reset on Path B. Last Failure Parameter[0] contains the node number of the resetting node.
424D2080	Ci_isr found Fetch parity on Transmit Path A.
424E2080	Ci_isr found Fetch parity on Transmit Path B.
424F0100	HP could not allocate buffers to repopulate dds when we close path.
42590001	Ci_isr found Host Reset on Path A. Last Failure Parameter[0] contains the node number of the resetting node.
425A0001	Ci_isr found Host Reset on Path B. Last Failure Parameter[0] contains the node number of the resetting node.
425B2080	CI_ISR found the abort bit set with out any valid reason. Path A.
425C2080	CI_ISR found transmit parity error without abort bit set. Path A.
425D2080	CI_ISR found buffer underflow without abort bit set. Path A.
425E2080	Ci_isr found that the yaci hardware had invalid xmt status on path a, no bits set.
425F2080	CI_ISR found the abort bit set with out any valid reason. Path B.
42602080	CI_ISR found transmit parity error without abort bit set. Path B.
42612080	CI_ISR found buffer underflow without abort bit set. Path B.
42622080	Ci_isr found that the yaci hardware had invalid xmt status on Path B, no bits set.
42632080	Receive_main found destination address in the rcv packet does not match node address.

Table C–44 Disk and Tape MSCP Server Last Failure Codes

Code	Description
60050100	Invalid return value from routine HIS$CONNECT, while DCD attempting to establish connection to a remote subsystem.
60070100	Invalid return value from routine HIS$MAP, while dmscp_dcd_allocate_bh attempting to map a buffer.
60090100	Invalid return value from routine HIS$DISCONNECT, while dmscp_dcd_comm_path_event attempting to disconnect a remote source connection.
600C0100	Invalid return value from routine RESMGR$ALLOCATE_DATA_SEGMENT.
600D0100	Opcode field in command being aborted is not valid.
60120100	Opcode of TMSCP command to be aborted is invalid.
60130100	tmscp_clear_sex_cdl_cmpl_rtn detected an unexpected opcode.
60140100	tmscp_clear_sex_cdl_cmpl_rtn detected an unexpected opcode.
60150100	VA$CHANGE_STATE failed to change the SW Write protect when requested to do so as part of the Disk Set Unit Characteristics command.
60160100	VA$CHANGE_STATE failed to change the SW Write protect when requested to do so as part of the Tape Set Unit Characteristics command.
60170100	Invalid type in entry of long interval work queue.

(continued on next page)

Table C–44 (Cont.) Disk and Tape MSCP Server Last Failure Codes

Code	Description
60180100	mscp_short_interval found an Invalid type in entry of long interval work queue.
60190100	dmscp_dcd_send_cmd found that the SIWI Work Item code supplied is unrecognized or invali in this context during DCD inhibited processing.
601B0100	Invalid EVENT_CODE parameter in call to dmscp_connection_event.
601C0100	Invalid EVENT_CODE parameter in call to tmscp_connection_event.
601D0100	Invalid EVENT_CODE parameter in call to dmscp_dcd_comm_path_event.
601E0100	Invalid EVENT_CODE parameter in call to dmscp_dcd_comm_path_event.
60250100	An attempt was about to be made to return a progress indicator to the host that was 0xFFFFFFFF, the only invalid value.
60260100	An WH_DAF command was requested to be performed by the wrong process.
60270100	A nonimmediate WHM operation was passed to the dmscp_exec_whm_ immediate routine.
60280100	This routine found an invalid xfer_state so cannot continue.
60290100	HIS did not allocate an HTB when there should have been one reserved for this connection as determined by mscp_rcv_listen.
602A0100	HIS did not allocate an HTB when there should have been one reserved for this connection as determined by dmscp_dcd_src_gcs_send.
602B0100	HIS did not allocate an HTB when there should have been one reserved for this connection as determined by dmscp_dcd_comm_path_event.
602C0100	When trying to put THE extra send-HTB on the connections send_htb_list there was already one on the queue.
602D0100	The VA$CHANGE_STATE service did not set the Software write protect as requested (for disk).
602E0100	The VA$CHANGE_STATE service did not set the Software write protect as requested (for tape).
603B0100	Initial HIS$LISTEN call for MSCP$DISK was unsuccessful.
603C0100	Initial HIS$LISTEN call for MSCP$TAPE was unsuccessful.
60400100	Unrecognized or invalid in this context return value from routine RESMGR$ALLOCATE_DATA_SEGMENT, while dmscp_dcd_allocate_ dseg attempting to allocate a data segment.
60410100	Unrecognized or invalid in this context return value from routine RESMGR$ALLOCATE_DATA_BUFFERS, while dmscp_dcd_allocate_ dbuf attempting to allocate a data buffer.
60420100	dmscp_dcd_rmte_end_msg was unable to find a command message that corresponds to end message it is currently processing.
60440100	dmscp_dcd_src_gcs_cmpl found the command being GCSed is no longer at the head of the remote connection's queue.
60450100	dmscp_dcd_errlog_rvc found that an error log is not associated with a command, internal miscellaneous error logs are assumed to not be associated with a connection and remote miscellaneous error logs generation was not requested.

(continued on next page)

Table C–44 (Cont.) Disk and Tape MSCP Server Last Failure Codes

Code	Description
60460100	dmscp_dcd_elrt_scc_send was entered to issue a remote source connection SCC but was unable to find an available HTB on the connection's htb_list. With no active DCDs the connection should always have HTBs available.
60480100	tmscp_suc_avl_cmpl_rtn found the unit not in the available state.
60490100	tmscp_clear_sex_cdl_cmpl_rtn found the state change failed.
604A0100	tmscp_clear_sex_cdl_cmpl_rtn found the state change failed.
604B0100	Subroutine process_event returned a value to dmscp_dcd_comm_path_ event that indicates that an internal disconnect request occurred while processing an immediate communications event.
604D0100	Subroutine process_event returned a value to dmscp_dcd_comm_path_ event that indicates that a connection established event occurred while no DCD commands were active.
604F0100	tmscp_set_sex_cmpl_rtn found the state change failed.
60500100	dmscp_dcd_op_cmpl found an unrecognized P_STS value in a DCD HTB status field.
60550100	mscp_initialize unable to get LOCAL STATIC memory from exec for use as a local connection ITB.
60560100	mscp_initialize unable to get LOCAL STATIC memory from exec for use as an AVAILABLE ITB.
60570100	mscp_initialize unable to get LOCAL STATIC memory from exec for use as an AVAILABLE state change ITB.
60580100	mscp_initialize unable to get LOCAL STATIC memory from exec for use as a state change ITB.
605D0100	tmscp_onl_cleanup_rtn detected a failure in enabling variable speed mode suppression.
605E0100	tmscp_suc_cmpl_rtn detected a failure in enabling variable speed mode suppression.
605F0100	tmscp_suc_cmpl_rtn detected a failure in enabling variable speed mode suppression.
60610100	mscp_initialize unable to get BUFFER STATIC memory from exec for use as Write History Logs.
60620100	mscp_initialize unable to get LOCAL STATIC memory from exec for use as Write History Log Allocation Failure Lists.
60640100	Invalid condition when there exists no unused Write History Log Entries.
60650100	Attempting to block incoming requests for the tape/loader when it was unnexpectedly found already blocked.
60660100	Loader boundary block request to stall incoming requests to the tape/loader unit was not setup as expected.
60680100	VA$ENABLE_NOTIFICATION failed with insufficient resources at init time.
606B0100	mscp_foc_receive_cmd detected that the message sent from the other controller had an illegal usb index.
606C0100	mscp_foc_receive_cmd detected that the message sent from the other controller had an illegal exclusive access state.

(continued on next page)

Table C–44 (Cont.) Disk and Tape MSCP Server Last Failure Codes

Code	Description
606D0100	FOC provided mscp_foc_send_cmpl_rtn with an invalid status for the FOC$SEND transmit command completion.
606E0100	FOC provided mscp_foc_send_rsp_done with an invalid transmit status for the FOC$SEND transmit response completion.

Table C–45 Diagnostics and Utilities Protocol Server Last Failure Codes

Code	Description
61010000	Controller crash was intentionally caused by the execution of the CRASH program. This bugcheck does not indicate the occurance of a controller failure.
61020100	HIS$LISTEN call failed with INSUFFICIENT_RESOURCES.
61090100	LISTEN_CONNECTION_ESTABLISHED event from HIS specified a connection ID for a connection we already know about.
610B0100	Code Load or Code Patch utility in CLCP local program forced controller restart to force new code or patch to take effect.
610C0100	HIS has reported a connection event that should not be possible.

Table C–46 System Communication Services Directory Service Last Failure Codes

Code	Description
62000100	HIS$LISTEN call failed with INSUFFICIENT_RESOURCES.
62020100	Failure to allocate associated timer queue.
62030100	Failure to allocate connection id timers.

Table C–47 Disk Inline Exerciser (DILX) Last Failure Codes

Code	Description
80010100	An HTB was not available to issue an IO when it should have been..
80020100	A unit could not be dropped from testing because an available cmd failed.
80030100	DILX tried to release a facility that was not reserved by DILX.
80040100	DILX tried to change the unit state from MAINTENANCE_MODE to NORMAL but was rejected because of insufficient resources.
80050100	DILX tried to change the usb unit state from MAINTENANCE_MODE to NORMAL but DILX never received notification of a successful state change.
80060100	DILX tried to switch the unit state from MAINTENANCE_MODE to NORMAL but was not successful.
80070100	DILX aborted all cmds via va$d_abort( ) but the htbs have not been returned.
80080100	While DILX was deallocating his eip buffers, at least one could not be found.

(continued on next page)

Table C–47 (Cont.) Disk Inline Exerciser (DILX) Last Failure Codes

Code	Description
80090100	DILX received an end msg which corresponds to an op code not supported by DILX.
800A0100	DILX was was not able to restart his timer.
800B0100	DILX tried to issue an IO for an opcode not supported.
800C0100	DILX tried to issue a oneshot IO for an opcode not supported.
800D0100	A DILX device control block contains an unsupported unit_state.
800E0100	While trying to print an Event Information Packet, DILX discovered an unsupported MSCP error log format.
80100100	DILX could not compare buffers because no memory was available from EXEC$ALLOCATE_MEM_ZEROED.
80120100	DILX expected an eip to be on the receive eip q but no eips were there.
80130100	DILX was asked to fill a data buffer with an unsupported data pattern.
80140100	DILX could not process an unsupported answer in dx$reuse_params( ).

Table C–48 Tape Inline Exerciser (TILX) Last Failure Codes

Code	Description
81010100	An HTB was not available to issue an IO when it should have been
81020100	A unit could not be dropped from testing because an available cmd failed
81030100	TILX tried to release a facility that was not reserved by TILX
81040100	TILX tried to change the unit state from MAINTENANCE_MODE to NORMAL but was rejected because of insufficient resources
81050100	TILX tried to change the usb unit state from MAINTENANCE_MODE to NORMAL but TILX never received notification of a successful state change
81060100	TILX tried to switch the unit state from MAINTENANCE_MODE to NORMAL but was not successful
81070100	TILX aborted all cmds via va$d_abort( ) but the htbs have not been returned
81080100	While TILX was deallocating his eip buffers, at least one could not be found
81090100	TILX received an end msg which corresponds to an opcode not supported by TILX
810A0100	TILX was was not able to restart his timer
810B0100	TILX tried to issue an IO for an opcode not supported
810C0100	TILX tried to issue a oneshot IO for an opcode not supported
810D0100	A TILX device control block contains an unsupported unit_state
810E0100	TILX received an unsupported Value Added status in a Value added completion message
810F0100	TILX found an unsupported device control block substate while trying to build a cmd for the Basic Function test
81100100	TILX found an unsupported device control block substate while trying to build a cmd for the Read Only test

(continued on next page)

Table C–48 (Cont.) Tape Inline Exerciser (TILX) Last Failure Codes

Code	Description
81110100	TILX found an unsupported device control block substate while trying to build a cmd for the User Defined test
81120100	TILX received an EOT encountered while in a substate where EOT encountered should not occur
81130100	TILX calculated an illegal position type value while trying to generate a cmd for the position intensive phase of the Basic Function test
81140100	While trying to print an Event Information Packet, TILX discovered an unsupported MSCP error log format
811A0100	TILX expected a deferred error to be on the receive deferred error q but no deferred errors were there.
811B0100	TILX was asked to fill a data buffer with an unsupported data pattern.
811C0100	TILX could not process an unsupported answer in tx$reuse_params( ).

Table C–49 Device Configuration Utilities (CONFIG/CFMENU) Last Failure Codes

Code	Description
83020100	An unsupported message type or terminal request was received by the CONFIG virtual terminal code from the CLI
83030100	Not all alter_device requests from the CONFIG utility completed within the timeout interval
83050100	An unsupported message type or terminal request was received by the CFMENU utility code from the CLI
83060100	Not all alter_device requests from the CFMENU utility completed within the timeout interval

Table C–50 Clone Unit Utility (CLONE) Last Failure Codes

Code	Description
84010100	An unsupported message type or terminal request was received by the CLONE virtual terminal code from the CLI

Table C–51 Controller Restart Codes

Code	Description
0	Full firmware restart
1	No restart
2	Automatic hardware restart

C.5 Event Notification/Recovery Threshold

An Event Notification/Recovery Threshold value is assigned to each significant event that can be reported by an HS1CP controller. The Event Notification/Recovery Threshold values and their meanings are shown in Table C–52.

Table C–52 Event Notification/Recovery Threshold Classifications

Threshold Value	Classification	Description
01	IMMEDIATE	Failure or potential failure of a component critical to proper controller operation is indicated; immediate attention is required.
02	HARD	Failure of a component that affects controller performance or precludes access to a device connected to the controller is indicated.
0A	SOFT	An unexpected condition detected by a controller firmware component (such as protocol violations, host buffer access errors, internal inconsistencies, and so forth) is indicated.
64	INFORMATIONAL	An event having little or no effect on proper controller or device operation is indicated.

With the exception of events reported via the Disk Copy Data Correlation Event Log, the Event Notification/Recovery Threshold value assigned to a particular event is supplied in the NR Threshold subfield of the "instance code" field of the event log used to report the event. See Section C.3 for "instance code" field details.

Disk Copy Data Correlation Event Log Conditions

The Event Notification/Recovery Threshold Classification assigned to the following conditions reported via a Disk Copy Data Correlation Event Log is SOFT (see Table C–52):

- Subcommand Error (subcode "Destination—Command Timed Out")
- Subcommand Error (subcode "Source—Command Timed Out")
- Subcommand Error (subcode "Destination—Inconsistent State"), cases A, B, C, D, E, and F
- Controller Error (subcode "Local Connection Request Failed, Insufficient Resources to Request Local Connection")
- Controller Error (subcode "Remote Connection Request Failed, Insufficient Resources to Request Local Connection")

All other conditions that can be reported via the Disk Copy Data Correlation Event Log are not assigned a specific Event Notification/Recovery Threshold Classification because they can be correlated with the associated condition specific event log.

C.6 Recommended Repair Action

A Recommended Repair Action code is assigned to each significant event that can be reported by an HS1CP controller. The Recommended Repair Action codes and their meanings are shown in Table C–53.

Table C–53 Recommended Repair Action Codes

Code	Description
00	No action necessary.
01	An unrecoverable hardware detected fault occurred or an unrecoverable firmware inconsistency was detected, proceed with controller support avenues. Contact Digital Multivendor Customer Service.
02	Inconsistent/erroneous information received from the operating system, proceed with operating system software support avenues. Contact Digital Multivendor Customer Service.
03	Follow the recommended repair action contained in the "last failure code" field.
04	Two possible problem sources are indicated: • In the case of a shelf with dual power supplies, one of the power supplies has failed. Follow repair action 07 for the power supply with the Power LED out. • One of the shelf blowers has failed. Follow repair action 06.
05	Four possible problem sources are indicated: • Total power supply failure on a shelf. Follow repair action 09. • A device inserted into a shelf that has a broken internal SBB connector. Follow repair action 0A. • A standalone device is connected to the controller with an incorrect cable. Follow repair action 08. • A controller hardware failure. Follow repair action 20.
06	Determine which blower has failed and replace it.
07	Replace power supply.
08	Replace the cable. Refer to the specific device documentation.
09	Determine power failure cause.
0A	Determine which SBB has a failed connector and replace it.
0B	The other controller in a dual-redundant configuration has been reset with the "Kill" line by the controller that reported the event. To restart the "Killed" controller enter the CLI command RESTART OTHER on the "Surviving" controller and then depress the RESET button on the "Killed" controller. If the other controller is repeatedly being "Killed" for the same or a similar reason, follow repair action 20.

(continued on next page)

Table C–53 (Cont.) Recommended Repair Action Codes

Code	Description
0C	Both controllers in a dual-redundant configuration are attempting to use the same SCSI ID (either 6 or 7 as indicated in the event report). Note that the other controller of the dual-redundant pair has been reset with the "Kill" line by the controller that reported the event. Two possible problem sources are indicated. • A controller hardware failure • A controller backplane failure First, follow repair action 20 for the "Killed" controller. If the problem persists, then follow repair action 20 for the "Surviving" controller. If the problem still persists, then replace the controller backplane.
20	Replace the controller module.
22	Replace the indicated cache module.
23	Replace the indicated write cache battery. CAUTION: BATTERY REPLACEMENT MAY CAUSE INJURY.
24	Check for the following invalid write cache configurations: • If the wrong write cache module, replace with the matching module or clear the invalid cache error via the CLI (refer to Appendix E). • If the write cache module is missing, re-seat cache if it is actually present, or add the missing cache module or clear the invalid cache error via the CLI (refer to Appendix E). • If in a dual-redundant configuration and one of the write cache modules is missing, match write cache boards with both controllers. • If in a dual-redundant configuration and both caches are not of the same type (i.e, both write cache), replace a cache module to assure both are compatible. • If in a dual-redundant configuration and both write caches are not of the same size, replace a cache module to assure both are compatible.
25	An unrecoverable Memory System failure occurred. Upon restart the controller will generate one or more Memory System Failure Event Error Logs; follow the repair actions contained therein.
26	The Master DRAB detected a Cache Timeout condition. The cache regions in effect are identified in the Master DRAB RSR register: bits 8 through 11 identify the CACHEA memory region, bits 12 through 15 identify the CACHEB memory region, bits 20 through 23 the CACHEA DRAB registers region, and bits 24 through 27 identify the CACHEB DRAB registers region. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. Follow repair action 36.
27	The Master DRAB detected an Nbus Transfer Error Acknowledge (TEA) condition. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. Follow repair action 36.

(continued on next page)

Table C–53 (Cont.) Recommended Repair Action Codes

Code	Description
28	A Multiple Bit ECC error was detected by the Master DRAB. The Master DRAB DER register bits 0 through 6 contain the syndrome value. The Master DRAB EAR register combined with Master DRAB ERR bits 0 through 3 (address region) yields the affected memory address. The Master DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. Follow repair action 34.
29	A Multiple Bit ECC error was detected by the CACHEA0 or CACHEA1 DRAB. The CACHEAn DRAB DER register bits 0 through 6 contain the syndrome value. The CACHEAn DRAB EAR register combined with the Master DRAB RSR register bits 8 through 11 (CACHEA memory region) yields the affected memory address. The CACHEAn DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. Follow repair action 34.
2A	A Multiple Bit ECC error was detected by the CACHEB0 or CACHEB1 DRAB. The CACHEBn DRAB DER register bits 0 through 6 contain the syndrome value. The CACHEBn DRAB EAR register combined with the Master DRAB RSR register bits 12 through 15 (CACHEB memory region) yields the affected memory address. The CACHEBn DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. Follow repair action 34.
2B	The Master DRAB detected an Ibus to Nbus Timeout condition. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. If Master DRAB CSR register bits 10 through 12 contains the value 1 and "wdr1" register bit 26 is clear or Master DRAB CSR register bits 10 through 12 contains the value 2 and "wdr1" register bit 27 is clear or Master DRAB CSR register bits 10 through 12 contains the value 3 and "wdr1" register bit 28 is clear or Master DRAB CSR register bits 10 through 12 contains the value 4 and "wdr1" register bit 29 is clear or Master DRAB CSR register bits 10 through 12 contains the value 5 and "wdr1" register bit 30 is clear or Master DRAB CSR register bits 10 through 12 contains the value 6 and "wdr1" register bit 31 is clear, a firmware fault is indicated; follow repair action 01. Otherwise, follow repair action 36.
2C	The Master DRAB detected a Nonexistent Memory Error condition. The Master DRAB EAR register combined with Master DRAB ERR bits 0 through 3 (address region) yields the affected memory address. The Master DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. Follow repair action 36.

(continued on next page)

Table C–53 (Cont.) Recommended Repair Action Codes

Code	Description
2D	The CACHEA0 or CACHEA1 DRAB detected a Nonexistent Memory Error condition. The CACHEAn DRAB EAR register combined with the Master DRAB RSR register bits 8 through 11 (CACHEA memory region) yields the affected memory address. The CACHEAn DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. Follow repair action 36.
2E	The CACHEB0 or CACHEB1 DRAB detected a Nonexistent Memory Error condition. The CACHEBn DRAB EAR register combined with the Master DRAB RSR register bits 12 through 15 (CACHEB memory region) yields the affected memory address. The CACHEBn DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. Follow repair action 36.
2F	The Master DRAB detected an Address Parity Error or a Write Data Parity Error condition. The Master DRAB EAR register combined with Master DRAB ERR bits 0 through 3 (address region) yields the affected memory address. The Master DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. For Write Data Parity Error conditions Bits 0 through 3 of the Master DRAB CSR register identify the byte in error. For Address Parity Error conditions follow repair action 34. For Write Data Parity Error conditions follow repair action 35.
30	The CACHEA0 or CACHEA1 DRAB detected an Address Parity Error or a Write Data Parity Error condition. If the failure occurred during a memory refresh attempt, the CACHEAn DRAB EAR register combined with the Master DRAB RSR register bits 8 through 11 (CACHEA memory region) yields the affected memory address. If the failure occurred during a memory access attempt, the CACHEA0 DRAB EAR register combined with the Master DRAB RSR register bits 8 through 11 (CACHEA memory region) or bits 20 through 23 (CACHEA DRAB register region) yields the affected memory address. Unfortunately, no other information is available to distinguish a memory region access from a DRAB register region access. The CACHEAn DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is set, the failure was reported via the DRAB_INT. For Write Data Parity Error conditions bits 0 through 3 of the CACHEAn DRAB CSR register identify the byte in error. For Address Parity Error conditions follow repair action 34. For Write Data Parity Error conditions follow repair action 35.

(continued on next page)

Table C–53 (Cont.) Recommended Repair Action Codes

Code	Description
31	The CACHEB0 or CACHEB1 DRAB detected an Address Parity Error or a Write Data Parity Error condition. If the failure occurred during a memory refresh attempt, the CACHEBn DRAB EAR register combined with the Master DRAB RSR register bits 8 through 11 (CACHEB memory region) yields the affected memory address. If the failure occurred during a memory access attempt, the CACHEB0 DRAB EAR register combined with the Master DRAB RSR register bits 8 through 11 (CACHEB memory region) or bits 20 through 23 (CACHEB DRAB register region) yields the affected memory address. Unfortunately, no other information is available to distinguish a memory region access from a DRAB register region access. The CACHEBn DRAB EDR register contains the error data. If the failure involved a Device Port, the Master DRAB CSR register bits 10 through 12 identify that Device Port. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. For Write Data Parity Error conditions bits 0 through 3 of the CACHEBn DRAB CSR register identify the byte in error. For Address Parity Error conditions follow repair action 34. For Write Data Parity Error conditions follow repair action 35.
32	The Master DRAB detected an Ibus Parity Error condition. The Master DRAB EAR register combined with the Master DRAB ERR bits 4 through 7 (address region) yields the affected memory address. The Master DRAB EDR register contains the error data. If Master DRAB DSR register bit 14 is set, the failure was reported via the NMI. If Master DRAB DSR register bit 14 is clear, the failure was reported via the DRAB_INT. If bits 20 through 23 of the Master DRAB DCSR register contain a nonzero value, a firmware fault is indicated; follow repair action 01. Otherwise, follow repair action 36.
33	This event report contains supplemental information related to a Memory System Failure event report delivered earlier. Use the instance code contained in the "memory address" field of this event report to correlate this event report with the other event report.
34	If Bit 31 of the DCSR register of the DRAB that detected the failure is set, a firmware fault is indicated; follow repair action 01. Otherwise, follow repair action 36.
35	If bits 20 through 23 of the "wdr1" register contain a nonzero value, a firmware fault is indicated; follow repair action 01. Otherwise, follow repair action 36.
36	Unfortunately, no other information is available to aid in diagnosing the cause of the failure. If the Master DRAB detected the failure, follow repair action 20. If the CACHEAn or CACHEBn DRAB detected the failure, follow repair action 22. If the problem persists, follow repair action 01.
37	The Memory System Failure translator could not determine the failure cause. Follow repair action 01.
40	If the Sense Data FRU field is non-zero, follow repair action 41. Otherwise, replace the appropriate FRU associated with the device's SCSI interface or the entire device.
41	Consult the device's maintenance manual for guidance on replacing the indicated device FRU.
43	Update the configuration data to correct the problem.
44	Replace the SCSI cable for the failing SCSI bus. If the problem persists, replace the controller backplane, drive backplane, or controller module.
45	Interpreting the device supplied Sense Data is beyond the scope of the controller's firmware. See the device's service manual to determine the appropriate repair action, if any.

(continued on next page)

Table C–53 (Cont.) Recommended Repair Action Codes

Code	Description
60	Swap the transmit and receive cables for the indicated path.
61	Check indicated path cables for proper installation.
63	Check the CI adapter on the host system identified in the "remote node name" field for proper operation.
64	Excessive VC closures are occuring. Perform repair action 61 on both sets of path cables. If the problem persists, perform repair action 63.

Recommended Repair Action codes apply to each reportable event (except those reported via the Disk Copy Data Correlation Event Log) as identified by the value contained in the Repair Action subfield of the "instance code" field of the event logs described in Section C.3.

For events reported via the Last Failure Event Log the Recommended Repair Action code is contained in the Repair Action subfield of the "last failure code" field of that event log.

Disk Copy Data Correlation Event Log Conditions

The Recommended Repair Action Code assigned to the following conditions reported via a Disk Copy Data Correlation Event Log is 01 (see Table C–53):

- Subcommand Error (subcode "Destination—Command Timed Out")
- Subcommand Error (subcode "Source—Command Timed Out")
- Subcommand Error (subcode "Destination—Inconsistent State"), cases C, D, E, and F
- Controller Error (subcode "Local Connection Request Failed, Insufficient Resources to Request Local Connection")
- Controller Error (subcode "Remote Connection Request Failed, Insufficient Resources to Request Local Connection")

The Recommended Repair Action Code assigned to the following condition reported via a Disk Copy Data Correlation Event Log is 02 (see Table C–53):

- Subcommand Error (subcode "Source—Inconsistent State"), cases A and B.

All other conditions that can be reported via the Disk Copy Data Correlation Event Log are not assigned a specific Recommended Repair Action Code because they can be correlated with the associated condition specific event log.

C.7 Deskew Command Procedure

Example C–4 presents a command procedure to deskew the "CONTROLLER DEPENDENT INFORMATION" for a "CONTROLLER LOG" type error log.

Example C–4 Deskew Command Procedure Example

```
$! P1 = Input file name
$! P2 = Output file name
$ on warning then $exit
$ inew_entry = " ******************************** ENTRY"
$ ctrl_entry = "                                        CONTROLLER LOG"
$ lw_entry = "       LONGWORD"
$ ctrl_inp = "FALSE"
$ lw_string = ""
$ open/read inf 'p1'
$ open/write ouf 'p2'
$in_loop:
$ read/end=in_done inf inr
$ inrlen = f$length(inr)
$ if f$locate(new_entry,inr) .ne. inrlen
$ then
$  write sys$output inr
$  if ctrl_inp
$  then
$   gosub convert_longs
$   ctrl_inp = "FALSE"
$  endif
$ else
$  if f$locate(ctrl_entry,inr) .ne. inrlen
$  then
$   write sys$output inr
$   ctrl_inp = "TRUE"
$   lw_string = ""
$  endif
$  if f$locate(lw_entry,inr) .ne. inrlen .and. ctrl_inp
$  then
$   lw = f$element(2," ",f$edit(inr,"TRIM,COMPRESS"))
$   if lw_string .eqs. ""
$   then
$    lw_string = f$extract(0,4,lw)
$   else
$    lw_string = lw + lw_string
$   endif
$  endif
$ endif
$ write ouf inr
$ goto in_loop
$in_done:
$ close inf
$  if ctrl_inp
$  then
$   gosub convert_longs
$  endif
$ close ouf
$ exit
$convert_longs:
$ index = 1
$ write ouf "<FF>"
$ write ouf ""
```

(continued on next page)

Example C–4 (Cont.) Deskew Command Procedure Example

```
$ write ouf ""
$ write ouf ""
$ write ouf "LONGWORD DESKEW:"
$ write ouf ""
$ write ouf ""
$convert_longs_loop:
$ len = f$length(lw_string)
$ if len .le. 4 then goto convert_longs_done
$ lw = f$extract(len - 8,8,lw_string)
$ write ouf "        LONGWORD[''index'] = ",lw
$ lw_string = f$extract(0,len - 8, lw_string)
$ index = index + 1
$ goto convert_longs_loop
$convert_longs_done:
$ write ouf "<FF>"
$ return
```

Example C–5 shows an ERF error log before running the command procedure.

Example C–5 ERF Error Log Before Command Procedure

```
V A X / V M S        SYSTEM ERROR REPORT          COMPILED 16-MAR-1993 12:30:07
                                                                  PAGE 144.

******************************** ENTRY      11. ********************************
ERROR SEQUENCE 2820.                          LOGGED ON:         SID 05903914
DATE/TIME 16-MAR-1993 11:35:45.39                              SYS_TYPE 00000000
SYSTEM UPTIME: 2 DAYS 22:48:03
SCS NODE: CNOTE                                                VAX/VMS V5.5-2

ERL$LOGMESSAGE ENTRY  KA825  HW REV# B  PATCH REV# 28.  UCODE REV# 20.
                      BI NODE # 2.

I/O SUB-SYSTEM, UNIT _HS1CP2$DUA20:

      MESSAGE TYPE          0001
                                        DISK MSCP MESSAGE

      MSLG$L_CMD_REF    5B54001E
      MSLG$W_SEQ_NUM        0039
                                        SEQUENCE #57.

      MSLG$B_FORMAT           00
                                        CONTROLLER LOG

      MSLG$B_FLAGS            00
                                        UNRECOVERABLE ERROR

      MSLG$W_EVENT          01CA
                                        CONTROLLER ERROR
                                        POLICY PROCESS ERROR

      MSLG$Q_CNT_ID     00000021
                        01280001
                                        UNIQUE IDENTIFIER, 000100000021(X)
                                        MASS STORAGE CONTROLLER
                                        MODEL = 40.

      MSLG$B_CNT_SVR          FF
                                        CONTROLLER SOFTWARE VERSION #255.

      MSLG$B_CNT_HVR          00
                                        CONTROLLER HARDWARE REVISION #0.

CONTROLLER DEPENDENT INFORMATION

      LONGWORD 1.       01010000
                                        /..../

      LONGWORD 2.       044103CF
                                        /Ï.A./

      LONGWORD 3.       00000000
                                        /..../

      LONGWORD 4.       00470000
                                        /..G./

      LONGWORD 5.       00000000
                                        /..../

      LONGWORD 6.       00020000
                                        /..../

      LONGWORD 7.       00000000
                                        /..../
```

Example C–6 shows the same ERF error log after running the command procedure (notice the deskewed longwords).

Example C–6 ERF Error Log After Command Procedure

```
V A X / V M S        SYSTEM ERROR REPORT          COMPILED 16-MAR-1993 12:30:07
                                                                  PAGE 144.

******************************** ENTRY      11. ********************************
ERROR SEQUENCE 2820.                              LOGGED ON:        SID 05903914
DATE/TIME 16-MAR-1993 11:35:45.39                                SYS_TYPE 00000000
SYSTEM UPTIME: 2 DAYS 22:48:03
SCS NODE: CNOTE                                                   VAX/VMS V5.5-2

ERL$LOGMESSAGE ENTRY  KA825  HW REV# B  PATCH REV# 28.  UCODE REV# 20.
                      BI NODE # 2.

I/O SUB-SYSTEM, UNIT _HS1CP2$DUA20:

      MESSAGE TYPE          0001
                                        DISK MSCP MESSAGE

      MSLG$L_CMD_REF    5B54001E
      MSLG$W_SEQ_NUM        0039
                                        SEQUENCE #57.

      MSLG$B_FORMAT           00
                                        CONTROLLER LOG

      MSLG$B_FLAGS            00
                                        UNRECOVERABLE ERROR

      MSLG$W_EVENT          01CA
                                        CONTROLLER ERROR
                                        POLICY PROCESS ERROR

      MSLG$Q_CNT_ID     00000021
                        01280001
                                        UNIQUE IDENTIFIER, 000100000021(X)
                                        MASS STORAGE CONTROLLER
                                        MODEL = 40.

      MSLG$B_CNT_SVR          FF
                                        CONTROLLER SOFTWARE VERSION #255.

      MSLG$B_CNT_HVR          00
                                        CONTROLLER HARDWARE REVISION #0.

 CONTROLLER DEPENDENT INFORMATION

      LONGWORD 1.       01010000
                                        /..../

      LONGWORD 2.       044103CF
                                        /Ï.A./

      LONGWORD 3.       00000000
                                        /..../

      LONGWORD 4.       00470000
                                        /..G./

      LONGWORD 5.       00000000
                                        /..../

      LONGWORD 6.       00020000
                                        /..../

      LONGWORD 7.       00000000
                                        /..../

LONGWORD DESKEW:
```

(continued on next page)

Example C–6 (Cont.) ERF Error Log After Command Procedure

```
LONGWORD[1] = 03CF0101
LONGWORD[2] = 00000441
LONGWORD[3] = 00000000
LONGWORD[4] = 00000047
LONGWORD[5] = 00000000
LONGWORD[6] = 00000002
```

D

Configuration Rules and Restrictions

This appendix describes rules and restrictions associated with the electronic and mechanical configuration of the SW800 cabinet and components when configured as a FDDI Server. Its purpose is to provide a base-line reference for standard FDDI Server component placements and configurations.

Information about initial component installation or product upgrades can be found in the documentation that accompanies the upgrade installation kit. Access to cabinet component parts and component subassemblies is described in Chapter 8. Topics discussed in this chapter are:

- Standard FDDI Server configurations
- Electrostatic discharge protection
- Connecting a maintenance terminal
- Configuring the SW800 cabinet
- Configuring cabinet power
- Configuring internal bus cables for the controller and the server processor
- Configuring the HS1CP
- Configuring the server processor
- Running the EISA Configuration utility
- Upgrading existing FDDI Servers

D.1 Standard FDDI Server Configurations

There are generally two standard FDDI Server configurations hosted in the SW800 cabinet. The HS111 configuration is the minimal FDDI Server system. The basic HS111 system contains one HS1CP and a single cache memory module, one Cable Distribution Unit (CDU) and a single server processor. The HS121 configuration adds a second CDU, HS1CP and cache memory module, and a second server processor providing storage system redundancy. All configurations allow the customer to specify the number and type of storage devices installed in the FDDI Server.

Component parts of each of the HS111 or HS121 configurations are generally the same. Differences in specific component configurations and cable connections are provided when appropriate.

WARNING

Cabinet rail edges may be sharp and can slice or abrade skin and cable insulation.

D.2 Electrostatic Discharge Protection

Electrostatic discharge can damage subsystem components. Before working inside the SW800 cabinet or on any individual component, protect the subsystem components against ESD by using the following strategies:

- Maintain more than 40 percent humidity in the room where your subsystem resides.
- Place the subsystem cabinet away from heavy traffic paths.
- If possible, do not work in a carpeted area. If carpeting is required, choose an antistatic carpet. If any type of carpeting is installed, place antistatic mats around the system to decrease electrostatic discharge.

ESD Grounding Procedure

Prior to removing or replacing any module:

1. Obtain and attach an ESD wrist strap to your wrist. Ensure that the strap fits snugly around your wrist.
2. Plug or clip the other end of the ESD wrist strap to a ESD bolt or ground stud usually located on the cabinet's vertical rail that is common for both the device channel processor shelves and the storage shelves.
3. Obtain and use an approved antistatic bag and/or a grounded antistatic mat.

CAUTION

Do not attempt to replace or repair components within FRUs while power is applied or equipment damage may result.

D.3 Connecting a Maintenance Terminal

You do not need to have a locally connected maintenance terminal for normal operations, but a terminal is required during configuration and maintenance activities. Any terminal that supports ANSI control sequences can be used, including graphics displays that provide emulation of an ANSI compatible video terminal.

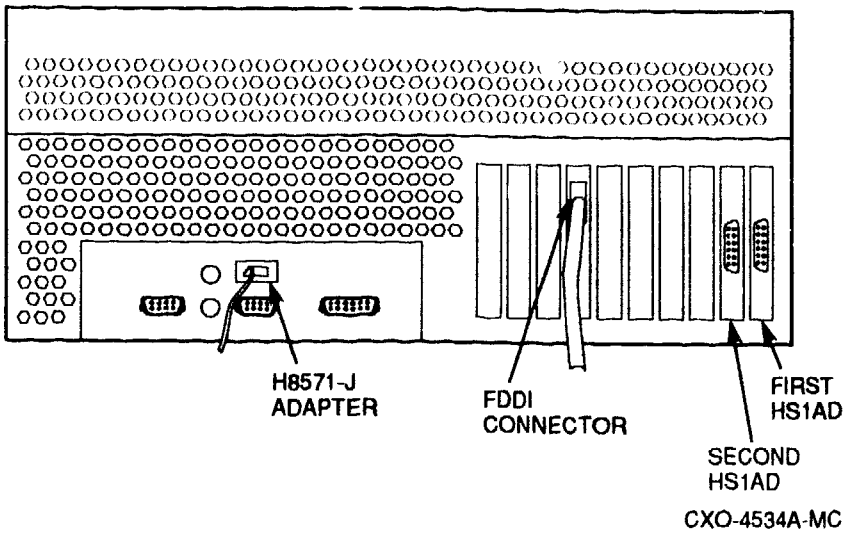
Note

Refer to your terminal documentation if you are unfamiliar with how to connect a terminal and set it up.

Use the following procedure to connect most EIA–423 compatible terminals:

1. Make sure the power switch on the back of the terminal is OFF (○).
2. Plug the terminal's power cord into a proper wall socket.
3. Connect one end of the terminal cable to the back of the terminal.
4. Connect the other end of the terminal cable to the EIA terminal port on the rear of the original server processor using the female MMJ adapter as shown in Figure D–1.

Figure D–1 Connecting a Maintenance Terminal to the Server Processor



5. Turn the terminal power switch to the ON (|) position.
6. Set the terminal's communication setup to 9600 baud, with 8 data bits, 1 stop bit, and no parity. Refer to your terminal documentation for terminal setup instructions.
7. Your terminal screen displays the DCL prompt $ when your server is up and running.
8. Press the Return key if no HS1CP> prompt is visible on the screen. This brings you to the HS1CP's command line interpreter (CLI) prompt.

D.4 Configuring the SW800 Cabinet

The information in this section describes physical configurations with respect to both standard and nonstandard (customized) FDDI Systems and subsystems. You can find additional information in the specific StorageWorks FDDI Server installation and upgrade documentation and in the specific cabinet, shelf, and configuration documentation.

D.4.1 General Configuration Rules for FDDI Server Cabinets

This section discusses the rules to apply to subsystem configurations in FDDI Servers. Refer to the *StorageWorks Solutions FDDI Server Data Center Cabinet Installation and User's Guide* for more details.

Figure D–2 shows the position of storage, server processor, and controller shelves (without tape drives) in a FDDI Server or SW800-series data center cabinet.

Note

In the illustrations in this chapter, the designator "S" indicates a BA350–S storage shelf and the designator "C" indicates a BA350–M controller shelf.

The following rules apply to the configuration of the FDDI Server cabinet:

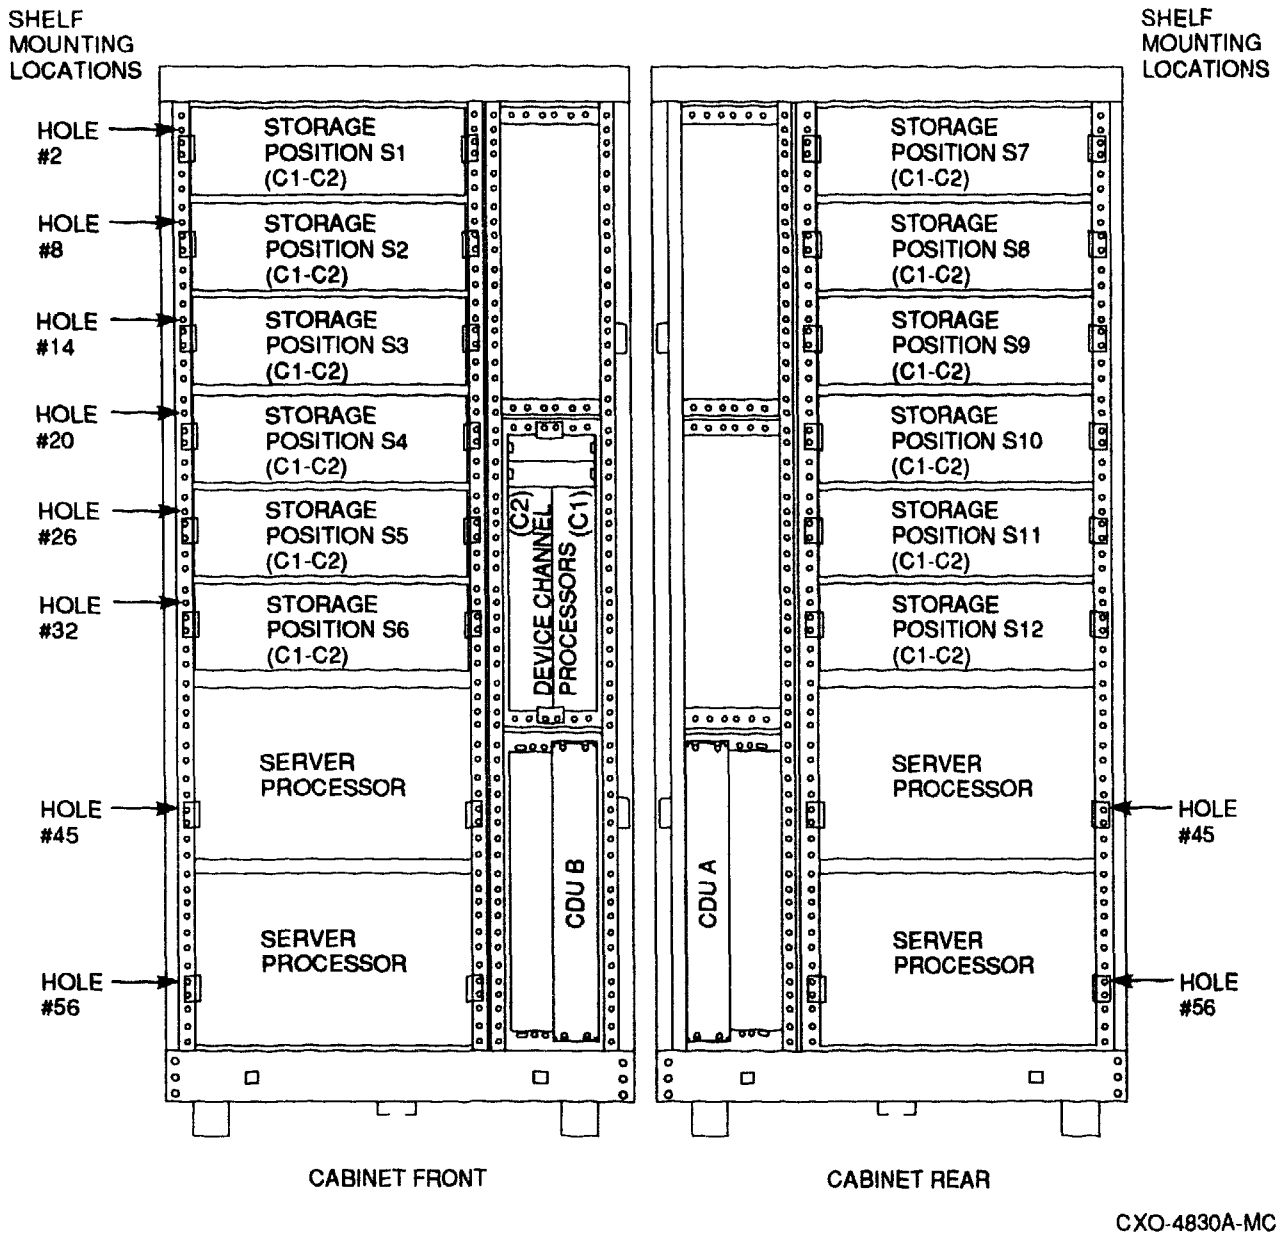
- Cabinet capacities are measured by the number of individual shelves that can be installed into the cabinet (for example, BA350–S storage shelves).

CAUTION

If the configuration of a cabinet is to be changed, abide by the cabinet-specific installation sequence to maintain cabinet stability. Failure to follow the sequence can result in tipping the cabinet.

- Shelf-specific configuration rules are in the *StorageWorks Solutions Shelf and SBB User's Guide*.
- Switch-controlled input power is required for each shelf.
- Bulkhead input/output (I/O) connections for the shelves, the host computer, and/or external controllers is available.
- A FDDI Server cabinet typically does not support dc power modules.
- Input power can be either 60 Hz or 50 Hz with the appropriate cabinet cable distribution unit.
- For the controller and storage shelves, a second cable distribution unit (CDU) is required to support more than a single power supply per shelf.

Figure D–2 HS121 Storage Configuration with Optimal Disk Storage and No TZ8*xx*–Series Tape Drives Installed



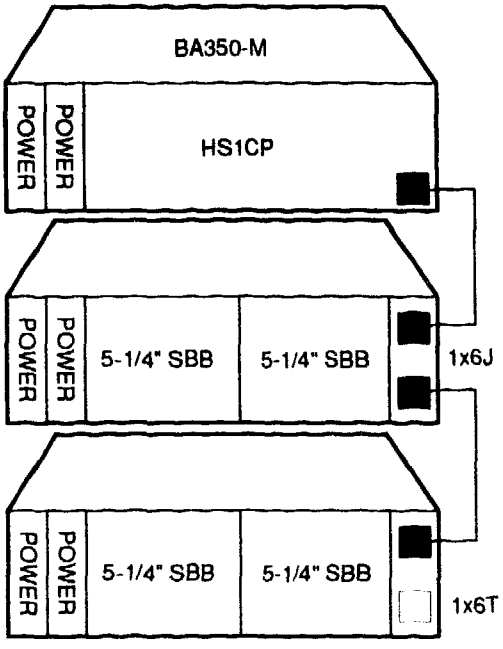
CXO-4830A-MC

The following rules apply to subsystem configurations in the FDDI Server cabinet:

- A standard of one or two BA350–M shelves connected to 6 to 12 BA350–S shelves in a single FDDI cabinet is suggested.
- Two storage shelves per port (jumpered-pairs)—Two BA350–S shelves can be joined on the same controller port with the following restrictions:
  - The SCSI–2 cable to the first BA350–S storage shelf is 1.0 meter or less. The associated BA350–M controller shelf must be positioned close enough to satisfy this restriction.

- The SCSI–2 cable from the first BA350–S shelf to the second shelf is 0.5 meters or less. This requires the two shelves to be immediately adjacent to each other.
- The first BA350–S storage shelf is configured for unterminated single SCSI. Figure D–3 shows an example of device shelves in a single extension configuration.

Figure D–3 Single Extension from Device Shelf to Device Shelf



CXO-4606A-MC

- Any TZ8*x*7 half-rack tape loader device must be located at the top front (each tape loader occupies the full cabinet depth) positions, filling the two or four top BA350–S shelf positions (front and back). Up to four tape drive loader devices can be loaded into a FDDI Server cabinet, leaving 4 BA350–S shelves.

  Figure D–4 and Figure D–5 show the position of the storage and controller shelves when two or four TZ8*xx*-series tape devices are installed in an FDDI Server cabinet.

Figure D–4 HS121 Storage Configuration with Four TZ8*xx*–Series Tape Drives Installed

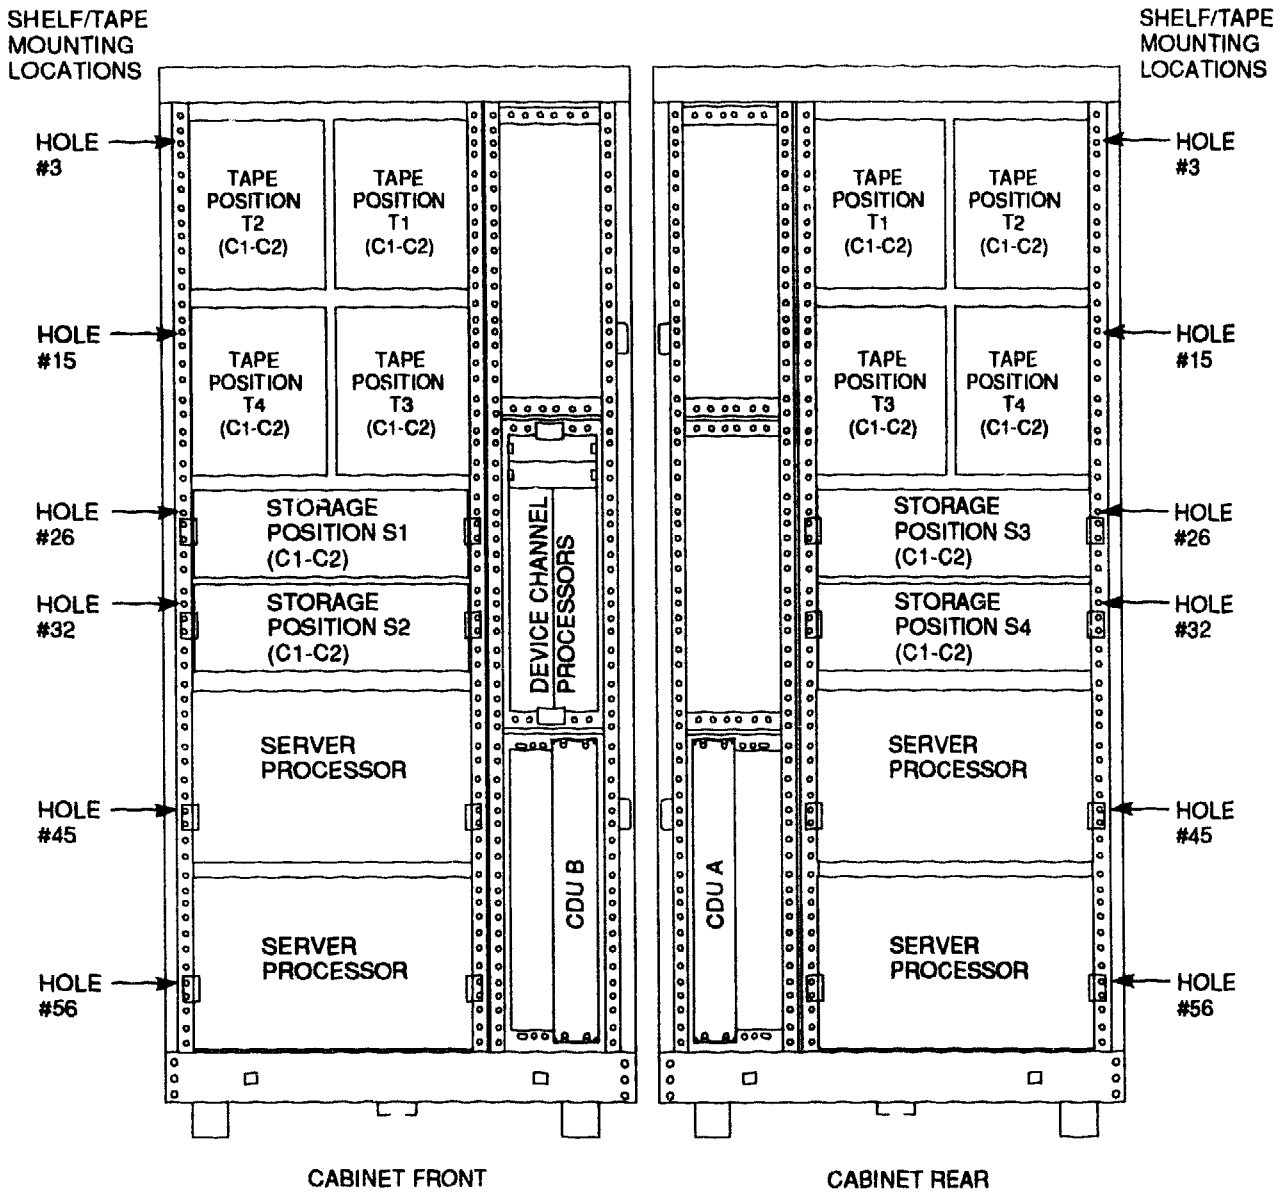
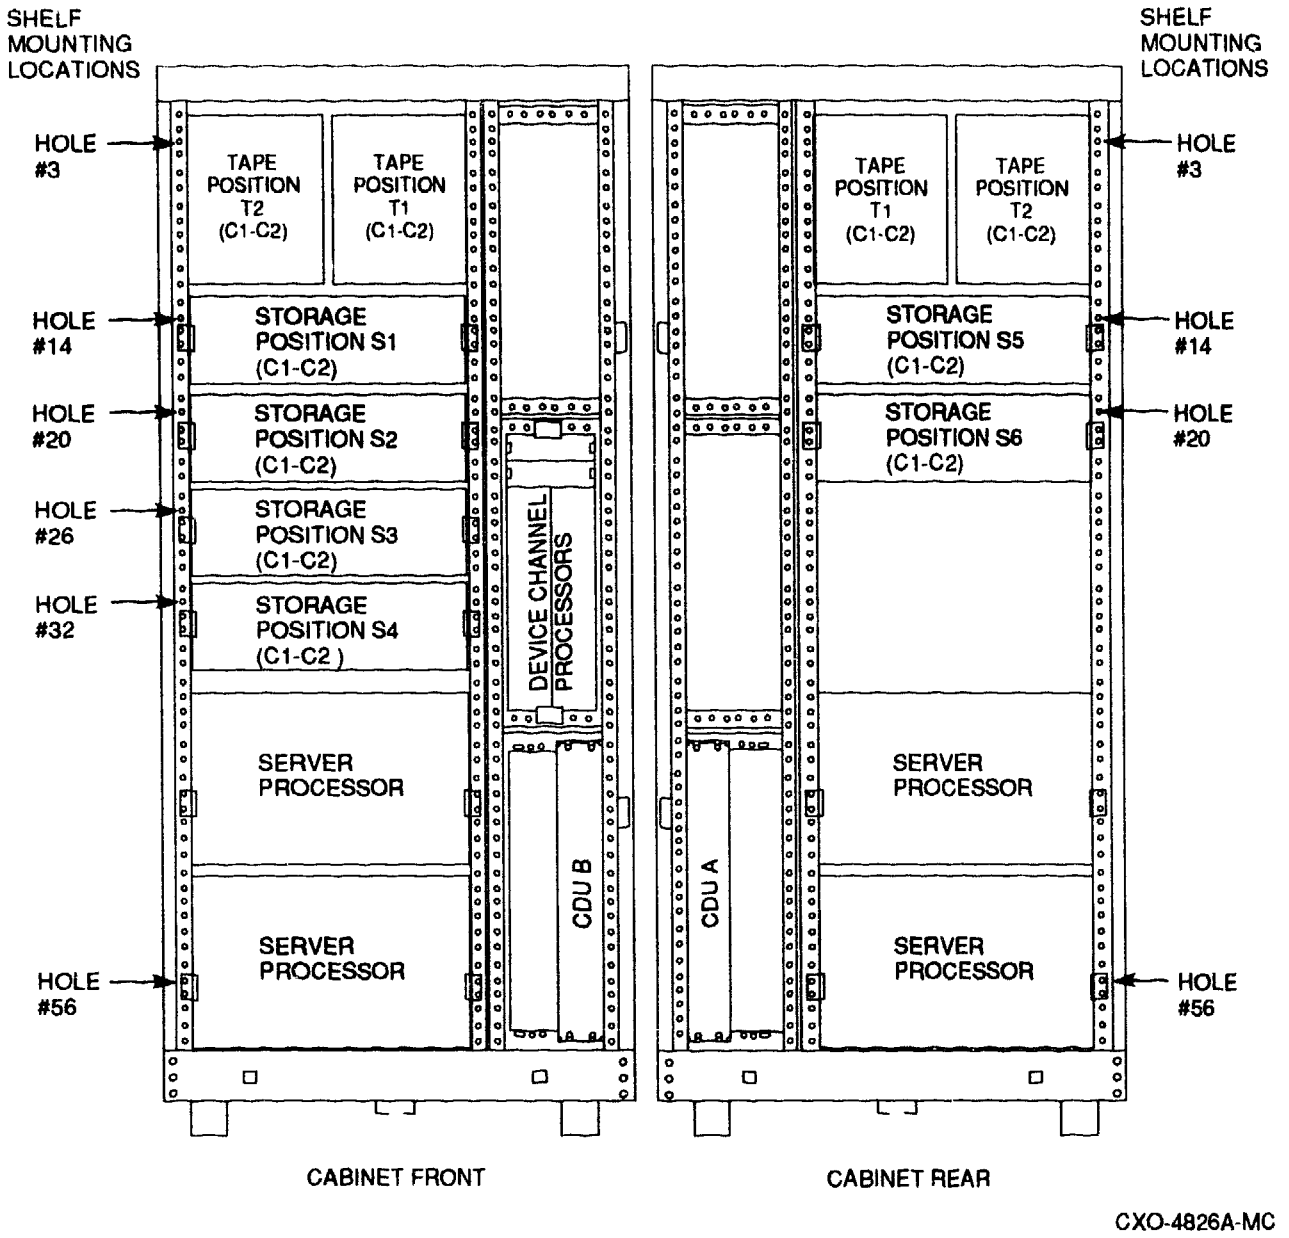


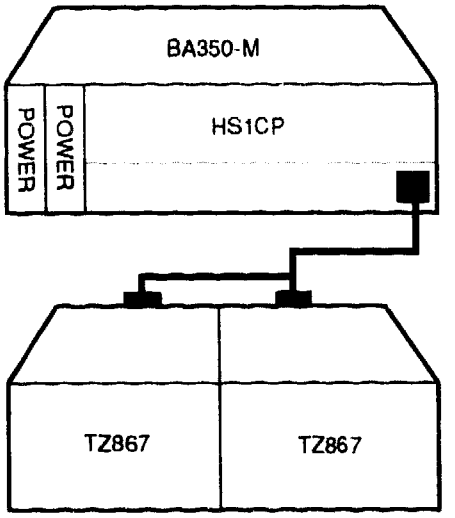
Figure D–5 HS121 Storage Configuration with Two TZ8*xx*–Series Tape Drives Installed



- Half-rack/full-depth devices, for example all TZ867 tapes, cannot share controller ports with devices that have a different form factor, and cannot be connected as an extension from a BA350–S shelf. A maximum of two such devices can be configured per controller port, and those devices must be physically adjacent to each other at the top of a cabinet.

  Figure D–6 shows two adjacent tape drives attached to a single port of the controller shelf.

Figure D–6 Adjacent Devices on a Single Port



- Two device shelves jumpered together must be connected to the first controller shelf with a 1.0 meter cable. Using a 2.0 meter cable does *not* permit shelf jumpering. (Required cable length varies depending on cabinet type, device shelf position, and controller shelf position.)

D.4.2 Device Placement

The following sections describe recommended device configurations for 3.5-inch and 5.25-inch SBBs.

- Disk and tape drives cannot be intermixed on the same controller port.

Note

You can intermix disk SBBs and tape SBBs on the same controller port, provided all other configuration rules are followed.

- Single (or paired) TZ8x7 devices must be connected with a 0.2 meter (8-inch) SCSI–1-to-StorageWorks transition cable, then to a 2.0 meter SCSI–2 cable that connects to one of the controller's SCSI–2 ports.

Note

Redundant power or a dual-redundant configuration is not supported when using 42 storage devices. This is not a Digital recommended configuration.

- Up to 42 storage devices can be attached, using seven 3.5-inch SBBs in each of six BA350–S shelves, attached to HS1CPs with six controller ports.
- Maximum number of storage shelves is four horizontal BA350–S storage shelves if four tape loaders are present (eight if only one or two TZ8x7 tape loaders are present).

- Vertical shelves are not used for storage shelves, because some devices require horizontal alignment. If desired, vertical shelf locations can be used for most disk drives. Refer to device-specific documentation for requirements.

D.4.3 Shelf Arrangement

BA350-series controller and storage shelves can be arranged in any SCSI–2 legal configuration, subject to the following rules:

- No more than a single extension from one BA350–S storage shelf is permitted.
- Half-rack/full-depth devices, for example all TZ8*xx*–series tapes, must be on their own port and cannot be connected as an extension from a BA350–S shelf. Only two such devices (maximum) can be configured per controller port, and those devices must be physically adjacent to each other at the top of the cabinet.
- When using a 1.0 meter cable with an HS1CP in the lower controller shelf position (C1) in the front of the cabinet, all front-mounted shelves can be reached. The 2.0 meter cable reaches all shelves, but does not permit shelf jumpering. (Refer to Figure D–2.)

D.4.4 Typical and Recommended 3.5-Inch and 5.25-Inch SBB Configurations

The following sections describe recommended device configurations for 3.5-inch and 5.25-inch SBBs.

Note

Intermixing disk SBBs and tape SBBs on the same controller port is permitted, provided all other configuration rules in the chapter also are obeyed.

Table Conventions

The following describes the designations used in the following sections. These designations help you to determine the possible devices in each shelf and the possible number of devices in similarly configured shelves.

(n)mxoT

or

(n)mxoJ

where:

n is the number (in parentheses) of storage shelves of this type.
m is the number of SCSI–2 connections to a storage shelf.
x is the symbol used for multiply
o is the number of devices on each SCSI–2 connection.
T indicates the storage shelf is terminated.
J indicates the storage shelf is jumpered.

According to the formula:

m * *o* is the number of possible devices in each shelf.

n * *m* * *o* is the possible number of devices in similarly configured shelves.

D.4.4.1 3.5-Inch SBB Restrictions

There are no restrictions for adding 3.5-inch SBBs to a configuration. Refer to your product-specific SPD and release notes for a list of specific supported device types.

Table D–1 lists some recommended configurations for 3.5-inch SBBs in 6-port configurations.

Table D–1 3.5-Inch SBB Configurations, 6-Port Controllers

Number of Devices	Number of BA350–S Shelves [1]	Configure as [2]	Available as 3.5-Inch SBBs [3]	Ports Used
1-2	1	(1)2x3T	5-4	1-2
3-4	2	(2)2x3T	9-8	3-4
5-18	3	(3)2x3T	13-0	5-6
19-24	4	(2)2x3T	5-0	6
		(2)1x6T		
25-30	5	(1)2x3T	5-0	6
		(4)1x6T		
31-36	6	(6)1x6T	5-0	6
37-42 [3]	6	(6)1x7T	5-0	6

[1]Consult *StorageWorks Solutions Shelf and SBB User's Guide* for BA350–S shelf information.

[2]Each BA350–S shelf's upper SCSI–2 port connector is cabled to a controller port. The lower SCSI–2 port connector is attached to a controller port for 2x3T configurations and is unused for a 1x6T or 1x7T.

[3]Available for future expansion.

Key for Table Conventions

2x3T refers to two (split) SCSI–2 connections, separately terminated in the shelf. The devices appear as IDs 0, 2, 4, and 1, 3, 5.

1x6T refers to a single path SCSI–2 connection terminated in the shelf. The devices appear as IDs 0 through 5.

1x7T refers to a single path SCSI–2 connection terminated in the shelf. The devices appear as IDs 0 through 6.

Parentheses ( ) around a number indicates the number of storage shelves.

T indicates that the shelf is terminated.

D.4.4.2 5.25-Inch SBB Restrictions

The following restrictions apply when using 5.25-inch SBBs in your configuration. Refer to your model-specific SPD and release notes for a list of specific supported device types.

- A maximum of two 5.25-inch SBBs are allowed per port (in a single shelf), or four 5.25-inch SBBs per port (in adjacent jumpered shelves).

  No more than four 5.25-inch SBBs are allowed on a single port (that would take three shelves, which cannot be configured within SCSI–2 cable limits).

- Intermixing 5.25-inch and 3.5-inch SBBs is permitted using up to six devices per port (maximum of two shelves), with no more than three 5.25-inch SBBs.

  You can use two 5.25-inch SBBs and four 3.5-inch SBBs in two BA350–S shelves, or one 5.25-inch SBB and four 3.5-inch SBBs in one BA350–S shelf.

- When using jumpered shelves, only five jumpered-pair shelves (for a total of 10 shelves) can be used within each FDDI Server cabinet. The sixth port is left unused. Alternately, four jumpered ports permit two single-shelf connections on the remaining two controller ports, which is preferable.

  This is permitted only in the lower front of the cabinet from the C1 controller position. Five such ports can take up to a maximum of 10 front shelf locations, with no allowance for cable access to shelves or devices in the rear of the cabinet. (Refer to Figure D–2.)

  A more balanced configuration consists of four 5.25-inch SBBs on each of four ports and two ports each with two 5.25-inch SBBs.

- When five ports have doubled shelves for 5.25-inch SBBs (4+2), TZ8*x*7 tapes cannot be connected or even mounted in the cabinet because all or most (front) shelf locations are needed for the 5.25-inch SBBs.

D.4.4.3 5.25-Inch SBB Recommended Configurations

Table D–2 lists some 5.25-inch SBB configurations.

Table D–2 5.25-Inch SBB Configurations for the HS1CP

Number of Devices	Number of BA350–S Shelves [1]	Configure as	Available for 5.25-Inch SBBs [2]	Ports Used
1-2	1	(1)2x3T	1-0	1-2
3-4	2	(2)2x3T	1-0	3-4
5-6	3	(3)2x3T	1-0	5-6
7-8	4	(2)1x6T	1-0	6
		(2)2x3T		
9-10	5	(4)1x6T	1-0	6
		(1)2x3T		
11-12	6	(6)1x6T	1-0	6
13-14 [3]	7	(6)1x6T	1-0	6
		(1)1x6J		
15-16 [3]	8	(6)1x6T	1-0	6
		(2)1x6J		
17-18 [3]	9	(6)1x6T	1-0	6
		(3)1x6J		
19-20 [3]	10	(6)1x6T	1-0	6
		(4)1x6J		

[1]Consult the *StorageWorks Solutions Shelf and SBB User's Guide* to configure BA350–S storage shelves.

[2]Available for additional 5.25-inch device.

[3]When used with the HS1CP in the C1 position in an FDDI cabinet. (Refer to Figure D–2 for cabinet positions.)

Key for Table Conventions

Each BA350–S shelf has its upper connector cable attached to either the adjacent BA350–S shelf's lower connector (1x6J), or a controller port connector (2x3T or 1x6T).

The lower connector cable is attached to either an adjacent BA350–S shelf's upper connector (1x6J, as in the first list item), controller port connector (2x3T), or is unused (1x6T).

Parentheses ( ) around a number indicate the number of storage shelves.

T indicates the shelf is terminated.

J indicates the shelf is not terminated and jumpered to the next shelf.

D.4.4.4 Intermixing 5.25-Inch and 3.5-Inch SBBs

Use these guidelines when 5.25-inch and 3.5-inch SBBs have been intermixed:

- Treat each 5.25-inch SBB as three 3.5-inch SBBs.
- Each 5.25-inch SBB must have its SCSI–2 ID set manually using the address switch on the rear of the SBB or by setting the switch to automatic and letting the slot connector dictate the device address. (Refer to the *StorageWorks Solutions Shelf and SBB User's Guide.*)
- A 5.25-inch SBB can be located in the same shelf with three or four 3.5-inch SBBs.

D.4.4.5 Atypical Configurations

By unbalancing the number of devices per controller port, configurations can be devised with a smaller shelf count. This results in lower performance and/or availability. Table D–3 lists the minimum shelf count for various numbers of 3.5-inch SBBs in the FDDI Server cabinet for 6-port controller configurations.

Table D–3 Small Shelf Count Configurations, 6-Port Controllers

Number of Devices	Number of BA350–S Shelves [1]	Configure as	Ports Used
1-6	1	1x6T [2]	1
7-12	2	1x6T	2
13-18	3	1x6T	3
19-24	4	1x6T	4
25-30	5	1x6T	5
31-36	6	1x6T	6
37-42 [3]	6	1x7T	6

[1]Consult the *StorageWorks Solutions Shelf and SBB User's Guide* for BA350–S shelf information.

[2]T indicates that the shelf is terminated.

[3]Nonredundant controller and power configurations (not recommended).

D.5 Configuring Cabinet Power

The degree of power redundancy is determined as follows:

- No power redundancy—Single cabinet CDU with single power supply per shelf
- Partial power redundancy—Single cabinet CDU with two power supplies per shelf
- Full power redundancy—Two cabinet CDUs with two power supplies per shelf

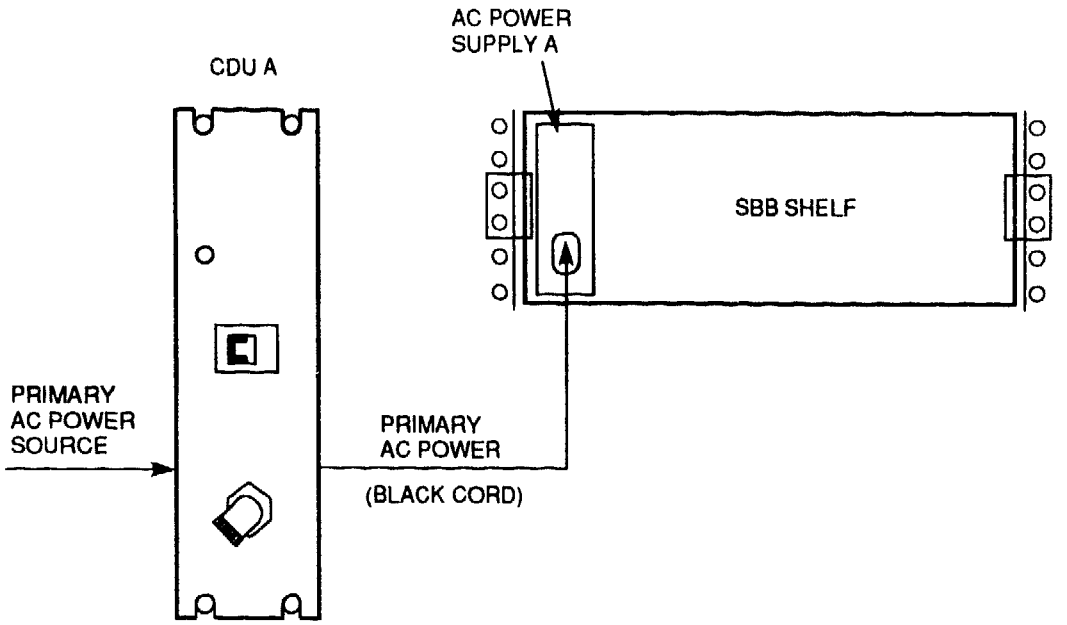
D.5.1 Configuring for No Power Redundancy

One CDU per cabinet and one power supply per StorageWorks shelf is all that is needed to supply power to all components of the FDDI Server if redundant power is not needed. This is the normal configuration for the cabinet. The disadvantage to this system is that a failure in any part of the power system results in a loss of all or part of the capabilities of your FDDI Server. See Figure D–2 for CDU locations.

FDDI Server cabinets are shipped with a single CDU and are cabled for dual shelf power supplies. The cabinet is supplied with 24 preinstalled, black ac power cables. Two of these cables are routed to each front and rear shelf position. These cables are used to supply ac power to shelves as they are added to the cabinet. The cables need to be plugged into each shelf's power supply to complete the shelves' installation. No additional materials are necessary for this configuration.

Figure D–7 shows a nonredundant power configuration. If any component in the power system fails, the StorageWorks devices mounted in the shelf become unavailable until you repair the power system failure.

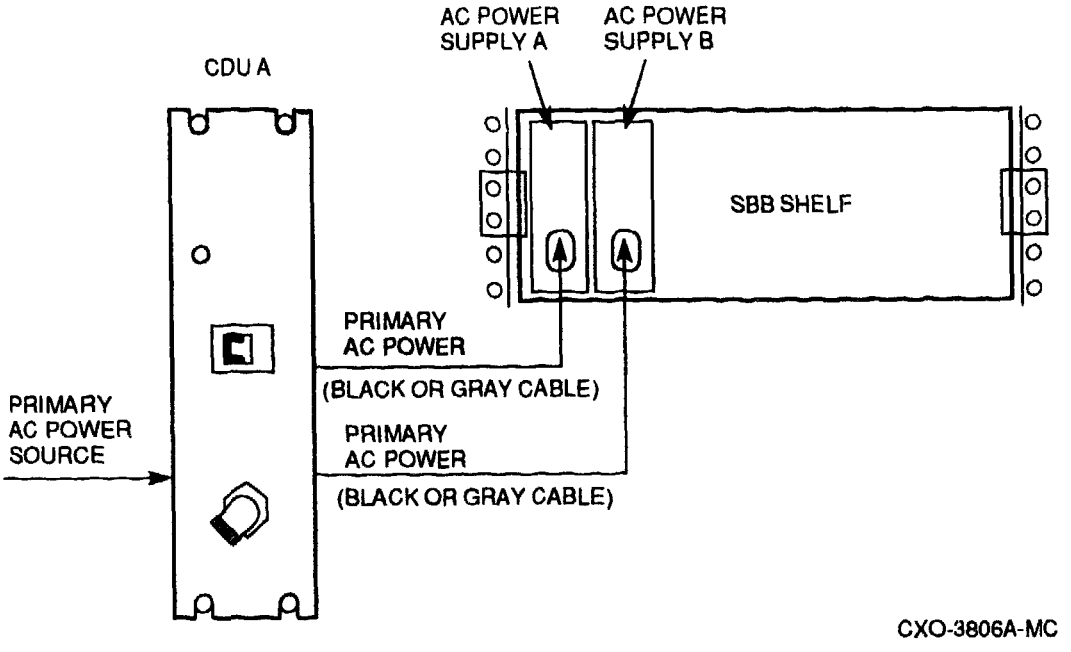
Figure D–7 Single-Cabinet Power Configuration



D.5.2 Partial Power Redundancy Configuration

If protection against individual shelf ac power supply failure is required, the system should have been configured with a power system for partial power redundancy. Each ac power supply provides backup for the other in case of power supply failure. Figure D–8 shows the power system configured with shelf power redundancy. One primary power source and CDU (A) supply ac power to two ac power supplies (A and B) in each shelf. If either shelf power supply fails, the shelf remains powered by the other supply. This configuration provides shelf power supply redundancy to the shelves. The 12 shelves of a fully configured FDDI Server are protected against shelf power supply failure using this configuration. Note, however, that a failure in the CDU or the primary power system results in a loss of the FDDI Server capabilities.

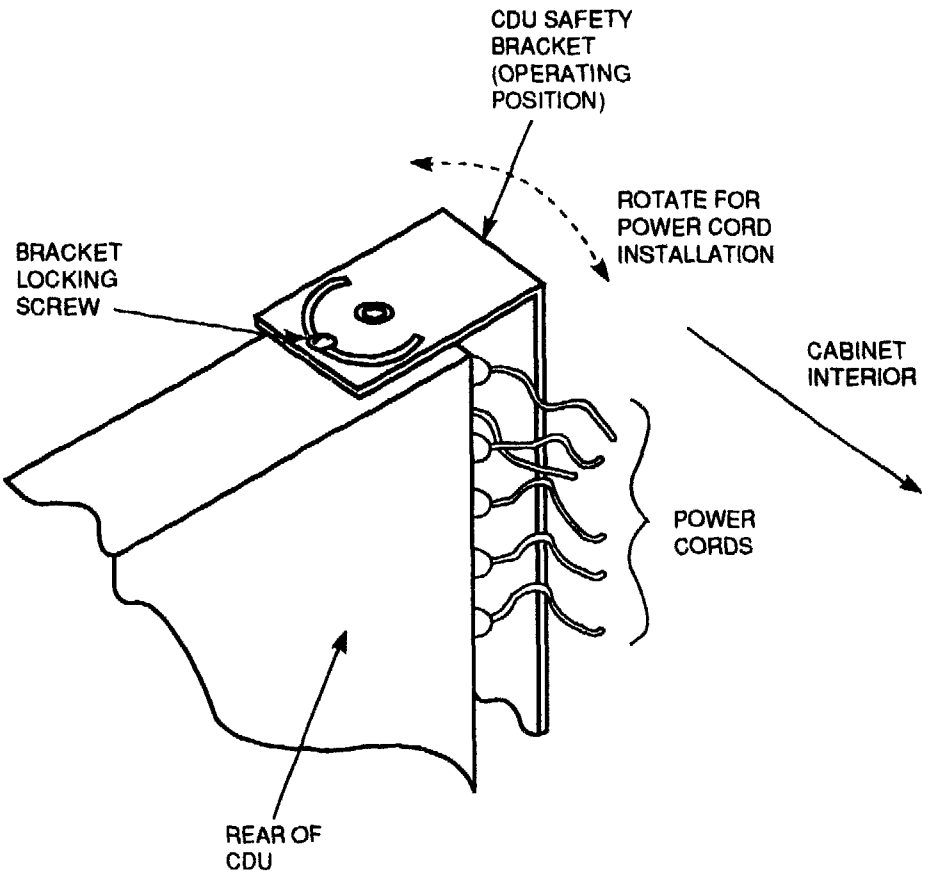
Figure D–8 Dual Shelf AC Power Supply Configuration



In the dual shelf power configuration, a second ac power cable and ac power supply is necessary for each shelf. (Refer to Figure D–8.) Shelves should be configured for dual ac power supplies as follows:

- The second ac power supply (B) should be plugged in next to the existing ac power supply (A) in each shelf.
- The appropriate end of each ac power cable should be plugged into a connector on the rear panel of the CDU.
- Make sure the safety bracket is properly secured by the bracket locking screws over the cable connectors (see Figure D–9).
- Each power cable should be secured by cable clamps or ties.

Figure D–9 Safety Bracket Rotation

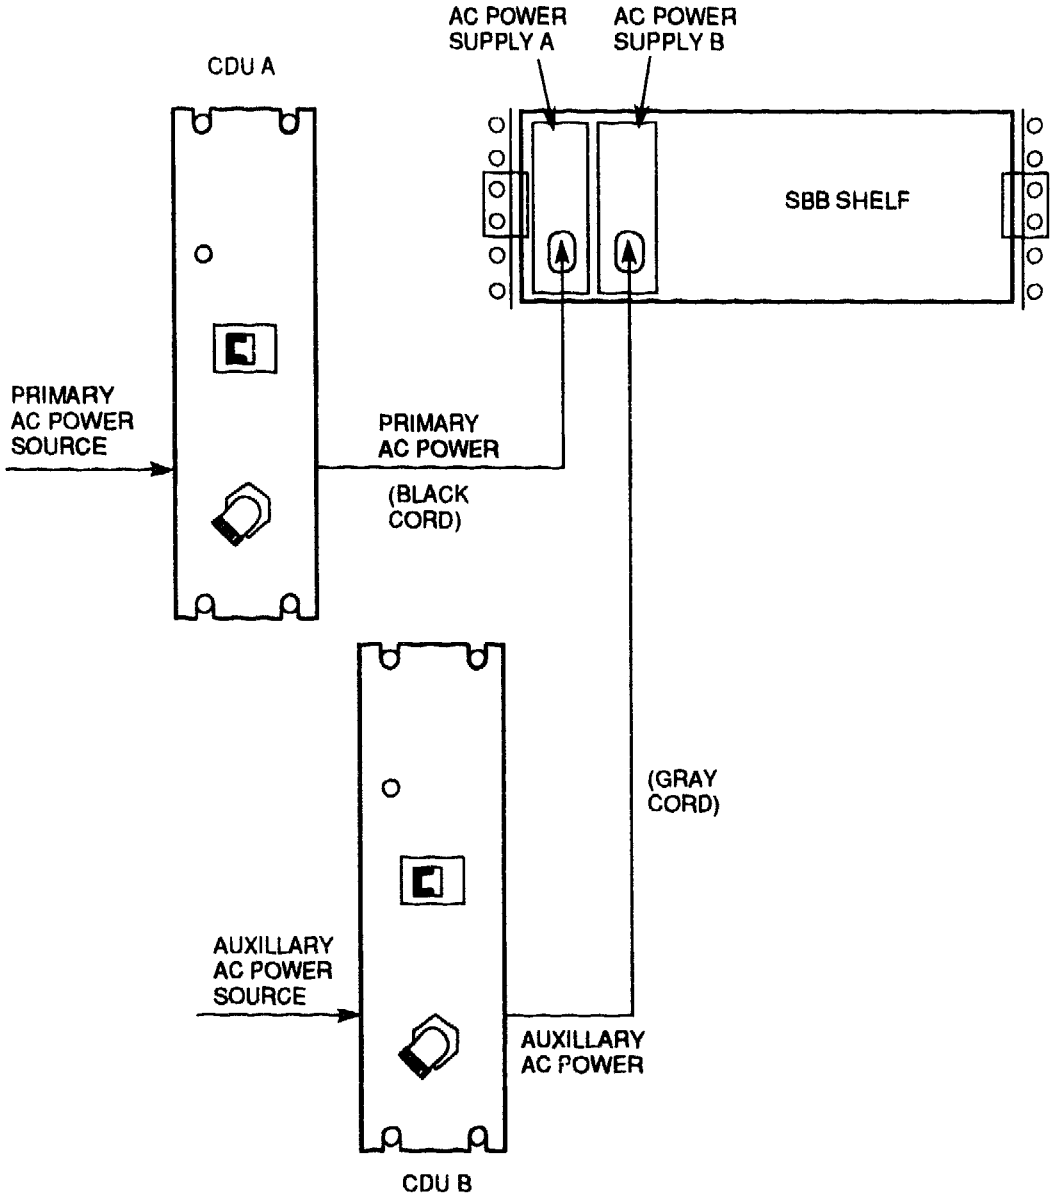


D.5.3 Full Power Redundancy Configuration

If your FDDI Server must continue to operate in all conditions except for a major loss of site power, it should be configured for full power redundancy.

The fully redundant configuration requires two ac power supplies per shelf, two CDUs per cabinet, and two primary power sources. Figure D–10 shows a full redundant power configuration. If the primary power source, CDU A, or ac power supply A fails, the shelf remains powered by ac power supply B. This configuration provides complete power redundancy to the shelves. The 12 shelves of a fully configured StorageWorks FDDI Server are powered in this manner.

Figure D–10 Dual Cabinet Power Configuration



The dual cabinet power configuration requires two CDUs for power distribution as shown in Figure D–2.

D.6 Configuring Internal Bus Cables for the HS1CP and Server Processor

The following sections describe how the internal bus cables between the HS1CP and the server processor are configured in various FDDI Servers.

Following are special guidelines for the internal bus cables going to and from the HS1CP:

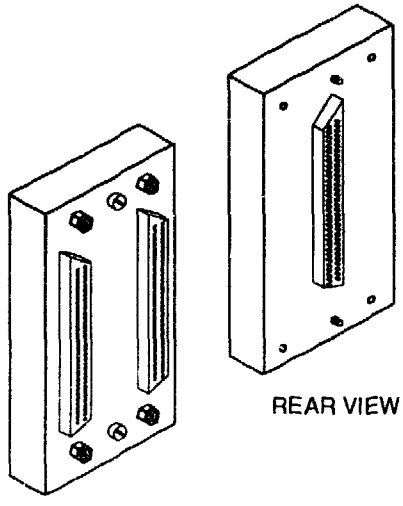
- Cable length between nodes/members on the internal bus must be no greater than 16 feet (4.9 meters).
- Total cable length (end-to-end) on one internal bus must be no greater than 60 feet (18.3 meters).

CAUTION

Currently, component damage can result if internal bus cables are connected or disconnected with power applied, *unless* the mating guide is installed around the outside edge of your HS1CP port connector. HS1CPs are shipped from the factory with the mating guide and the trilink connector block premounted on the HS1CP port connector.

Be aware that the trilink connector block (Figure D–11) is the interface between the HS1CP and other internal bus cable connections. It is the *trilink* that you can disconnect and connect safely when you have a mating guide installed. In a power-on situation, you must work around any internal bus cable or terminator connections to the trilink *without* disconnecting them.

Figure D–11 Trilink Connector



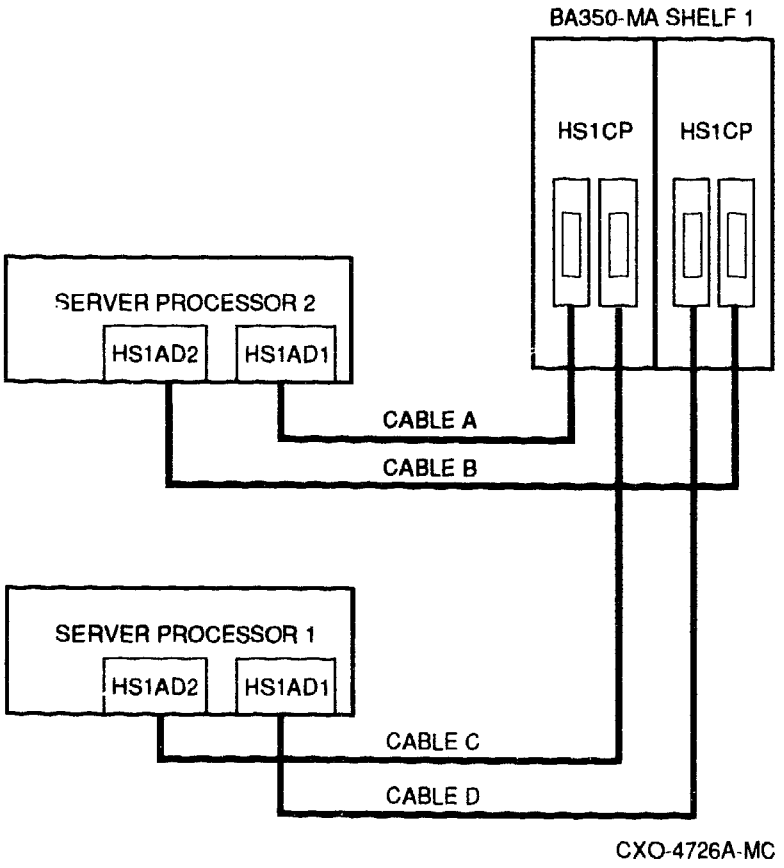
D.6.1 Configuring HS111/211 Internal Bus Cable and Terminator

For model HS111/211 FDDI Servers, the internal bus cable is connected between the left side of the trilink connector on the HS1CP device channel processor and the HS1AD device bus adapter in the server processor. See Cable D in Figure D–12. A terminator is installed in the right side of the trilink connector.

D.6.2 Configuring HS121/221 Internal Bus Cables

Figure D–12 shows the internal bus cable connections on a model HS121/221 FDDI Server.

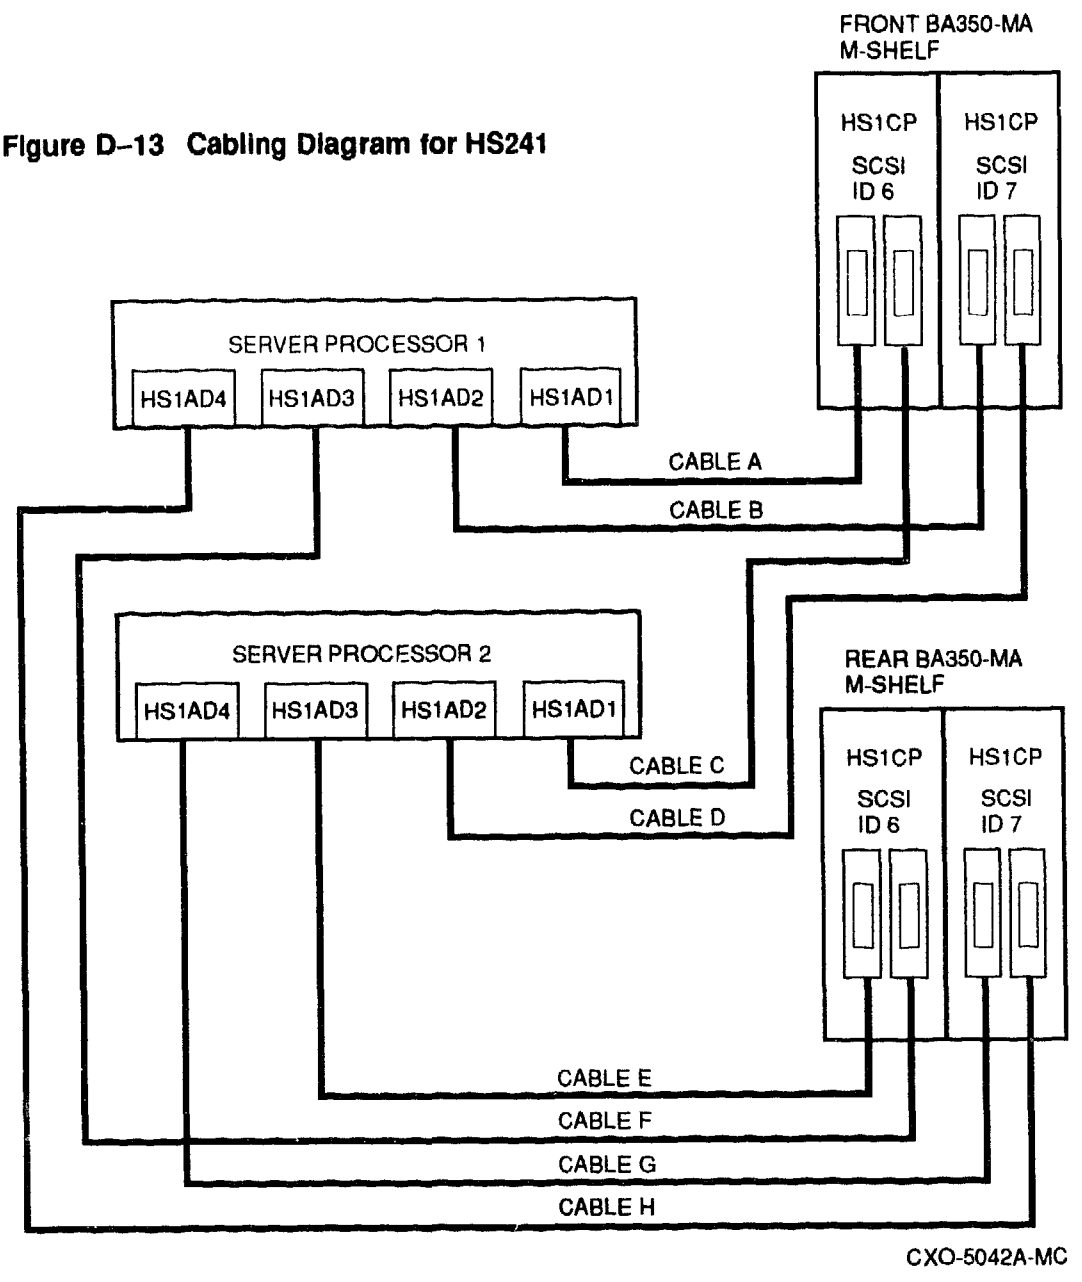
Figure D–12 Cabling Diagram for HS121/221



D.6.3 Configuring HS241 Internal Bus Cables

Figure D–13 shows the internal bus cable connections on a model HS241 FDDI Server.

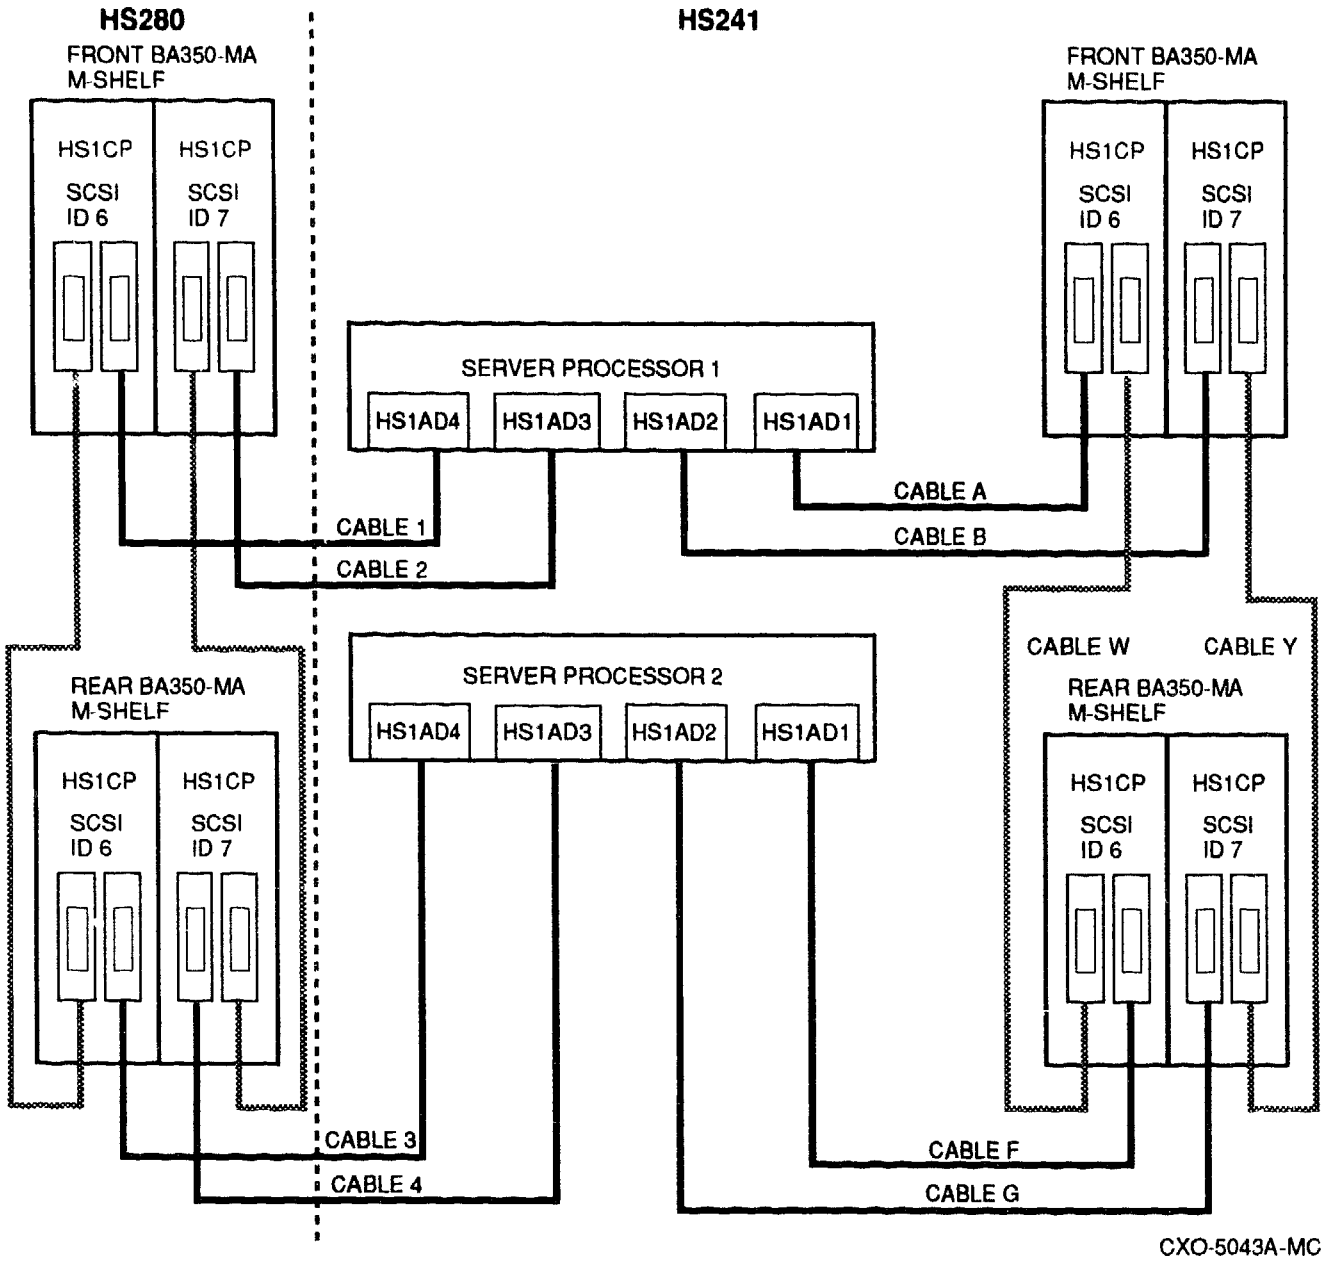
Figure D–13 Cabling Diagram for HS241



D.6.4 Configuring HS241/280 Internal Bus Cables

Figure D–14 shows the internal bus cable connections on a model HS241 FDDI Server with HS280 expansion cabinet.

Figure D–14 Cabling Diagram for HS241/280



D.7 Configuring the HS1CP

The following support concerns apply to configuring HS1CPs:

HS1CP Hardware Support

HS1CP operating firmware supports the following revisions or higher for the HS1CP and associated hardware:

- HS1CP Device Channel Processor, hardware revision A05
- BA350–M shelf, hardware revision A
- BA350–Sx shelf, hardware revision B
- BA35x–HA power unit, hardware revision H

HS1CP Operating System Support

The HS1CP should be installed only in a FDDI Server operating its embedded version of the OpenVMS Alpha™ Operating System.

HS1CP Device Support

HS1CP operating firmware supports the devices listed in Table D–4 at the indicated hardware and microcode levels or higher.

The following device installation rules apply:

- In an HS111 configuration, the single HS1CP occupies slot 7 (*right* slot). This slot is assigned SCSI ID 7. The cache memory module associated with the HS1CP is in the slot to the left and adjacent to the HS1CP.
- In an HS121 configuration, the write-back cache module and HS1CP are installed into the right side of the of the BA350–M shelf.
- In an HS121 configuration, the second HS1CP occupies slot 6 (*left* slot). This slot is assigned SCSI ID 6. The cache memory module associated with the HS1CP is in the slot to the left and adjacent to the HS1CP.
- In an HS121 configuration, a second write-back cache module and HS1CP are installed into the left side of the of the BA350–M shelf.

Table D–4 HS Operating Firmware Supported Devices

Device	Type	Capacity (gigabytes)	Microcode Version	Hardware Revision
RZ26–VA	Disk drive	1.05	T392	D02
RZ26L–VA	Disk drive	1.05	440C	A01
RZ28–VA	Disk drive	2.1	435E	B01
RZ28B–VA	Disk drive	2.1	0003	A01
RZ74–VA	Disk drive	3.57	T427B	B07
TKZ60–EC [1]	Tape drive	0.440	111	B01
TL820 [1,2] (Revision A01)	Tape library	4,000 [3]	1d3M	L1
TLZ06–VA	Tape drive	4 [3]	0491	A04
TLZ6L–VA	Tape loader	16 [3]	0491	A01
TSZ07–AA [1]	Tape drive	0.140	0309	A01
TZ86–VA	Tape drive	6	430B	A02
TZ87–VA	Tape drive	20 [3]	930A	A01
TZ87N–VA [4]	Tape drive	20 [3]	930A	A01
TZ87–TC [1,2]	Tape drive	20[3]	9514	B02
TZ867–AE/AF [1]	Tape drive	42	430B	A01
TZ875–NT [1,4]	Tape drive	100 [3]	930A	A01
TZ875–TA [1]	Tape drive	100 [3]	930A	A01
TZ877–AE/AF [1]	Tape drive	140 [3]	930A	A01
RRD42–VB/VU	CD-ROM reader	0.6	1.1a	A01
RRD43–AA	CD-ROM reader	0.6	0064	A02
RRD44–AA	CD-ROM reader	0.6	3493	A02
RW524–ZA [1]	Optical disk library	9.50/19.20	3404	A01
EZ51R–VA [5]	Solid state disk	0.1069	V096	D01
EZ54R–VA [5]	Solid state disk	0.4277	V096	A01
EZ58R–VA [5]	Solid state disk	0.8553	V096	A01

[1]Requires 0.2 meter SCSI–1-to-SCSI–2 transition cable (Digital internal part number 17–03831–01).

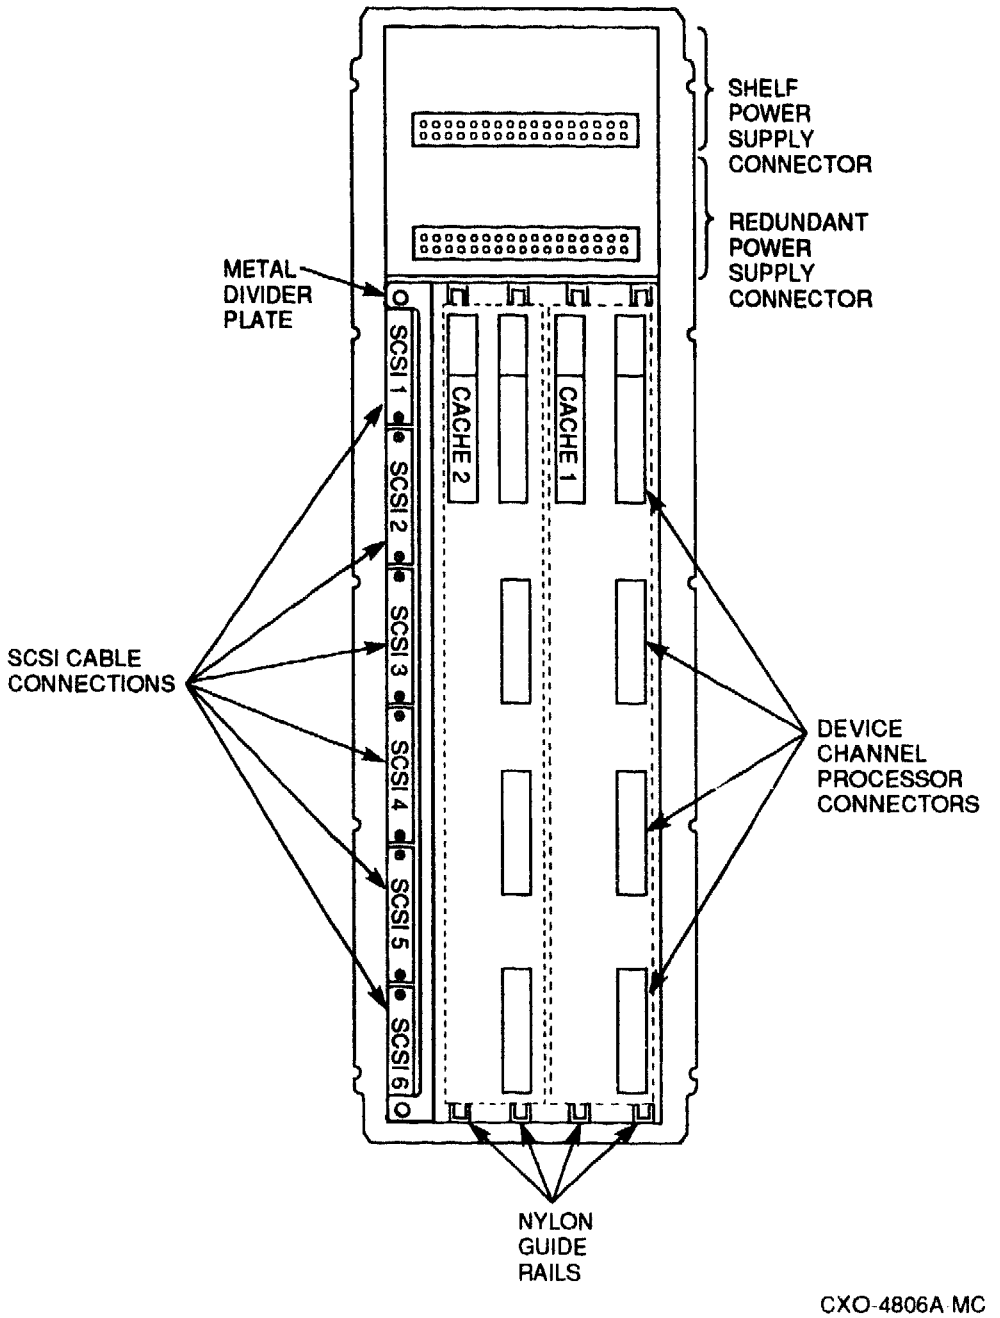
[2]Requires DWZZA–AA 8–bit-to-16–bit SCSI adapter (also supported).

[3]Values represent compressed data.

[4]Cannot read TK50™, TK70, or TZ30 format tapes.

[5]Do not warm swap this device. Use a cold swap method for device replacement.

Figure D–15 BA350–M Shelf Layout (Front View)



D.7.1 HS1CP Configuration Guidelines and Considerations

The following sections describe specifics for configuring the HS1CP:

Nonredundant HS1CP Configurations

The following considerations apply to nonredundant configurations:

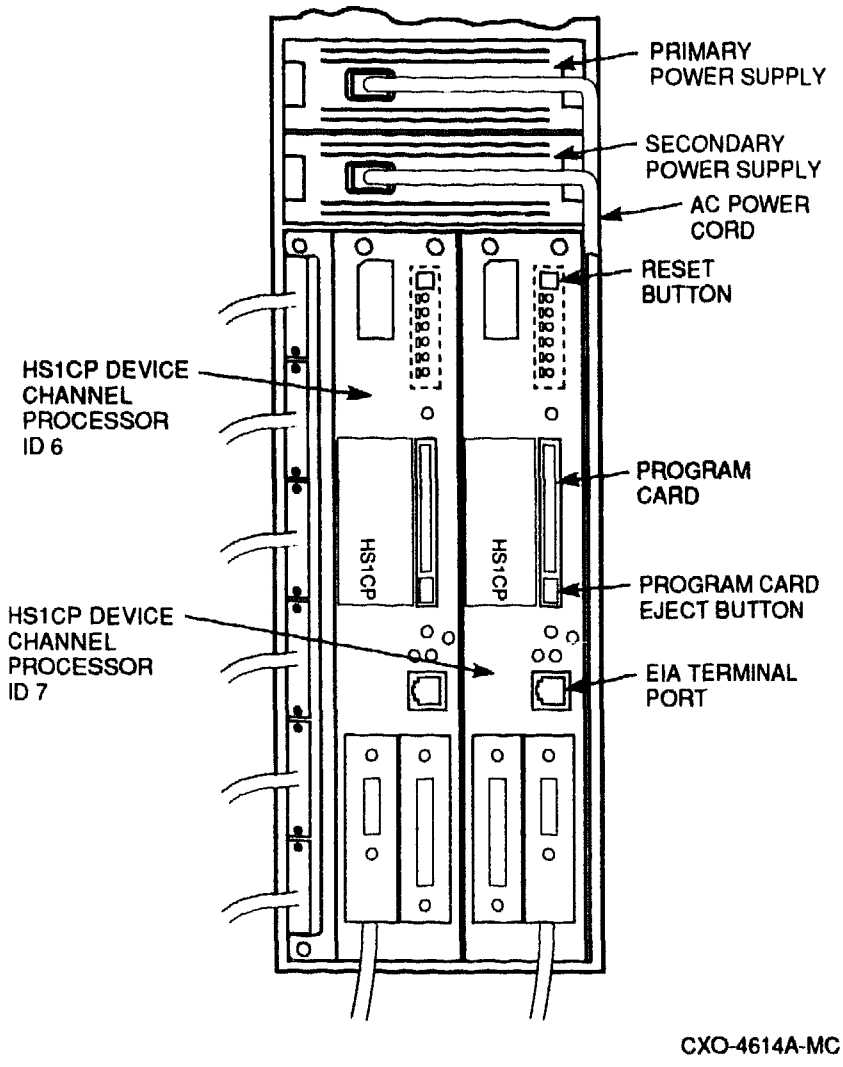
- A nonredundant HS1CP must be installed in the controller slot furthest from the BA350–M controller shelf's SCSI connectors. This slot is SCSI ID 7. By using SCSI ID 7, SCSI ID 6 (the other controller slot) is available as an additional ID on the storage shelf.
- Only two HS1CPs per internal bus are allowed.
- The maximum recommended HS1CP subsystem configuration using the BA350-series storage shelves is six devices per port (36 devices). This allows for the addition of another HS1CP and additional power supplies in the storage shelves.

Dual-Redundant HS1CP Configurations

The following considerations apply to dual-redundant HS1CP configurations:

- Dual-redundant controllers must be located in the same shelf and connected to each other through the shelf backplane. Both controllers have access to all the devices on each other's ports. This setup increases availability and provides for failover when one controller in the pair fails. (The surviving controller takes over service to all devices.)
- Dual-redundant configurations follow the same guidelines as nonredundant configurations.
- Both HS1CPs' firmware versions must be identical. If there is a mismatch, neither HS1CP can access any devices. Both controllers' cache modules must be the same type (read or write-back), have the same number of megabytes, and have the same firmware versions.
- Dual-redundant HS1CPs must be on the same internal bus.

Figure D–16 HS1CP Installation



Optimal Performance Configurations

For optimal performance, configure to the following guidelines:

- Balance the number of devices on each port of the HS1CP. This permits parallel activity on the HS1CP's available ports to the attached devices.
- Evenly distribute higher performance devices across separate ports so that higher and lower performance devices are intermixed on the same port. (For example, put multiple solid state disks on separate ports.) This intermixing of higher and lower performance devices on the same port benefits overall performance. Use the guidelines in Table D–5.
- Limit the number of devices per controller port to three in dual-redundant configurations. In doing so, both controllers access three devices per *each other's* port, maintaining six SCSI–2 devices combined total.
- Maximize the amount of cache memory per controller with the 32 MB cache module option.

- Limit the number of devices per controller port to two in nonredundant configurations and four in dual-redundant configurations.

Table D–5 High-Performance Devices per Port

Number of Relatively Higher Performance Devices	Maximum Number of Devices Configured per Port
1 - 3	1
4 - 6	1
7 - 9	2
10 - 12	2
13 - 15	3
16 - 18	3

Highest Performance

To obtain the highest performance possible, use a dual-redundant configuration and balance the number of devices across the two HS1CPs. This is done through your operating system by setting the order of how the devices were mounted or sequenced, and by having set the path definitions.

The result is a requirement for approximately half of the number of devices normally accessed through each HS1CP. Should one HS1CP fail, its devices automatically fail over to the other HS1CP.

Optimal Availability Configurations

For optimal availability, the following configurations should be observed:

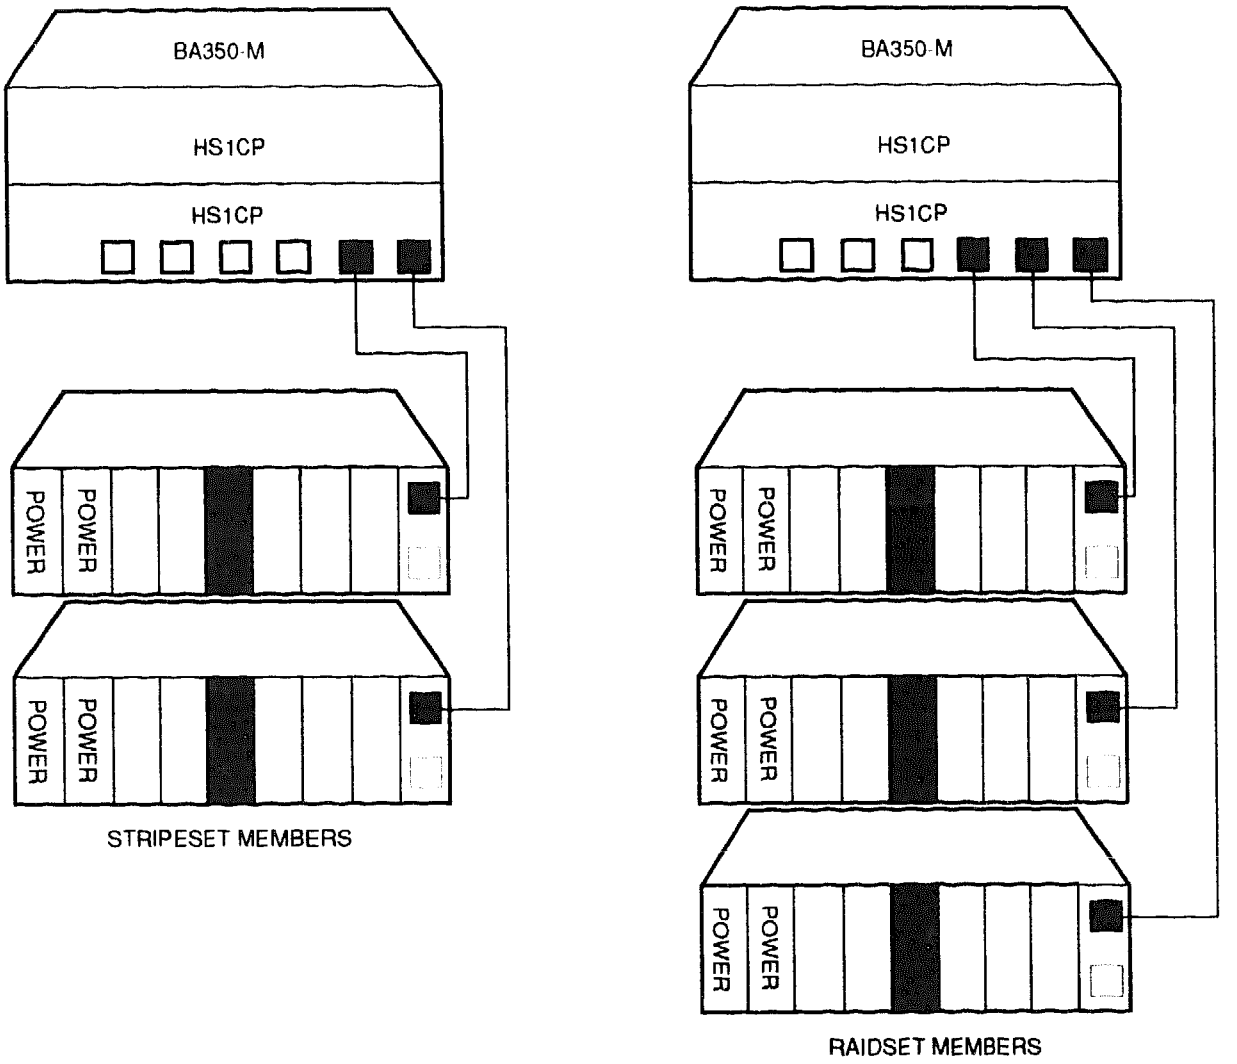
- Use dual-redundant controller and redundant power supplies in all shelves.
- Place storageset members on different controller ports and different storage shelves.
- Use predesignated spares on separate controller ports and storage shelves.
- Place storageset members on separate HS1CPs when using host-based RAID implementations (for example, shadowing).

Highest Availability

- For host-based RAID implementations, split the normal access path between HS1CPs.
- Use redundant power supplies in all shelves.

Figure D–17 shows examples of optimal configurations for RAIDset members and designated spares on separate controller ports.

Figure D–17 Optimal Availability Configuration Example



D.7.2 Preconfigured and CTO Subsystems

Preconfigured and configure to order (CTO) subsystems have all subsystem components preinstalled (HS1CPs, SCSI–2 device cables, storage shelves, devices, power supplies, and so forth).

Therefore, to check the configuration of the component, you need to do only the following:

1. Visually inspect all subsystem components to determine the following:
   - All cables are seated properly.
   - All SBBs are seated properly.
   - All modules are seated properly.
   - All shelf fans are seated properly.
   - All program cards are loaded properly.

2. Plug a maintenance terminal's EIA cable into the EIA–423 terminal port on the front bezel of your HS1CP (for defining your subsystems's initial configuration parameters). See Section D.3 for maintenance terminal setup information.
3. Turn on the power to your maintenance terminal.
4. Turn on the circuit breakers (CB1) on your subsystem cabinet's cable distribution units to the ON (|) position. The HS1CP and devices begin their normal initialization sequence.
5. From your maintenance terminal, check to see which configuration parameters have been preset by entering the SHOW THIS_CONTROLLER command. A hardcopy printout of subsystem configuration parameter settings may also be available. Refer to Section D.7.3 for additional information.
6. Set the configuration parameters *before* connecting your host port cables. Set the parameters from a maintenance terminal connected to your HS1CP.
7. Connect the appropriate host port cables and/or terminators to the supplied trilink connector on the front bezel of your HS1CP.

Note

Do not connect the HS1CP end of any internal host port cables to your controller unless the HS1CP's initial parameters are set and the host ports are enabled.

D.7.3 Preset Configuration Parameters

The HS1CP's minimum initial parameters are defined at the factory. Refer to the most recent configuration printout. You can verify the preset parameters by entering the following commands at the HS1CP> prompt:

```
HS1CP> SHOW THIS_CONTROLLER
```

or

```
HS1CP> SHOW OTHER_CONTROLLER
HS1CP> SHOW DEVICES
```

If initial parameters are not set, set parameters for your nonredundant or dual-redundant configuration (such as setting the controller's ID, turning on host port paths, and so forth).

After setting your initial parameters and defining device configurations, make sure the current list is printed and kept available to assist in servicing the subsystem in the future. Make a new printout each time you change your configuration parameters or add or delete units or storagesets.

Note

Use the SHOW commands to display your device and storageset configurations (SHOW DEVICES, SHOW UNITS, SHOW STRIPESETS, SHOW RAIDSETS, and so forth).

D.8 Configuring the Server Processor

This section describes various configurations of the FDDI Server server processor in the SW800 cabinet. With two server processors installed in the HS121 configuration, the layout and function of the internal components is the same. When differences in configuration of the SCSI system disks and CD–ROM drives occur, the differences are noted.

The standard server processor contains a single power supply, one HS1AD device bus adapter, five SIMMs, a DEFEA module, and other modules and components required for operation. A fully configured server processor can contain two power supplies, up to four banks of memory each with 128 megabytes, and four KFESB modules. Only one Alpha processor daughterboard can be supported in each server processor. Figure D–18 shows the layout of the front section of the server processor. Figure D–19 shows the layout of the back section of the server processor.

Figure D–18 Layout of the Front Section of the Server Processor

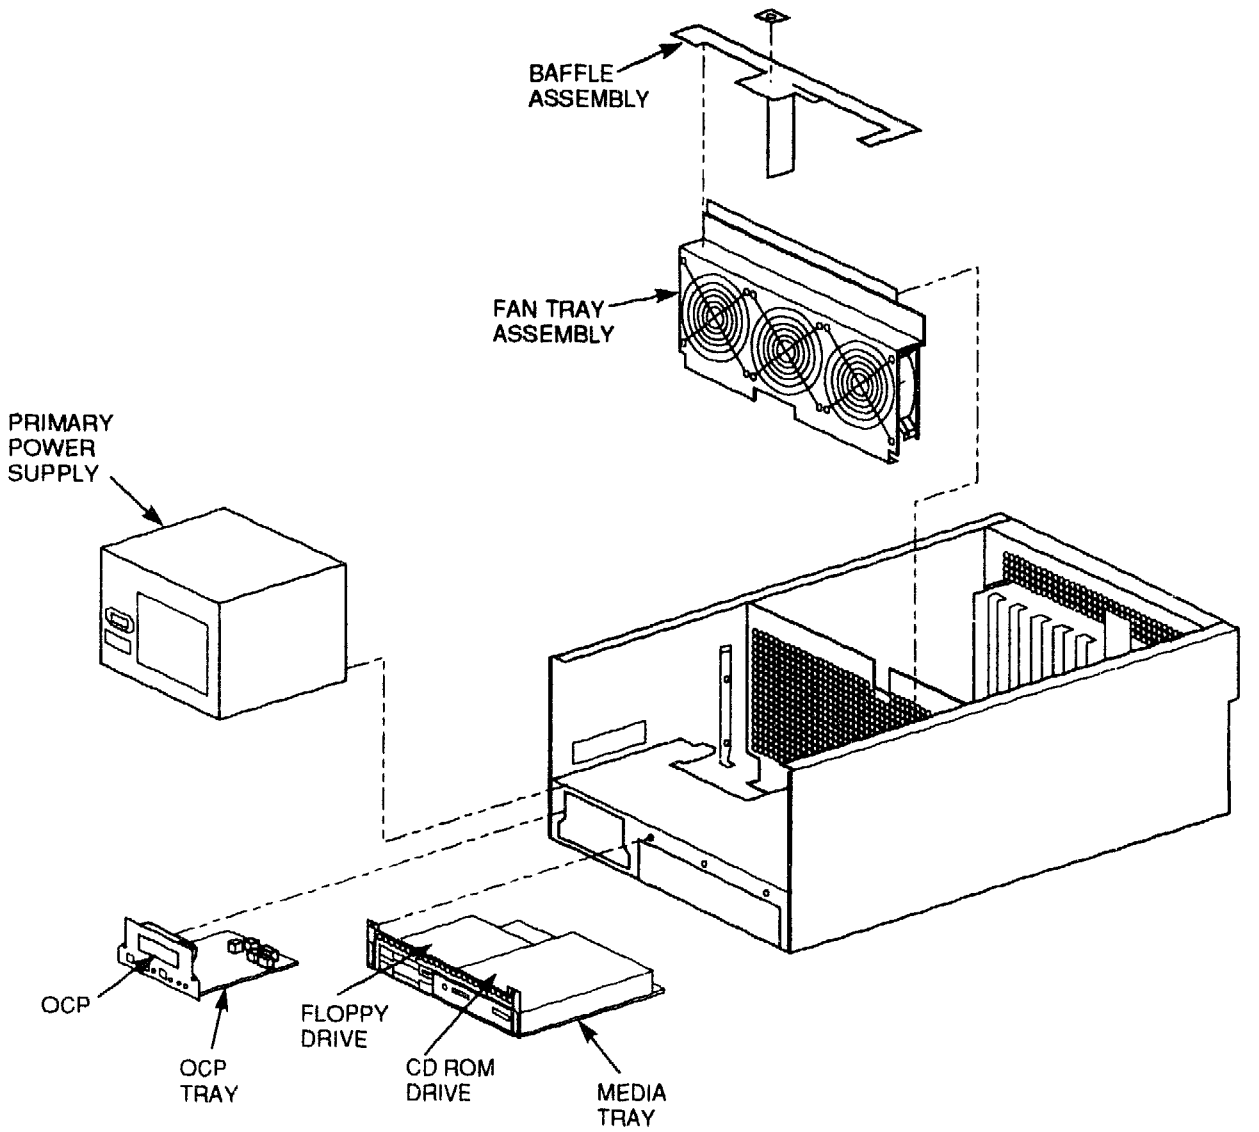
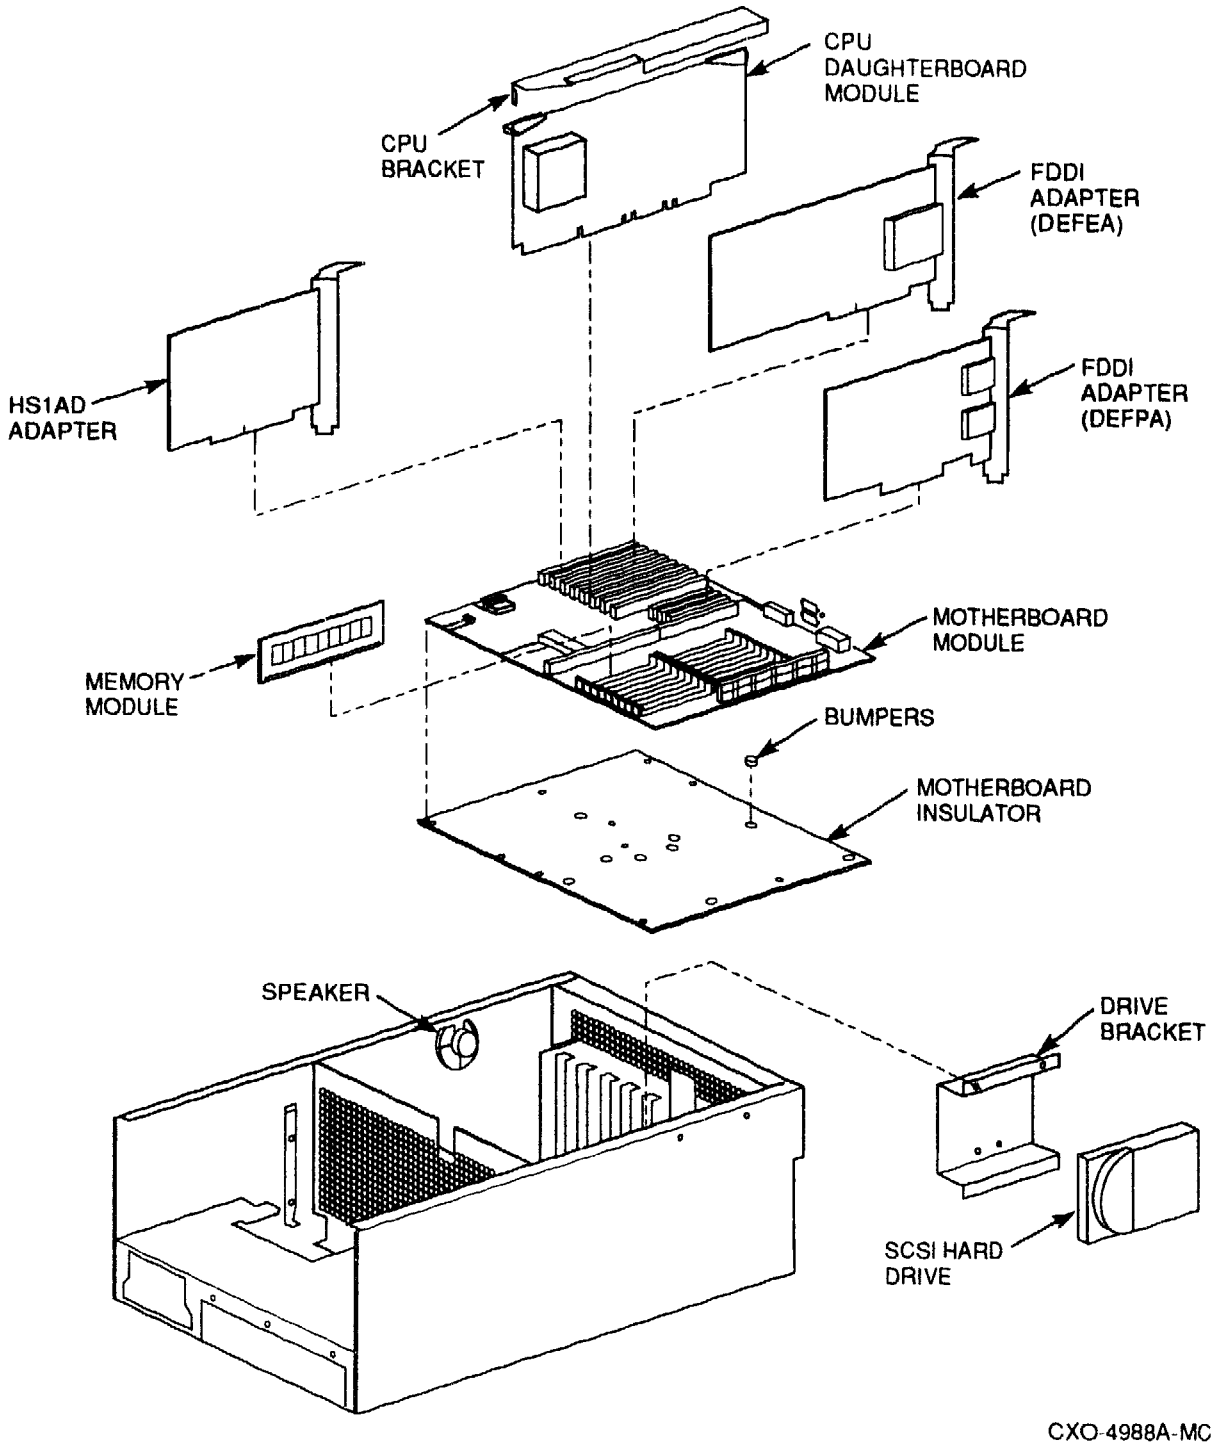


Figure D–19 Layout of the Back Section of the Server Processor



D.8.1 The HS1AD Adapter

If the system is configured as an HS111, the single server processor contains a single HS1AD adapter.

When the system is configured as an HS121, each server processor contains two HS1ADs. Refer to Chapter 8 for information on how to access the internal SW800 components. The number of HS1ADs in each server can be determined by looking at the back of the server processor. Figure D–20 illustrates two HS1ADs installed in the right-most EISA/ISA option slots.

Figure D–20 Server Processor–Rear View

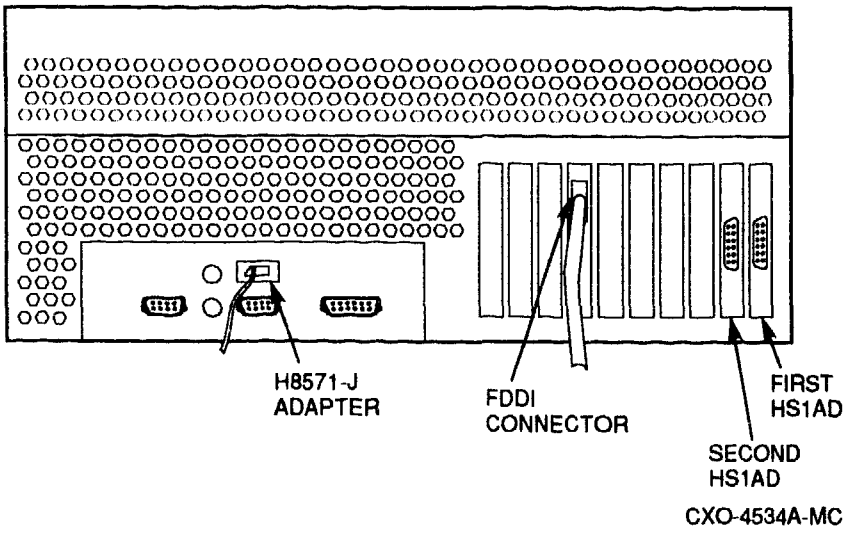
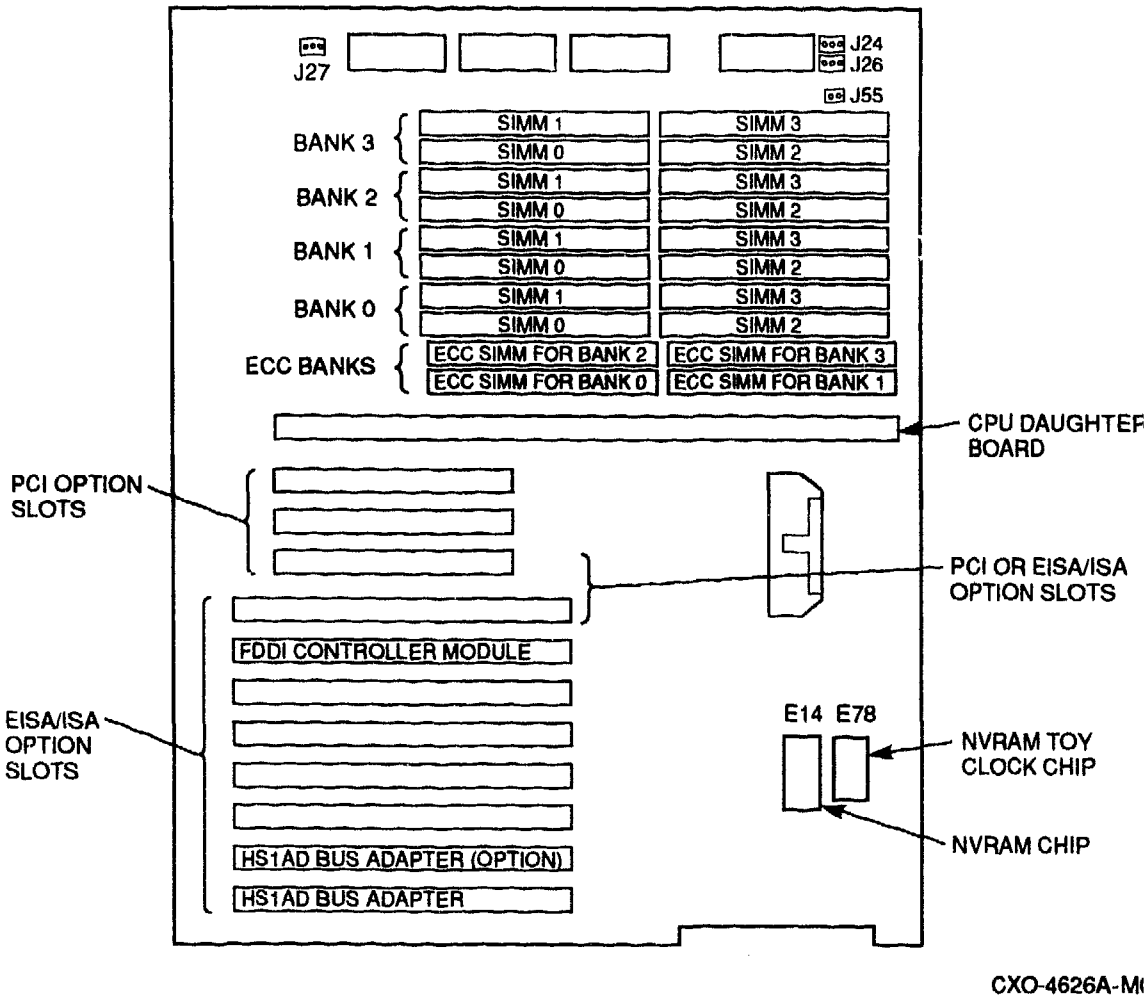


Figure D–21 shows the location for the device bus adapters in the server processor.

Note

It is possible to use any EISA slot from 2 through 4 when an additional HS1AD is installed. However, Digital recommends that the second and subsequent HS1AD be installed in numerical order (for example, the first HS1AD should be in slot 1; the second in slot 2, next to the first HS1AD, and so on). In Figure D–21, the slot position of the second HS1AD is "HS1AD Bus Adapter (option)."

Figure D–21 Device Bus Adapter Location



If access is required to either remove, replace, or install an additional device bus adapter, refer to Chapter 8 or your upgrade manual for detailed information.

D.8.2 System Disk SCSI Identifier

Each device on a SCSI bus must have a unique address. Typically, the first device on the bus is assigned address 0, the second device is assigned address 1, and so on through address 7.

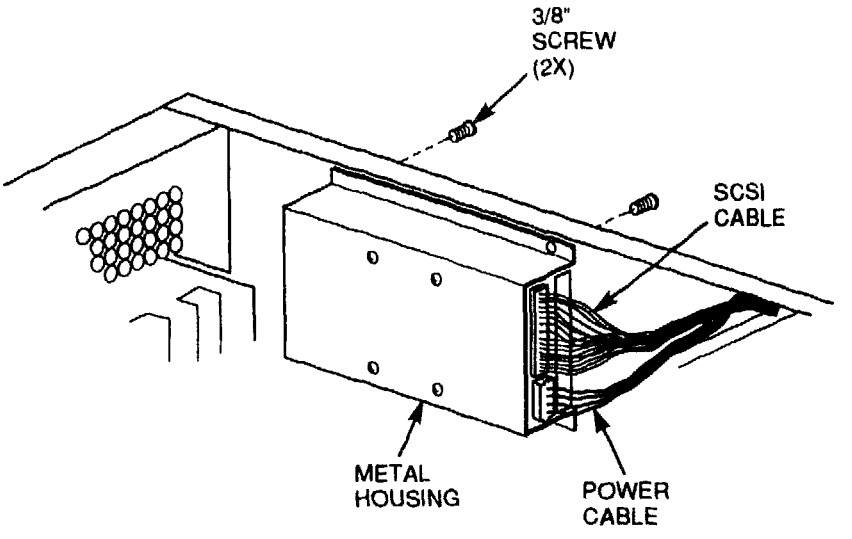
Every server processor in an HS111 is shipped from the factory with the system disk SCSI ID (identifier) set to ID 1. In an HS121, the system disk ID in the lower server processor is set to ID 1, and the system disk ID in the upper server processor is set to ID 0.

Use the following procedure to check whether the ID on SCSI system disk 1 is 1 and the second SCSI ID is 0:

1. Access the server processor and system disk using the procedures listed in Chapter 8.

2. Disconnect the SCSI cable and the power cable that are plugged into the back of the system disk (see Figure D–22).

Figure D–22 System Disk Cabling



3. Carefully remove the system disk and its drive bracket from the inside of the server processor.
4. Locate the option connector on the system disk (see Figure D–24). This option connector is a small black 20-pin connector with two white jumpers installed.

CAUTION

The first three pairs of pins on the option connector are used to set the drive SCSI ID. Other jumpers on the option connector should not be changed.

5. If checking the system disk in Server Processor 1 (the lower server), the SCSI ID should be set to 1. Refer to Figure D–23.
6. The system disk in Server Processor 1 should show device DKA0 when the SHOW DEVICE command is issued from the DCL prompt.
7. If checking the system disk in Server Processor 2, the SCSI ID should be set to 0. Refer to Figure D–24.
8. The jumper pins should have been removed from the first three pairs of SCSI ID pins. Removing this jumper sets the drive to SCSI ID 0 (see Figure D–24).

Figure D–23 System Disk Option Connector and SCSI ID 1

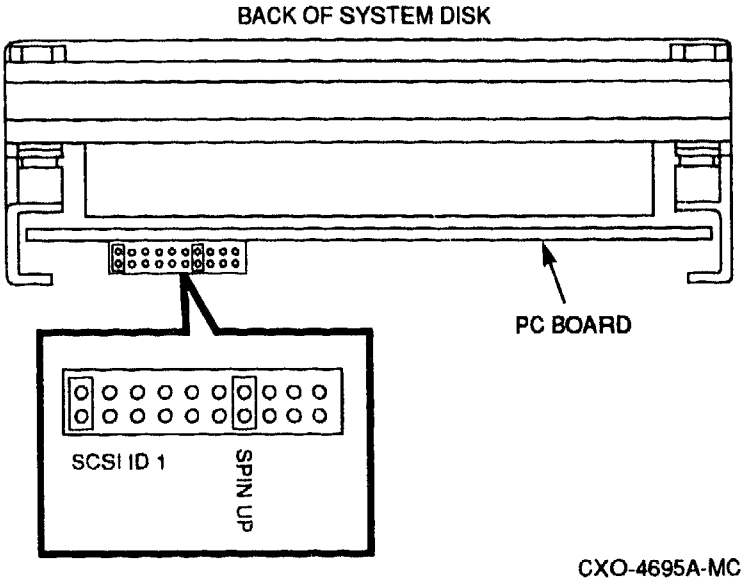
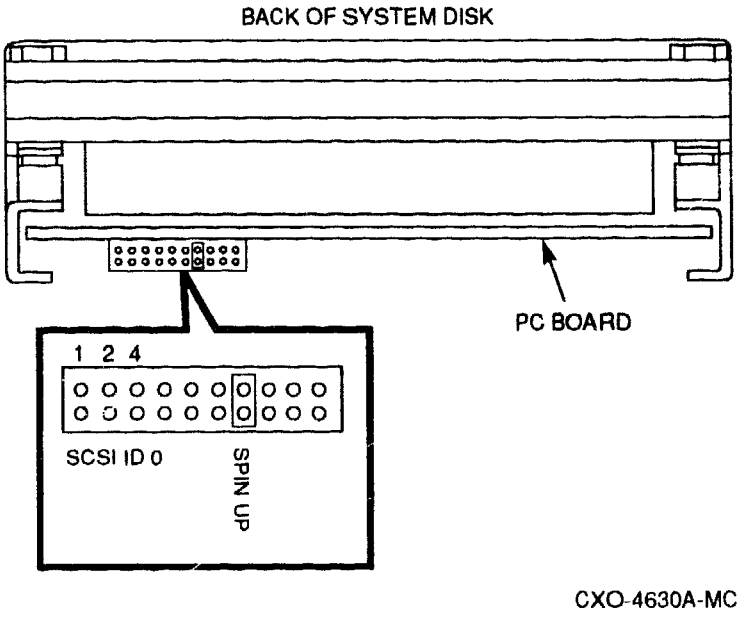


Figure D–24 System Disk Option Connector and SCSI ID 0



9. With SCSI ID jumpers removed from Server Processor 1, the system disk in the server processor shows device DKA0 when the SHOW DEVICE command is issued from the DCL prompt.

D.8.3 CD–ROM Drive SCSI Identifiers

The CD–ROM drive ID in Server Processor 2 should be set to an address different than Server Processor 1. Every server processor is shipped from the factory with its RRD43 CD–ROM drive ID set to SCSI ID 5. Normally, the CD–ROM SCSI ID in Server Processor 1, the lower server processor, is set to 5 and the CD–ROM SCSI ID of Server Processor 2 is set to 4.

If either SCSI ID needs to be changed, refer to the procedures found in Chapter 8 to access the desired CD–ROM drive and use the following procedure to change the SCSI ID.

1. Locate the option connector on the rear of the CD–ROM drive (see Figure D–25).

CAUTION

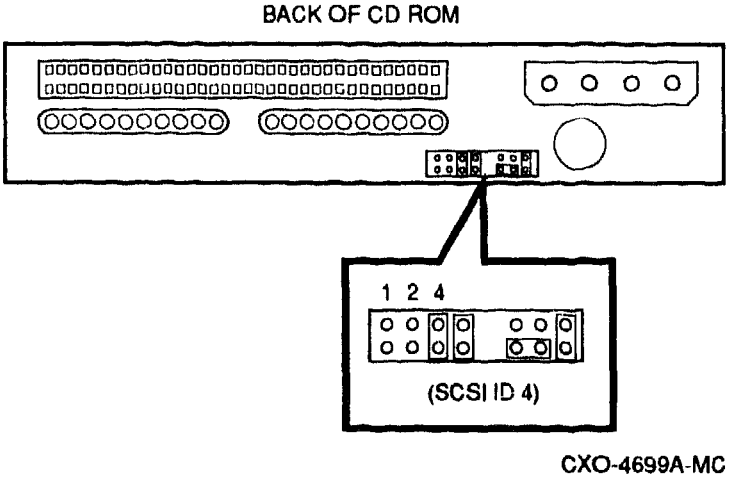
The first three pairs of pins on this option connector are used to set the drive SCSI ID. Do not change any other jumpers on the option connector.

2. If the SCSI ID is to be changed from 5 to 4, do the following:
   - Remove the black jumper from the pair of pins labeled SCSI ID 1 in Figure D–25. The CD–ROM drive SCSI ID now is set to SCSI ID 4.
   - Slide the unused jumper over the bottom two pins of ID 1 and ID 2 in case you want to set the CD–ROM drive to another ID at a later time.

     The RRD43 CD–ROM drive now appears as DKA400 when the SHOW DEVICE command is issued from the DCL prompt.
3. If the SCSI ID is to be changed from 4 to 5:
   - Install a black jumper onto the pair of pins labeled SCSI ID 1 in Figure D–25. The CD–ROM drive SCSI ID now is set to SCSI ID 4.

     The RRD43 CD–ROM drive now appears as DKA401 when the SHOW DEVICE command is issued from the DCL prompt.

Figure D–25 CD–ROM Drive Option Connector and SCSI ID 4



4. Reinstall the CD–ROM drive using the procedure found in Chapter 8.

D.9 Running the EISA Configuration Utility (ECU)

Whenever you add, remove, or move an EISA option, you need to run the EISA Configuration Utility (ECU). The ECU and a configuration file (CFG) for the HS1AD are located on the diskette included with the FDDI Server's software kit. The ECU uses the CFG file to allocate system resources and create a conflict-free configuration. This configuration information is saved to your server's nonvolatile memory.

Note

If you cannot find the file, the ECU program prompts you to insert the CFG diskette (AK–Q767A–CA) shipped with the HS1AD.

The ECU allows you to set device bus adapter parameters, as described in Section D.9.1.

D.9.1 How to Run the ECU

Refer to Section D.3 for information on attaching a maintenance terminal to your server processor.

Use the following procedure to run the ECU:

Note

If you do not set your terminal as described in the following procedure, the keyboard does not function correctly with the ECU. Your terminal mode setup must be set to 7-bit mode.

Check the communication setup and make sure it is set to 9600 baud, 8 bit, one stop bit, and no parity.

1. If the terminal is not displaying the >>> prompt, power cycle the server processor (depress the Power button on the server processor twice: once to turn it off and again to restart the system). When the >>> prompt appears on the terminal screen, the self test is complete.
2. Insert the ECU diskette into the floppy diskette drive.
3. Type ECU and press the Return key at the >>> prompt. The system loads the ECU software.

4. During loading, press Return to respond to system questions until the software displays the menu shown in the following example:

```
        EISA Configuration Utility
         Steps in configuring your computer

  STEP 1: Important EISA configuration Information

  STEP 2: Add or remove boards

  STEP 3: View or edit details

  STEP 4: Examine required details

  STEP 5: Save and exit

>Select=ENTER<  <Cancel=ESC>
```

5. Use the down arrow key to select the "STEP 3: View or edit details" option (shown in the previous example) and press the Enter key.

6. Use the down arrow key to scroll through the file until you find the device bus adapter option and its slot number. The display lists the current settings, as shown in the following example:

```
        Step 3: View or edit details

Slot 7 -- Digital device bus adapter internal bus EISA Host Adapter   Added

Host Adapter Interface.................Trigger EDGE
Host Adapter IRQ.......................IRQ 15
Host Adapter internal bus ID...........Device ID 7

Slot 8 -- Digital device bus adapter internal bus EISA Host Adapter   Added

Host Adapter Interface.................Trigger EDGE
Host Adapter IRQ.......................IRQ 14
Host Adapter internal bus ID...........Device ID 7
```

Note

The IRQ and device ID of the HS1AD are set automatically to the default values by the ECU. The previous display shows that the server processor contains one HS1AD with a device ID of 7 and an IRQ of 15, and another HS1AD with a device ID of 7 and an IRQ of 14.

7. To change the settings (Step 3, "View or edit details"), use the down arrow key to select the options you want to change and press the Return key.

8. Set both HS1ADs to SCSI ID 7.

9. When you have finished with the option settings, enter Ctrl/[(control left bracket). A main menu similar to the following displays:

```
        EISA Configuration Utility
         Steps in configuring your computer

  STEP 1: Important EISA configuration Information

  STEP 2: Add or remove boards

  STEP 3: View or edit details

  STEP 4: Examine required details


  STEP 5: Save and exit
```

```
>Select=ENTER<  <Cancel=ESC>
```

10. Use the arrow key to select "Step 5: Save and exit" (see the previous example) and press the Enter key. The display shows the configuration is saved.
11. Press Enter again, and the >>> prompt displays.
12. Remove the ECU diskette from the floppy diskette drive.
13. Power cycle the server processor (depress the Power button on the server processor twice: once to turn it off and again to restart the system). After the system performs the self test, the >>> prompt displays.

D.9.2 Verifying Your Installation

Use the following procedure to verify the installation of the HS1AD:

1. **At the >>> prompt, enter SHOW CONFIG, and press Return. A screen** display similar to the following appears:

```
Firmware
SRM Console:    V3.0-1
ARC Console:    4.6
PALcode:        VMS PALcode X5.48-49
Serial Rom:     V1.1

Processor
DECchip (tm) 21064-2

MEMORY
    128 Meg of System Memory
     Bank 0 = 64 Mbytes (16 MB Per Simm) Starting at 0x00000000
     Bank 1 = 64 Mbytes (16 MB Per Simm) Starting at 0x02000000
     Bank 2 = No Memory Detected
     Bank 3 = No Memory Detected

PCI Bus
     Bus 00  Slot 06: NCR    810 Scsi Controller
                                pka0.7.6.0          SCSI Bus ID 7
                                dka100.1.0.6.0      RZ28B
                                dka400.4.0.6.0      RRD43

     Bus 00  Slot 07: Intel  8275EB PCI to Eisa Bridge


EISA Bus Modules (installed)
     Slot 3   DEC3002            fra0.0.0.3.1        08-00-2B-A5-CC-93
     Slot 7   DEC2E00
     Slot 8   DEC2E00
```

2. If the second HS1AD appears in slot 07, the installation is correct.

   If slot 07 does not display on your terminal screen, the second HS1AD is not being recognized by the server processor.

3. Leave the server processor at the >>> prompt.

D.10 Upgrading Existing FDDI Servers

The StorageWorks FDDI Servers have a number of applicable upgrades: some are StorageWorks FDDI Server specific, others are applicable to the SW800 cabinet.

Table D–6 shows the recommended migration paths among the various StorageWorks FDDI Servers.

Table D–6 StorageWorks FDDI Server Migration Path Upgrade Matrix

From↓ To→	HS111	HS121	HS211	HS221	HS241	HS280[2]
SW800	HS110–AA	SW800→HS111 (HS110–AA), then HS111→HS121 (HS110–AA)[1]	HS210–AA	SW800→HS111 (HS110–AA), then HS211→HS221 (HS221–BX)	SW800→HS111 (HS110–AA), then HS211→HS221 (HS221–BX), then HS221→HS241 (HS241–BX)	NA
HS111	NA	HS110–AA[1]	HS211–AX	HS111→HS211 (HS211–AX), then HS211→HS221 (HS221–BX)	HS111→HS211 (HS211–AX), then HS211→HS221 (HS221–BX), then HS221→HS241 (HS241–BX)	NA
HS121	NA	NA	NA	HS221–AX	HS241–AX	NA
HS211	NA	NA	NA	HS221–BX	HS211→HS221 (HS221–BX), then HS221→HS241 (HS241–BX)	NA
HS221	NA	NA	NA	NA	HS241–BX	NA

[1]The HS111-to-HS121 upgrade requires purchase of additional parts not included in the HS110–AA kit. These parts include two (2) device bus adapters (HS1AD–AA), two internal bus cables (BC29S–09) and one (1) cable distribution unit (SW8XP–AX).

[2]StorageWorks FDDI Server expansion cabinet (HS280–AA) for expansion of model HS241 StorageWorks FDDI Server capacity.

E

Command Line Interpreter

This appendix provides the following information:

- A comprehensive list of all CLI commands
- CLI error messages the operator may encounter
- Examples of some common CLI-based procedures

An overview of how to use the CLI, as well as a description of how to access and exit the CLI, is provided in Chapter 5.

E.1 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.

Remember these two guidelines when using the CLI:

- You do not need to specify all configuration parameters on one line. You can enter them by using multiple SET commands.
- You need to enter only enough of each command to make the command unique (usually three characters). For example, SHO is equivalent to SHOW.

Note

Certain information related to the companion controller in a dual-redundant configuration does not always appear. When device or parameter information does not appear after entering a command from one controller, enter the same command from the other controller.

ADD CDROM

Adds a CD–ROM drive to the list of known CD-ROM drives.

Format

ADD CDROM *container-name SCSI-location*

Parameters

container-name
Specifies the name that is used to refer to this CD–ROM 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 CD–ROM drive to be added in the form PTL where **P** designates the port (1–6 or 1–3, depending on the controller model), **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 (must be 0).

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

Description

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

Examples

1.
```
CLI> ADD CDROM CD_PLAYER 1 0 0
```
Adds a CD–ROM 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 or 1–3, depending on the controller model), **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 (must be 0).

When entering the PTL, at least one space must separate 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 controller 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.
```
CLI> 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.
```
CLI> ADD DISK DISK0 2 3 0 NOTRANSPORTABLE
```
Adds a non transportable disk to port 2, target 3, LUN 0 and names it DISK0.

3.
```
CLI> 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 or 1–3, depending on the controller model), **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 (must be 0).

When entering the PTL, at least one space must separate 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. `CLI> ADD LOADER LOAD0 2 0 0`

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

ADD MIRRORSET

Binds a set of physical devices to a mirrorset specified by a container name.

Format

ADD MIRRORSET *container-name disk-device-name1 [disk-device-nameN]*

Parameters

container-name

Specifies the name that is used to refer to this mirrorset container. This name is referred to when creating mirrorsets. 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.

disk-device-name1 disk-device-nameN
The disk drives that make up this mirrorset. A mirrorset is made up of from 1 to 6 disk drives.

Description

Adds a mirrorset to the list of known mirrorsets and names the mirrorset. The number of members is set to the number of devices specified in the command. If the physical devices have never been initialized as a mirrorset, then a CLI INITIALIZE command must be issued for the mirrorset prior to binding the mirrorset to a higher level storage container.

Qualifiers

COPY=*copy_speed*

The COPY qualifier allows you to specify the speed at which mirrorset copies are performed. You may specify either NORMAL or FAST.

NORMAL uses relatively few controller resources to perform the copy, and has little impact on controller performance.

FAST uses more controller resources, which reduces the time it takes to complete the copy, but also reduces overall controller performance.

POLICY=BEST_FIT
POLICY=BEST_PERFORMANCE (Default)
NOPOLICY

The POLICY qualifier specifies the replacement policy to be used when a mirrorset member within the mirrorset 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 mirrorset. If more than one device in the spareset is the correct size, 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 mirrorset (the device should be on a different port). If more than one device in the spareset has the best performance, the device that most closely matches the size of the remaining members of the mirrorset is selected.

NOPOLICY retries a failing device from the mirrorset without selecting a replacement. This causes the mirrorset to run with less than the nominal number of members until a BEST_FIT or BEST_PERFORMANCE policy is selected, or a member is manually replaced in the mirrorset.

READ_SOURCE=*read-source*

The READ_SOURCE qualifier allows you to control the read algorithm for the specified mirrorset. The following choices are allowed for read-source:

ROUND_ROBIN—Each NORMAL mirrorset member is the target of a read in sequential membership order. No preference is given to any NORMAL member. This is the default read-source.

LEAST_BUSY—The NORMAL mirrorset member with the least busy work queue is the target of the read.

device-container-name—All reads are done on *device-container-name*. If *device-container-name* fails out of the mirrorset, the READ_SOURCE algorithm reverts to LEAST_BUSY.

Examples

1.
```
CLI> ADD MIRRORSET MIRR1 DISK100 DISK210 DISK320
```

   Adds DISK100, DISK210, and DISK320 as a mirrorset with the name MIRR1.

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 or 1–3, depending on the controller model), **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 (must be 0).

When entering the PTL, at least one space must separate 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 controller 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. `CLI> 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. `CLI> 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.

Note

A warning message (Warning 3010) may appear when you add a passthrough container. You can safely ignore the warning.

Examples

1.
```
CLI> ADD PASSTHROUGH PASS0 LDR421
```

Creates a passthrough container to loader LDR421 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 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 reconstructes when a new member is added to the RAIDset or immediately after the RAIDset is initialized.

RECONSTRUCT=NORMAL (default) balances overall performance of the controller 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 controller 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 controller to another. NOREDUCED (default) identifies that all RAIDset members that make up the RAIDset are being specified.

Examples

1. `CLI> 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. `CLI> 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. `CLI> 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 is *not* selected.

4. `CLI> 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 controller to replace failing members of RAIDsets and mirrorsets. The ADD SPARESET command adds disk drives to the spareset and initializes the metadata on the drives.

Examples

1. `CLI> ADD SPARESET DISK220`

   Adds one disk to the spareset.

2. `CLI> ADD SPARESET DISK110 DISK210 DISK320 DISK430 DISK540`

   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 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. `CLI> ADD STRIPESET STRIPE0 DISK100 DISK110 DISK220 DISK340`

   Creates a STRIPESET with four disks (DISK100, DISK110, DISK220, and DISK340).

2. `CLI> ADD STRIPESET STRIPE1 MR1 MR2 MR3`

   Creates a STRIPESET with three members, each of which is a mirrorset.

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 or 1–3, depending on the controller model), **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 (must be 0).

When entering the PTL, at least one space must separate 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.
```
CLI> 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 controller.

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 CD–ROMs) 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.

Note

The only target numbers specified in the unit number *must* have been previously specified in the SET THIS_CONTROLLER ID=(*n*1, *n*2, ...) command. You cannot specify a target number that has not been specified previously by the SET THIS_CONTROLLER ID=(*n*1, *n*2, ...) command.

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 CD–ROM Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit spin up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit is still spun up, but the unit is not 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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single controller, however, the qualifier does not take effect until a second controller is added and the two controllers are configured for dual-redundancy. The

second controller inherits any PREFERRED_PATH settings, and the two controllers operate using the preset PREFERRED_PATH options.

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller'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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single controller, however, the qualifier does not take effect until a second controller is added and the two controllers are configured for dual-redundancy. The

second controller inherits any PREFERRED_PATH settings, and the two controllers operate using the preset PREFERRED_PATH options.

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller'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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single controller, however, the qualifier does not take effect until a second controller is added and the two controllers are configured for dual-redundancy. The

second controller inherits any PREFERRED_PATH settings, and the two controllers operate using the preset PREFERRED_PATH options.

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 disables the writing of any new data.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the controller'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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller'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 are 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

Examples

1. `CLI> ADD UNIT D0 DISK0`

   Creates disk unit number 0 from container DISK0.

2. `CLI> ADD UNIT T0 TAPE12`

   Creates tape unit number 0 from container TAPE12.

3. `CLI> 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 controller firmware are displayed before the CLI prompt. These errors are displayed even after the error condition is rectified, until the controller 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.
```
CLI>

All NVPM components initialized to their default settings.
CLI> CLEAR_ERRORS CLI
CLI>
```

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 controller.

Format

CLEAR_ERRORS INVALID_CACHE *controller*

Parameters

controller
Specifies which controller clears 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 controller is installed on this controller, or if the write-back cache module with unwritten cache data is removed from this controller, 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.

Note

Entering the CLEAR_ERRORS INVALID_CACHE command on one controller causes the other controller to restart.

Examples

1.
```
CLI> CLEAR_ERRORS INVALID_CACHE THIS_CONTROLLER
```

Clears all cache information from this controller'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, and so forth) that has 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.
```
CLI> 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 controller 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 rectify a condition that caused a device to be marked UNKNOWN, this command must be issued for the controller to recognize the device.

Examples

1. `CLI> CLEAR_ERRORS UNKNOWN DISK300`

   Causes the controller 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, and so forth) that has 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.
```
CLI> 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. `CLI> DELETE DISK0`

   Deletes DISK0 from the list of known containers.

2. `CLI> DELETE STRIPE0`

   Deletes STRIPE0 from the list of known containers.

3. `CLI> DELETE RAID9`

   Deletes RAID9 from the list of known containers.

DELETE FAILEDSET

Delete 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 and mirrorsets, either because they failed or were manually removed via the SET 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. `CLI> DELETE FAILEDSET DISK220`

   Deletes one disk from the failedset.

2. `CLI> DELETE FAILEDSET DISK100 DISK210 DISK220 DISK330 DISK400`

   Deletes five disks from the failedset.

DELETE SPARESET

Delete 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 controller to replace failing members of RAIDsets and mirrorsets. The DELETE SPARESET command removes disk drives from the spareset.

Examples

1. `CLI> DELETE SPARESET DISK230`

   Deletes one disk from the spareset.

2. `CLI> DELETE SPARESET DISK110 DISK210 DISK240 DISK320 DISK400`

   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, and so forth) 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 controller, 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.
```
CLI> DELETE D12
```
Deletes disk unit number 12 from the list of known units.

2.
```
CLI> 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 controller is displayed.

Examples

1.
```
CLI> DIRECTORY
TILX    V25J   D
DILX    V25J   D
VTDPY   V25J   D
FLS     V25J   D
ECHO    V25J   D
DIRECT  V25J   D
CLI     V25J   D
ECHO    V25J   D
CFMENU  V25J   D
CHVSN   V25J   D
CLCP    V25J   D
CLONE   V25J   D
CONFIG  V25J   D
CRASH   V25J   D
FMU     V25J   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 controller type, and the last fail packet.

Examples

1.
```
CLI> EXIT

Copyright (C) Digital Equipment Corporation 1994
HSICP Firmware version V020-0, Hardware version 0000

Last fail code: 01800080

Press " ?" at any time for help.

CLI>
```

An EXIT command issued on a maintenance terminal.

2.
```
CLI> 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.
```
CLI> HELP
Help may be requested by typing a question mark (?) at the CLI prompt.  This
prints 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.
```
CLI> 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. During initialization, a small amount of disk space is used for controller metadata and is made inaccessible to the host.

If a single-disk container was set as TRANSPORTABLE, 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, stripeset, or mirrorset.

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, stripeset, or mirrorset)

The INITIALIZE command 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 for RAIDsets and stripesets. The chunksize may be specified in blocks (CHUNKSIZE=*n*), or you can let the controller determine the optimal chunksize (CHUNKSIZE=DEFAULT). The CHUNKSIZE qualifier does not apply to mirrorsets.

DESTROY(Default)
NODESTROY
This qualifier prevents the user data and forced error metadata from being destroyed during the initialize. This allows the data on the container to be reused for a disk, stripeset, or mirrorset unit. (The NODESTROY qualifier is ignored for RAIDsets.) NODESTROY is only used when creating a unit out of devices that have been reduced from mirrorsets.

Examples

1. `CLI> INITIALIZE DISK0`

   Initializes container DISK0. If NOTRANSPORTABLE was specified (or allowed to default), metadata is written on the disk.

2. `CLI> INITIALIZE STRIPE0 CHUNKSIZE=20`

   Initializes container STRIPE0 and writes metadata on it. The default chunksize for stripesets is 128 blocks for HSOF V2.5.

3. `CLI> INITIALIZE RAID9 CHUNKSIZE=20`

   Initializes container RAID9 with a chunksize of 20 and writes metadata on it. The default chunksize for RAIDsets is 128 blocks for HSOF V2.5.

4. `CLI> INITIALIZE MIRROR1 NODESTROY`

   Initializes container MIRROR1 and does not write over the forced error metadata.

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.

LOADERS
The LOCATE LOADERS command turns on the amber device fault LEDs of all configured loader devices. See LOCATE CANCEL to turn off the LEDs.

An error is displayed if no loader devices have been configured.

OPTICALS
The LOCATE OPTICALS command turns on the amber device fault LEDs of all configured optical drives. See LOCATE CANCEL to turn off the LEDs.

An error is displayed if no optical drives have been configured.

TAPES
The LOCATE TAPES command turns on the amber device fault LEDs of all configured tape devices. See LOCATE CANCEL to turn off the LEDs.

An error is displayed if no tape devices 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 the device 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 or 1–3, depending on the controller model), **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.
```
CLI> LOCATE DISK0
```

   Turns on the device fault LED on device DISK0.

2.
```
CLI> LOCATE D12
```

   Turns on the device fault LEDs on all devices that make up disk unit number 12.

3.
```
CLI> LOCATE DISKS
```

   Turns on the device fault LEDs on all configured disk devices.

MIRROR *disk-device-name1 container-name*

Allows you to convert a physical device to a one-member mirrorset.

Format

MIRROR *disk-device-name1 container-name*

Parameters

disk-device-name1
Specifies the name of the physical device that you wish to convert to a one member mirrorset. The device must be part of a unit.

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

Description

The MIRROR *disk-device-name1 container-name* command allows you to convert a physical device specified by *disk-device-name* to a mirrorset with one member. The mirrorset name is specified by *container-name*. This command can be used on devices that are already members of higher level containers (stripesets or units).

After you convert the device to a mirrorset, increase the nominal number of members with the SET *mirrorset-container-name* MEMBERSHIP=*number-of-members* command. Use the SET *mirrorset-container-name* REPLACE=*disk-device-name* command to actually add more members to the mirrorset. Refer to SET *mirror-container-name* for details for using the MEMBERSHIP= *number-of-members* qualifier versus the REPLACE= *disk-device-name* qualifier.

NOTE

When this command is used to create mirrorsets from stripeset members that were created before HSOF Version 2.5, the stripeset is then incompatible with prior firmware versions.

Qualifiers

COPY=*copy_speed*

The COPY qualifier allows you to specify the speed at which mirrorset copies are performed. You may specify either NORMAL or FAST.

NORMAL uses relatively few controller resources to perform the copy, and has little impact on controller performance.

FAST uses more controller resources, which reduces the time it takes to complete the copy, but also reduces overall controller performance.

POLICY=BEST_FIT
POLICY=BEST_PERFORMANCE
NOPOLICY (Default)

The POLICY qualifier specifies the replacement policy to be used when a mirrorset member within the mirrorset 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 mirrorset. If more than one device in the spareset is the correct size, 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 mirrorset (the device should be on a different port). If more than one device in the spareset has the best performance, the device that most closely matches the size of the remaining members of the mirrorset is selected.

NOPOLICY retries a failing device from the mirrorset without selecting a replacement. This causes the mirrorset to run with less than the nominal number of members until a BEST_FIT or BEST_PERFORMANCE policy is selected, or a member is manually replaced in the mirrorset.

Examples

1.
```
CLI> MIRROR DISK210 MIRROR5
```

   Creates a one member mirrorset from a DISK210 (a single disk).

REDUCE *disk-device-name1 [disk-device-nameN]*

Allows you to remove members from an existing mirrorset.

Format

REDUCE *disk-device-name1 [disk-device-nameN]*

Parameters

disk-device-name1
Specifies the name of the NORMAL mirrorset member to be removed.

[disk-device-nameN]
Specifies the name of the second mirrorset member to be removed.

Description

This command allows you to remove members from mirrorsets. For a consistent copy of a stripeset whose members are mirrorsets, all mirrorsets must be reduced at the same time with one command. It is similar to the SET *mirrorset-container-name* REMOVE=*disk-device-name* command, except that the nominal number of members in the mirrorset is decreased by the number of members removed, and the devices are not placed in the failedset.

Note

Mirrorsets have both a nominal number of members and an actual number of members. The nominal number of members in a mirrorset is the number given in the SET *mirrorset- name* MEMBERSHIP= *number-of-members* command. If devices have not been added to the mirrorset, or if a member is removed, the actual number of members may be less than the nominal number. The actual number of members can never be greater than the nominal number of members.

The disk devices to be removed need not be members of the same mirrorset. The devices MUST be part of the same unit (for example, the same stripeset). This is an atomic operation because the I/O to the unit associated with the given mirrorset members is stalled while the specified mirrorset members are removed. No auto-sparing occurs and each mirrorset membership is set to the new reduced number of members. The removed devices specified by *disk-device-name1* through *disk-device-nameN* are not placed in the failedset, but are left as unused devices.

Note that for each mirrorset that you reduce, the mirrorset must have at least one remaining NORMAL member after the reduction. If this is not true for all *disk-devices-names* specified, then none of the specified mirrorsets are reduced.

Note

A NORMAL member is a mirrorset member whose entire contents are guaranteed to be the same as all other NORMAL members. All NORMAL members have exactly the same contents.

Examples

1.
```
CLI> REDUCE DISK210 DISK110
```
DISK210 and DISK110 are removed from their respective mirrorsets.

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.
```
CLI> RENAME DISK0 DISK100
```

Renames container DISK0 to DISK100.

RESTART OTHER_CONTROLLER

Restarts the other controller.

Format

RESTART OTHER_CONTROLLER

Description

The RESTART OTHER_CONTROLLER command flushes all user data from the other controller's write-back cache (if present), then restarts the other controller.

If any disks are online to the other controller, the controller does not restart unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the controller does not restart unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes the other controller 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 controller in a dual-redundant configuration. The other controller restarts and resumes operations where it was interrupted.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the controller is not restarted unless IGNORE_ERROR is specified.

If the IGNORE_ERRORS qualifier is specified, the controller 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_SHUTDOWN
NOIMMEDIATE_SHUTDOWN (Default)
If IMMEDIATE is specified, the controller 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 controller, the controller is not restarted unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the controller 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. `CLI> RESTART OTHER_CONTROLLER`

   Restarts the other controller as long as the other controller does not have any units online.

2. `CLI> RESTART OTHER_CONTROLLER OVERRIDE_ONLINE`

   Restarts the other controller even if there are units online to the other controller.

RESTART THIS_CONTROLLER

Restarts this controller.

Format

RESTART THIS_CONTROLLER

Description

The RESTART THIS_CONTROLLER command flushes all user data from this controller's write-back cache (if present), then restarts this controller.

If any disks are online to the other controller, the controller does not restart unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the controller does not restart unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes this controller 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 controller in a dual-redundant configuration. This controller 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 controller, the connection is lost when the controller restarts.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the controller is not restarted unless IGNORE_ERROR is specified.

If the IGNORE_ERRORS qualifier is specified, the controller 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_SHUTDOWN
NOIMMEDIATE_SHUTDOWN (Default)
If IMMEDIATE is specified, the controller 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 controller, the controller is not restarted unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the controller 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. `CLI> RESTART THIS_CONTROLLER`

   Restarts this controller as long as this controller does not have any units that are online.

2. `CLI> RESTART THIS_CONTROLLER OVERRIDE_ONLINE`

   Restarts this controller even if there are units online to this controller.

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, and so forth) 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.
```
CLI> 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 CLONE are examples of utilities and diagnostics that can be run from the CLI.

Description

The RUN command starts various diagnostics and utilities on THIS_CONTROLLER. Diagnostics and utilities can be run *only* on the controller where the terminal or DUP connection is connected.

For specific information about available diagnostics and utilities, refer to Chapter 7.

Examples

1.
```
CLI> RUN DILX
Disk Inline Exerciser - version 2.0
             .
             .
             .
```

Runs the DILX diagnostic.

SELFTEST OTHER_CONTROLLER

Runs a self-test on the other controller.

Format

SELFTEST OTHER_CONTROLLER

Description

The SELFTEST OTHER_CONTROLLER command flushes all user data from the other controller's write-back cache (if present), shuts down the other controller, then restarts it in DAEMON loop-on-self-test mode. The OCP reset (//) button must be pressed to take the other controller out of loop-on-self-test mode.

If any disks are online to the other controller, the controller does not self-test unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the controller does not self-test unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes the other controller 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 controller does not start the self-test unless IGNORE_ERRORS is specified.

If the IGNORE_ERRORS qualifier is specified, the controller 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 controller immediately starts a 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 controller, the controller does not self-test unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the controller 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. `CLI> SELFTEST OTHER_CONTROLLER`

   Starts the self-test on the other controller, as long as the other controller does not have any units online.

2. `CLI> SELFTEST OTHER_CONTROLLER OVERRIDE_ONLINE`

   Starts the self-test on the other controller even if there are units online to the other controller.

SELFTEST THIS_CONTROLLER

Runs a self-test on this controller.

Format

SELFTEST THIS_CONTROLLER

Description

The SELFTEST THIS_CONTROLLER command flushes all user data from this controller's write-back cache (if present), shuts down this controller, then restarts it in DAEMON loop-on-self-test mode. The OCP reset (//) button must be pressed to take this controller out of loop-on-self-test mode.

If any disks are online to this controller, the controller does not self-test unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the controller does not self-test unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes this controller 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 controller, the connection is lost when this controller starts the self-test.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the controller does not start the self-test unless IGNORE_ERRORS is specified.

If the IGNORE_ERRORS qualifier is specified, the controller 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 controller immediately starts a 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 controller, the controller does not self-test unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the controller 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. `CLI> SELFTEST THIS_CONTROLLER`

   Starts the self-test on this controller as long as this controller does not have any units online.

SET *disk-container-name*

Changes the transportable characteristics of a disk drive.

Format

SET *disk-container-name*

Parameters

disk-container-name
The name of the disk drive that has its characteristics changed.

Description

Changes the characteristics of a disk drive.

Qualifiers

TRANSPORTABLE
NOTRANSPORTABLE (Default)
In normal operations, the controller 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. `CLI> 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 one that either the maintenance terminal is connected to or the virtual terminal is connected to) is copied to the other controller.

If OTHER_CONTROLLER is specified for *configuration-source*, all the device configuration information on the OTHER_CONTROLLER (the controller that either the maintenance terminal or the virtual terminal connection is *not* connected to) is copied to this controller.

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 controllers fail, the devices and cache (if any) attached to the failed controller become available to and accessible through the operating controller.

CAUTION

All device configuration information on the controller *not* specified by the COPY= parameter is destroyed and overwritten by the configuration information found in the controller 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 controller is specified by the COPY= parameter.

Also note that due to the amount of information that must be passed between the two controllers, this command may take up to 1 minute to complete.

Examples

1. `CLI> SET FAILOVER COPY=THIS_CONTROLLER`

   Places two controllers into a dual-redundant configuration, where the "good" data was on the controller that the maintenance terminal or virtual terminal connection was connected to.

2. `CLI> SET FAILOVER COPY=OTHER_CONTROLLER`

   Places two controllers into a dual-redundant configuration, where the "good" data was on the controller that the maintenance terminal or virtual terminal connection was *not* connected to.

SET *mirrorset-container-name*

Changes the characteristics of a mirrorset.

Format

SET *mirrorset-container-name*

Parameters

mirrorset-container-name
The name of the mirrorset that has its characteristics modified. 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

Changes the characteristics of a mirrorset.

Qualifiers

MEMBERSHIP=*number-of-members*
Allows you to increase or decrease the nominal membership of a mirrorset to the number specified by *number-of-members*.

If the mirrorset membership is increased, and auto-sparing is turned on (by specifying POLICY=BEST_FIT or POLICY=BEST_PERFORMANCE), the mirrorset automatically brings in spares until either the new number of members is reached, or there are no more suitable spares.

If auto-sparing is turned off (by specifying NOPOLICY), then you must use the SET *mirrorset-container-name* REPLACE=*disk-device-name* command to bring the mirrorset up to the new nominal number of members.

You may not set the nominal number of members to be lower than the number of members physically present. Use the REMOVE=*disk-container-name* qualifier to reduce the number of devices that are part of the mirrorset.

REPLACE=*disk-device-name*
Specifies the replacement of a disk member into an existing mirrorset provided the following two conditions are met:

- The replacement policy must be set to NOPOLICY.
- The mirrorset must be missing at least one member.

If these two conditions are met, the device specified by *disk-device-name* is added to the mirrorset specified by *mirrorset-container-name*. The nominal number of members does not change.

Note

No other qualifiers to the SET *mirrorset-device-name* command may be specified if the REPLACE qualifier is specified.

REMOVE=*disk-container-name*

The SET *mirrorset-container-name* REMOVE=*disk-device-name* CLI command allows you to remove members from an existing mirrorset. The device specified by *disk-device-name* is removed from the mirrorset specified by *mirrorset-container-name*. If the physical device is not a member of the mirrorset, or if the mirrorset does not have a remaining NORMAL or NORMALIZING member, then an error is reported and no action is taken. On successful removal, the removed device is added to the failedset and a new member is auto-spared into the mirrorset (if applicable).

Note

A **NORMAL** member is a mirrorset member whose entire contents is guaranteed to be the same as all other NORMAL members. All NORMAL members have exactly the same contents.

NORMALIZING members only exist when a mirrorset is first created. One member is identified as NORMAL, and all other *original* mirrorset members are marked as NORMALIZING. All new data that is written to the mirrorset is written to all members. All data on the NORMAL member that existed before the mirrorset was created is copied to the NORMALIZING members. When all the blocks on the members are the same, the NORMALIZING members are marked as NORMAL.

(Members may also be marked as NORMALIZING if cache data is lost.)

Note that the nominal number of members in the mirrorset does not change. If auto-sparing does not occur, then when an acceptable spare becomes available or when the replacement policy changes, the mirrorset automatically adds the spare.

Note

No other qualifiers to the SET *mirrorset-device-name* command may be specified if the REMOVE qualifier is specified.

POLICY=BEST_FIT
POLICY=BEST_PERFORMANCE (Default)
NOPOLICY

The SET *mirrorset-container-name* POLICY=*policy-type* command specifies the replacement policy to be used when a mirrorset member within the mirrorset 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 mirrorset. If more than one device in the spareset is the correct size, 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 mirrorset (the device should be on a different port). If more than one device in the spareset has the best performance, the device that most closely matches the size of the remaining members of the mirrorset is selected.

NOPOLICY retries a failing device from the mirrorset without selecting a replacement. This causes the mirrorset to run with less than the nominal number of members until a BEST_FIT or BEST_PERFORMANCE policy is selected, or a member is manually replaced in the mirrorset.

READ_SOURCE=*read-source*

The SET *mirrorset-container-name* READ_SOURCE=*read-source* command allows you to control the read algorithm for the specified mirrorset. The following choices are allowed for read-source:

ROUND_ROBIN—Each NORMAL mirrorset member is the target of a read in sequential membership order. No preference is given to any NORMAL member. This is the default read-source.

LEAST_BUSY—The NORMAL mirrorset member with the least busy work queue is the target of the read.

device-container-name—All reads are done on *device-container-name*. If *device-container-name* fails out of the mirrorset, the READ_SOURCE algorithm reverts to LEAST_BUSY.

COPY=*copy_speed*

The SET *mirrorset-container-name* COPY=*copy_speed* command allows you to specify the speed at which mirrorset copies are performed. You may specify either NORMAL or FAST.

NORMAL uses relatively few controller resources to perform the copy, and has little impact on controller performance.

FAST uses more controller resources, which reduces the time it takes to complete the copy, but also reduces overall controller performance.

Examples

1. ```
   CLI> SET MIRR9 POLICY=BEST_FIT
   ```

   Changes mirrorset MIRR9's policy to BEST_FIT.

2. ```
   CLI> SET MIRR9 REMOVE=DISK0
   ```

   Removes mirrorset MIRR9's member DISK0 from the mirrorset. If there is a replacement policy, a new disk is taken from the spareset and placed in the mirrorset automatically.

3. ```
   CLI> SET MIRR9 REPLACE=DISK320
   ```

   Adds disk DISK320 to the reduced mirrorset, MIRR9. A reconstruct operation begins immediately on DISK320.
   ```

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 controller should be physically removed because the sharing of devices is not supported by single controller configurations.

The controller on which the command was entered is always removed from a dual-redundant state, even if the other controller is not currently reachable. No configuration information is lost when leaving a dual-redundant state.

Examples

1. `CLI> SET NOFAILOVER`

   Removes the two controllers from a dual-redundant configuration.

SET OTHER_CONTROLLER

Changes the other controller's parameters (in a dual-redundant configuration, the controller that the maintenance terminal is *not* connected to, or the controller that is *not* the target of the DUP connection.

Format

SET OTHER_CONTROLLER

Description

The SET OTHER_CONTROLLER command allows you to modify the controller parameters of the other controller 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. You must restart the controller before changes to the CACHE_FLUSH_TIMER take effect.

CACHE_POLICY=A (Default)
CACHE_POLICY=B
Allows selection of the write-back cache battery policy used by the controller. The policy affects the availability of RAIDsets and mirrorsets when the battery condition is low during controller initialization.

Changing this policy on one controller in a dual-redundant controller configuration automatically changes it on the other controller. The parameter change takes effect immediately.

The CACHE_POLICY setting affects RAIDsets and mirrorsets as follows:

- If the CACHE_POLICY on the controller is set to A and the batteries are low when the controller initializes, any RAIDset or mirrorset that does not have access to good batteries is made inoperative.
- If the CACHE_POLICY on the controller is set to B and the batteries are low when the controller initializes, any RAIDset or mirrorset that does not have access to good batteries is accessed in write-through (read cache) mode.

Regardless of the CACHE_POLICY setting, a low or bad battery affects controller operation:

- Stripesets and disk-based units with write-back caching enabled are accessed in write-through (read cache) mode, as long as unwritten cache data has not been lost.
- RAIDsets and mirrorsets fail over to the companion cache in a dual-redundant configuration, provided the other cache's batteries are fully charged.

  If the batteries go low after controller initialization, unwritten cache data is flushed from the cache and any RAIDset or mirrorset that does not have access to good batteries is made inoperative, regardless of the cache policy.

- If the batteries are bad or missing, RAIDsets and mirrorsets are made inoperative.
- Write-back caching automatically resumes when the cache batteries are fully recharged or replaced.

CAUTION

There is some risk in setting CACHE_POLICY=B to allow access to mirrorsets and RAIDsets when the batteries are low. Because the batteries may be in an unknown state, there is no guarantee as to how long they maintain data in the cache should a power failure occur. Fully charged batteries maintain the data for a minimum of 100 hours.

ID=*n*
Specifies the DSSI node number (0–7).

MSCP_ALLOCATION_CLASS=*n*
Specifies the allocation class (0–255 in a single controller configuration or 1–255 in a dual-redundant configuration).

When first installed, the controller'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 is displayed when the controller's CLI prompts for input. Only printable ASCII characters are valid.

When first installed, the CLI prompt is set to the first three letters of the controller's model number (for example, 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 controller to not check for, or transmit any parity on the terminal lines.

When first installed, the controller'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 controller's terminal speed is set to 9600 baud.

TMSCP_ALLOCATION_CLASS=*n*
Specifies the allocation class (0–255 in a single controller configuration or 1–255 in a dual-redundant configuration).

When first installed, the controller's TMSCP_ALLOCATION_CLASS is set to 0.

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 has 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 is 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 controller 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 controller 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. `CLI> SET RAID9 POLICY=BEST_FIT`

   Changes RAIDset RAID9's policy to BEST_FIT.

2. `CLI> 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. `CLI> SET RAID9 REPLACE=SPAREDISK`

   Adds disk SPAREDISK to the reduced RAIDset, RAID9. A reconstruct operation begins immediately on SPAREDISK.

SET THIS_CONTROLLER

Changes this controller's parameters (the controller 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 controller 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.

CACHE_POLICY=A (Default)
CACHE_POLICY=B
Allows selection of the write-back cache battery policy used by the controller. The policy affects the availability of RAIDsets and mirrorsets when the battery condition is low during controller initialization.

Changing this policy on one controller in a dual-redundant controller configuration automatically changes it on the other controller. The parameter change takes effect immediately.

The CACHE_POLICY setting affects RAIDsets and mirrorsets as follows:

- If the CACHE_POLICY on the controller is set to A and the batteries are low when the controller initializes, any RAIDset or mirrorset that does not have access to good batteries is made inoperative.
- If the CACHE_POLICY on the controller is set to B and the batteries are low when the controller initializes, any RAIDset or mirrorset that does not have access to good batteries is accessed in write-through (read cache) mode.

Regardless of the CACHE_POLICY setting, a low or bad battery affects controller operation:

- Stripesets and disk-based units with write-back caching enabled are accessed in write-through (read cache) mode, as long as unwritten cache data has not been lost.
- RAIDsets and mirrorsets fail over to the companion cache in a dual-redundant configuration, provided the other cache's batteries are fully charged.

  If the batteries go low after controller initialization, unwritten cache data is flushed from the cache and any RAIDset or mirrorset that does not have access to good batteries is made inoperative, regardless of the cache policy.
- If the batteries are bad or missing, RAIDsets and mirrorsets are made inoperative.

- Write-back caching automatically resumes when the cache batteries are fully recharged or replaced.

CAUTION

There is some risk in setting CACHE_POLICY=B to allow access to mirrorsets and RAIDsets when the batteries are low. Because the batteries may be in an unknown state, there is no guarantee as to how long they maintain data in the cache should a power failure occur. Fully charged batteries maintain the data for a minimum of 100 hours.

ID=*n*
Specifies the DSSI node number (0–7).

MSCP_ALLOCATION_CLASS=*n*
Specifies the allocation class (0–255 in a single controller configuration or 1–255 in a dual-redundant configuration).

When first installed, the controller'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 is displayed when the controller's CLI prompts for input. Only printable ASCII characters are valid.

When first installed, the CLI prompt is set to the first three letters of the controller's model number (for example, 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 controller to not check for, or transmit any parity on the terminal lines.

When first installed, the controller'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 controller's terminal speed is set to 9600 baud.

TMSCP_ALLOCATION_CLASS=*n*
Specifies the allocation class (0–255 in a single controller configuration or 1–255 in a dual-redundant configuration).

When first installed, the controller's TMSCP_ALLOCATION_CLASS is set to 0.

SET *unit-number*

Changes the unit parameters.

Format

SET *unit-number*

Parameters

unit-number
Specifies the logical unit number (D0–D4094 or T0–T4094, and so forth) 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 CD–ROM Drive

MAXIMUM_CACHED_TRANSFER=*n*
MAXIMUM_CACHED_TRANSFER=32 (Default)
Specifies the maximum size transfer in blocks to be cached by the controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller'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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller'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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 disables the writing of any new data.

WRITEBACK_CACHE
NOWRITEBACK_CACHE (Default)
Enables and disables the controller'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 controller. 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

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

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

RUN (Default)
NORUN
Enables and disables a unit's availability to the host. When RUN (default) is specified, the devices that make up the unit are spun up and the unit is made available to the host. If NORUN is specified, the devices that make up the unit are still spun up, but the unit is not 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 controller'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 are 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 controller that the unit should be accessed through (PREFERRED_PATH=) or whether the unit may be accessed through either controller (NOPREFERRED_PATH).

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

When the failed controller is restarted, the drives automatically return to the controller specified by the PREFERRED_PATH qualifier.

Note

The PREFERRED_PATH qualifier may be specified on a single controller, however, the qualifier does not take effect until a second controller is added and the two controllers are configured for dual-redundancy. The

second controller inherits any PREFERRED_PATH settings, and the two controllers operate using the preset PREFERRED_PATH options.

Examples

1. `CLI> SET D1 WRITE_PROTECT NOREAD_CACHE`

   Sets the write protect and turns off the read cache on unit D1.

2. `CLI> SET T47 DEFAULT_FORMAT=1600BPI_9TRACK`

   Sets unit T47 to 1600 bpi.

SHOW CDROMS

Shows all CD–ROM drives and drive information.

Format

SHOW CDROMS

Description

The SHOW CD–ROMS command displays all the CD–ROM drives known to the controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
CLI> SHOW 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 CD-ROMs.

2.
```
CLI> SHOW 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 CD–ROMs.

SHOW *cdrom-container-name*

Shows information about a CD–ROM.

Format

SHOW *cdrom-container-name*

Parameters

cdrom-container-name
The name of the CD–ROM drive to be displayed.

Description

The SHOW *cdrom-container-name* command is used to show specific information about a particular CD–ROM drive.

Examples

1.
```
CLI> SHOW CDROM230
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------

CDROM230      cdrom            2    3    0                    D623
          DEC      RRD44   (C) DEC  3593
```

A listing of CD–ROM 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 controller. First disks are shown, then tapes, then CD–ROMs 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.
```
CLI> SHOW DEVICES
Name            Type           Port Targ  Lun                Used by
-------------------------------------------------------------------------

DIO             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 controller.

2.
```
CLI> SHOW DEVICES FULL
Name            Type           Port Targ  Lun                Used by
-------------------------------------------------------------------------

DIO             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 controller.

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 controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
CLI> 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 controller.

2.
```
CLI> SHOW DISKS FULL
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------

DI0           disk             1    0    0                    D100
          DEC      RZ26     (C) DEC X388
        Switches:
          NOTRANSPORTABLE
        Size: 2050353 blocks
DI1           disk             1    1    0                    0110
          DEC      RZ26     (C) DEC T386
        Switches:
          NOTRANSPORTABLE
        Size: 2050353 blocks
```

Shows a full listing of disks attached to the controller.

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.
```
CLI> SHOW DI3
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------

DI3           disk             1    3    0                   D130
          DEC      RZ26     (C) DEC X388
        Switches:
          NOTRANSPORTABLE
        Size: 2050353 blocks
```

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.
```
CLI> 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 controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
CLI> SHOW LOADERS
Name            Type           Port Targ  Lun                    Used by
-------------------------------------------------------------------------------
LDR511          loader            5    1    1                    CMD511
```

Basic listing of loaders.

2.
```
CLI> 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

1.
```
CLI> SHOW LDR511
Name          Type          Port Targ  Lun                    Used by
-------------------------------------------------------------------------

LDR511        loader           5    1    1                    CMD511
          DEC      TZ Media Changer 930A
```

Listing of loader LDR511

SHOW MIRRORSETS

Shows all configured mirrorsets and any mirrorset-specific data related only to mirrorsets.

Format

SHOW MIRRORSETS

Description

The SHOW MIRRORSETS command displays all the mirrorsets known to the controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each mirrorset.

Examples

1.
```
CLI> SHOW MIRRORSETS
Name       Storageset          Uses                 Used by
-------------------------------------------------------------------------------
MIRR2      mirrorset           DISK110              STR0
                               DISK210
                               DISK310
                               DISK350

MIRR3      mirrorset           DISK120              STR0
                               DISK250

MIRR4      mirrorset           DISK330              STR0
```

MIRR2 is a four-member mirrorset. MIRR3 is a two-member mirrorset. MIRR4 is a one-member mirrorset. All three mirrorsets are members of stripeset STR0.

2.
```
CLI>SHOW MIRRORSETS FULL
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------------
MIRR2         mirrorset                     DISK110          STR0
                                            DISK210
                                            DISK310
                                            DISK350

        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          COPY (priority) = NORMAL
          READ_SOURCE = LEAST_BUSY
          MEMBERSHIP = 4, 4 members present
        State:
          DISK210    (member  0) is NORMAL
          DISK310    (member  1) is NORMAL
          DISK350    (member  2) is NORMAL
          DISK110    (member  3) is NORMAL
        Size: 2050353 blocks
```

```
MIRR3        mirrorset                          DISK120          STR0
                                                DISK250
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          COPY (priority) = FAST
          READ_SOURCE = LEAST_BUSY
          MEMBERSHIP = 2, 2 members present
        State:
          DISK250    (member  0) is NORMAL
          DISK120    (member  1) is NORMAL
        Size: 4109470 blocks

MIRR4        mirrorset                          DISK330          STR0
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          COPY (priority) = NORMAL
          READ_SOURCE = LEAST_BUSY
          MEMBERSHIP = 1, 1 member present
        State:
          DISK330    (member  0) is NORMAL
        Size: 2050353 blocks
```

Shows extended information for all mirrorsets known to the controller.

SHOW *mirrorset-container-name*

Shows the same information as SHOW MIRRORSETS FULL except that it only displays information on the mirrorset specified by *mirrorset-container-name*.

Format

SHOW *mirrorset-container-name*

Parameters

mirrorset-container-name
The name of the mirrorset to be displayed.

Description

The SHOW *mirrorset-container-name* command is used to show specific information about a particular mirrorset.

Examples

1.
```
CLI>SHOW MIRR0
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------
MIRR0           mirrorset                   DISK130          D290
        Switches:
          NOPOLICY (for replacement)
          COPY (priority) = NORMAL
          READ_SOURCE = LEAST_BUSY
          MEMBERSHIP = 2, 1 member present
        State:
          DISK130   (member  0) is NORMAL
        Size: 4109470 blocks
```

Shows a complete listing of the mirrorset named MIRR0.

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 controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
CLI> 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 controller.

2.
```
CLI> 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 controller.

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.
```
CLI> 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 controller.

Format

SHOW OTHER_CONTROLLER

Description

Shows all controller, port, and terminal information for the other controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information is displayed after the basic controller information.

Examples

1.
```
CLI> SHOW OTHER_CONTROLLER
Controller:
        HS1CP    (C) DEC ZG33400022 Firmware V2.5, Hardware 0000
        Configured for dual-redundancy with ZG33400026
            In dual-redundant configuration
        SCSI address 6
        Time: 14-FEB-1995 18:23:27
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)

        CACHE_POLICY = A
```

Shows the basic HS1CP controller 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 controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each storageset.

Examples

1.
```
CLI> 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.
```
CLI> SHOW CMD100
MASS> sho pass
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------

CMD100        passthrough                   DISK100          D610
```

Shows a listing of passthrough container CMD100.

xxx

SHOW RAIDSETS

Shows RAIDsets and RAIDset information.

Format

SHOW RAIDSETS

Description

The SHOW RAIDSETS command displays all the RAIDsets known by the controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each storageset.

Examples

1.
```
CLI> 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.
```
CLI> SHOW RAIDSETS FULL
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------

R0            raidset                       DISK110          D401
                                            DISK220
                                            DISK310
                                            DISK400

        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
          CHUNKSIZE = 63 blocks
        State:
          RECONSTRUCT 3% complete
          DISK220   (member  0) is RECONSTRUCTING
          DISK310   (member  1) is RECONSTRUCTING
          DISK400   (member  2) is RECONSTRUCTING
        Size: 2050353 blocks
```

```
R1              raidset                          DISK130
                                                 DISK240
                                                 DISK330
                                                 DISK420
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
          CHUNKSIZE = 63 blocks
        State:
          NORMAL
          DISK130    (member  0) is NORMAL
          DISK240    (member  1) is NORMAL
          DISK330    (member  2) is NORMAL
          DISK420    (member  3) is NORMAL
        Size: 2050353 blocks
```

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.
```
CLI> SHOW RAID9
Name          Storageset                          Uses             Used by
-------------------------------------------------------------------------------

RAID9         raidset                             DISK130
                                                  DISK240
                                                  DISK330
                                                  DISK420
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          RECONSTRUCT (priority) = NORMAL
          CHUNKSIZE = 63 blocks
        State:
          NORMAL
          DISK130    (member  0) is NORMAL
          DISK240    (member  1) is NORMAL
          DISK330    (member  2) is NORMAL
          DISK420    (member  3) is NORMAL
        Size: 2050353 blocks
```

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.
```
CLI> 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 controller. A storageset is any collection of containers, such as stripesets, mirrorsets, RAIDsets, the spareset and the failedset.

Stripesets are displayed first, followed by mirrorsets, 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.
```
CLI> SHOW STORAGESETS
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------------

S0            stripeset                     DISK500          D1
                                            DISK510
                                            DISK520
```

Shows a basic listing of all storagesets.

2.
```
CLI> SHOW STORAGESETS FULL
Name          Storageset                    Uses             Used by
-------------------------------------------------------------------------------

S0            stripeset                     DISK530
                                            DISK550
                                            DISK600

        Switches:
          CHUNKSIZE = 24 blocks
        State:
          NORMAL
          DISK530    (member  0) is NORMAL
          DISK550    (member  1) is NORMAL
          DISK600    (member  2) is NORMAL
        Size: 2050353 blocks
```

```
S1              stripeset                       DISK620
                                                DISK640
          Switches:
            CHUNKSIZE = 24 blocks
          State:
            NORMAL
            DISK620    (member  0) is NORMAL
            DISK640    (member  1) is NORMAL
          Size: 31304354 blocks

R0              raidset                         DISK110         D401
                                                DISK220
                                                DISK310
                                                DISK400
          Switches:
            POLICY (for replacement) = BEST_PERFORMANCE
            RECONSTRUCT (priority) = NORMAL
            CHUNKSIZE = 63 blocks
          State:
            NORMAL
            DISK110    (member  0) is NORMAL
            DISK220    (member  1) is NORMAL
            DISK310    (member  2) is NORMAL
            DISK400    (member  3) is NORMAL
          Size: 2050353 blocks

R1              raidset                         DISK130
                                                DISK240
                                                DISK330
                                                DISK420
          Switches:
            POLICY (for replacement) = BEST_PERFORMANCE
            RECONSTRUCT (priority) = NORMAL
            CHUNKSIZE = 63 blocks
          State:
            NORMAL
            DISK130    (member  0) is NORMAL
            DISK240    (member  1) is NORMAL
            DISK330    (member  2) is NORMAL
            DISK420    (member  3) is NORMAL
          Size: 2050353 blocks

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 controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each storageset.

Examples

1.
```
CLI> 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.
```
CLI> SHOW STRIPESETS FULL
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------

S0            stripeset                     DISK530
                                            DISK550
                                            DISK600

        Switches:
          CHUNKSIZE = 24 blocks
        State:
          NORMAL
          DISK530    (member  0) is NORMAL
          DISK550    (member  1) is NORMAL
          DISK600    (member  2) is NORMAL
        Size: 2050353 blocks
```

```
S1              stripeset                       DISK620
                                                DISK640
        Switches:
          CHUNKSIZE = 24 blocks
        State:
          NORMAL
          DISK620    (member  0) is NORMAL
          DISK640    (member  1) is NORMAL
        Size: 2050353 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.
```
CLI> SHOW STRIPE0
Name          Storageset                          Uses             Used by
------------------------------------------------------------------------------

STRIPE0       stripeset                           DISK530
                                                  DISK550
                                                  DISK600

        Switches:
          CHUNKSIZE = 24 blocks
        State:
          NORMAL
          DISK530   (member  0) is NORMAL
          DISK550   (member  1) is NORMAL
          DISK600   (member  2) is NORMAL
        Size: 2050353 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 controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information may be displayed after each device.

Examples

1.
```
CLI> SHOW TAPES
Name            Type           Port Targ  Lun                    Used by
-------------------------------------------------------------------------------

TAPE200         tape             2    0    0
```

Shows a basic listing of tape drives.

2.
```
CLI> 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.
```
HS1CP> 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 controller.

Format

SHOW THIS_CONTROLLER

Description

Shows all controller, port, and terminal information for this controller.

Qualifiers

FULL
If the FULL qualifier is specified, additional information is displayed after the basic controller information.

Examples

1.
```
CLI> SHOW THIS_CONTROLLER
Controller:
        HS1CP    (C) DEC ZG33400026 Firmware XMDV-0, Hardware 0000
        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)
```

Shows the basic HS1CP controller information.

SHOW UNITS

Shows all units and unit information.

Format

SHOW UNITS

Description

The SHOW UNITS command displays all the units known by the controller. First disks (including CD–ROMs) 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.
```
CLI> SHOW UNITS
MSCP unit                                          Uses
-------------------------------------------------------------

  D401                                             R0
  T41                                              TAPE200
```

Shows a basic listing of units available on the controller.

2.
```
CLI> 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
        Size: 2050353 blocks
  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 controller.

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.
```
CLI> 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
        Size: 2050353 blocks
```

Shows a listing of a specific disk unit.

2.
```
CLI> 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 controller.

Format

SHUTDOWN OTHER_CONTROLLER

Description

The SHUTDOWN OTHER_CONTROLLER command flushes all user data from the other controller's write-back cache (if present), then shuts down the other controller.

If any disks are online to the other controller, the controller does not shut down unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the controller does not shut down unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes the other controller 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 controller is not shut down unless IGNORE_ERROR is specified.

If the IGNORE_ERRORS qualifier is specified, the controller 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_SHUTDOWN
NOIMMEDIATE_SHUTDOWN (Default)
If IMMEDIATE is specified, the controller 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 controller, the controller is not shutdown unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the controller 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. `CLI> SHUTDOWN OTHER_CONTROLLER`

   Shuts down the other controller as long as the other controller does not have any units online.

2. `CLI> SHUTDOWN OTHER_CONTROLLER OVERRIDE_ONLINE`

   Shuts down the other controller even if there are units online to the other controller.

SHUTDOWN THIS_CONTROLLER

Shuts down and does not restart this controller.

Format

SHUTDOWN THIS_CONTROLLER

Description

The SHUTDOWN THIS_CONTROLLER command flushes all user data from this controller's write-back cache (if present), then shuts down this controller.

If any disks are online to this controller, the controller does not shut down unless the OVERRIDE_ONLINE qualifier is specified. If any user data cannot be flushed to disk, the controller does not shut down unless the IGNORE_ERRORS qualifier is specified.

Specifying IMMEDIATE causes this controller 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 controller is lost when this controller shuts down.

Qualifiers

IGNORE_ERRORS
NOIGNORE_ERRORS (Default)
If errors result when trying to write user data, the controller is not shut down unless IGNORE_ERROR is specified.

If the IGNORE_ERRORS qualifier is specified, the controller 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_SHUTDOWN
NOIMMEDIATE_SHUTDOWN (Default)
If IMMEDIATE is specified, the controller 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 controller, the controller is not shutdown unless OVERRIDE_ONLINE is specified.

If the OVERRIDE_ONLINE qualifier is specified, the controller 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. `CLI> SHUTDOWN THIS_CONTROLLER`

   Shuts down this controller as long as this controller does not have any units online.

2. `CLI> SHUTDOWN THIS_CONTROLLER OVERRIDE_ONLINE`

   Shuts down this controller even if there are units online to this controller.

UNMIRROR *disk-device-name*

Converts a one member mirrorset back to a single device and deletes the mirrorset from the list of known mirrorsets.

Format

UNMIRROR *disk-device-name*

Description

Allows you to convert a mirrorset with one disk, specified by the disk name, *disk-device-name*, to a physical device. This command can be used on mirrorsets that are already members of higher level containers (stripesets or units).

Examples

1.
```
CLI> UNMIRROR DISK130
```

   Converts a DISK130 back to a single device. An error message prints if the disk is not a mirrorset.

E.2 CLI Messages

The following sections describe messages you can encounter during interactive use of the CLI.

E.2.1 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.

E.2.2 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.

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 THIS_CONTROLLER supports without restarting the controller, 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 controller.

Error 1160: A restart of OTHER_CONTROLLER is required before units may be added

Explanation: You changed the target IDs that the OTHER_CONTROLLER supports without restarting the controller, 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 controller.

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 controller 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 3020: Cannot add <disk name>; This causes too many members in MIRRORSET <mirrorset name>. Use a SET <mirrorset name> MEMBERSHIP=<new count> to command increase the membership count first

Explanation: You can not add more disks to a mirrorset than is specified by the mirrorset's membership count. In order to successfully do a SET <mirrorset name> REPLACE=<disk name>, to a "full" mirrorset, you must first increase the mirrorset count.

Error 3030: <disk name> is not a member of <mirrorset name>, cannot be used for the read source

Explanation: When specifying a specific disk as a read source for a mirrorset, the disk specified must be a member of the mirrorset.

Choose a disk that's a member of the mirrorset and retry the command.

Error 3040: Cannot set read source to a specific device if mirrorset not configured as a unit

Explanation: Setting a mirrorset's read source to a specific device cannot be accomplished if the mirrorset is not configured as a unit.

Create a unit from the mirrorset and try the command again.

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 controller

Explanation: You attempted to use a feature that requires a license, and the license was not enabled on this controller.

Error 3100: <licensable feature> support is not enabled on other controller

Explanation: You attempted to use a feature that requires a license, and the license was not enabled on the other controller.

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 controller 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 controller 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 or mirrorset that is not configured as a unit

Explanation: A REPLACE operation may not be done on a RAIDset or MIRRORset 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> may not 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 or mirrorset that is not configured as a unit

Explanation: A RAIDset or MIRRORset must be configured as a unit before a disk may be removed.

Create a unit from the RAIDset or MIRRORset 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 3270: <disk name> not in NORMAL state. Only NORMAL state units may be specified as a read source

Explanation: You may not specify a MIRRORset member as a read source unless it's in NORMAL state.

Either wait for the desired disk to enter NORMAL state or choose another disk in the MIRRORset that's already in NORMAL state.

Error 3280: Cannot determine if <disk name> is in the NORMAL state

Explanation: Unable to determine at this point in time if the disk specified is in normal state.

Retry the command. If this error persists, call Digital Multivendor Customer Services.

Error 3290: Mirrorsets are limited to 6 members maximum

Explanation: A MIRRORset may only have 6 disks as members. If you try to exceed this amount the above message is printed.

Error 3300: Currently there are <current count> members in this mirrorset. You must specify a membership greater than or equal to <current count>.

Explanation: You attempted to set the number of MIRRORset members to less than the actual number of disk drives that make up this MIRRORset.

Either specify a number greater than or equal to the number of physical disk drives that make up this MIRRORset or remove one or more disk disk drives, then reduce the member count.

Error 3310: No other switches may be specified on a MEMBERSHIP operation

Explanation: When specifying the number of members in a MIRRORset, no other switches may be specified.

Retry the command only specifying the membership count.

Error 3320: Unable to change membership on mirrorset <mirrorset name>

Explanation: Unable to set the membership count on the specified MIRRORset at this point in time.

Retry the command. If this error persists, call Digital Multivendor Customer Services.

Error 3330: Can only MIRROR disks. <disk name> is not a disk

Explanation: Only disks may be MIRRORed. If you specify a device other than a disk drive, the above error is printed.

Try the command again specifying a disk drive.

Error 3340: Can only UNMIRROR disks. <disk name> is not a disk

Explanation: Only disks may be unMIRRORed. If you specify a device other than a disk drive, the above error is printed.

Try the command again specifying a disk drive.

Error 3350 <disk name> must be configured under a mirrorset to UNMIRROR

Explanation: Only disk drives that are configured under mirrorsets may be unMIRRORed.

Specify a disk that is configured under a MIRRORset.

Error 3360: To UNMIRROR a disk, the mirrorset must contain only one member, the disk to be UNMIRROR. Mirrorset <mirrorset name> contains more than one member

Explanation: To unmirror a disk drive, only one disk may be a member of the MIRRORset.

REMOVE all disk drives but one and retry the command.

Error 3370: A MIRROR may not be done on a disk that is not configured as a unit

Explanation: To MIRROR a disk drive either it or the storageset that it belongs to must be configured as a unit.

Configure the disk drive or the storageset that it belongs to as a unit and retry the command.

Error 3380: An UNMIRROR may not be done on a disk that is not configured as a unit

Explanation: To UNMIRROR a disk drive either it or the storageset that it belongs to must be configured as a unit.

Configure the disk drive or the storageset that it belongs to as a unit and retry the command.

Error 3390: REDUCE of specified disks failed

Explanation: The requested REDUCE did not complete successfully.

Check the configuration to assure that it was not partially completed (some disks removed from the specified MIRRORset(s); if so, reconfigure and wait for all disks to return to the NORMAL state. If this error persists, call Digital Multivendor Customer Services.

Error 3400: <device name> is not a disk. Can only REDUCE disks

Explanation: Only disk drives may be specified on the REDUCE command.

Retry the command only specifying disks.

Error 3410: <disk name> was specified twice in the REDUCE command

Explanation: The disk drive name called out was specified twice in the REDUCE command line.

Retry the command with only specifying disk names once.

Error 3420: <disk name> is not part of the same unit as previous disks specified

Explanation: All disks specified on the REDUCE command must be used by one common unit. If the disks specified are in use by more than one unit, the above message is printed.

Retry the command specifying disks that all have one common unit as a parent.

Error 3430: A REDUCE may not be done on disks not configured as a unit

Explanation: A REDUCE command is only valid when the storageset is configured as a unit.

Configure the top most storageset as a unit and retry the command.

Error 3440: <disk name> is not a member of a mirrorset

Explanation: The disk drive specified is not a member of a MIRRORset.

Retry the command only specifying disks that are members of MIRRORsets.

Error 3450: No NORMAL or NORMALIZING members would be left in MIRRORSET <mirrorset name>

Explanation: When reducing, at least one member that is in the NORMAL state must remain in each MIRRORset. This error results when the MIRRORset called out would not be left with at least one NORMAL member.

Make sure at least one NORMAL member remains in each MIRRORset specified and retry the command.

Error 3460: !AC cannot be mirrored due to bad cache on <controller> controller

Explanation: In order to mirror a disk drive the writeback cache must be correctly operating on this, or if in a dual-redundant controller, both controllers. If it's not, the above error results.

Repair the cache on the problem controller(s) and retry the command.

Error 3470: Only 20 mirrorsets and raidsets total can be supported on a controller. This storageset not added.

Explanation: The total number of RAIDsets and MIRRORsets supported on the controller is 20; this error results if you attempt to add a 21st MIRROR or RAIDset.

Delete unused RAID and MIRRORsets and retry the command.

Error 3480: Only 30 stripesets, mirrorsets and raidsets total can be supported on a controller. This storageset not added.

Explanation: The total number of RAIDsets, MIRRORsets and STRIPEsets supported on the controller is 30; this error results if you attempt to add an 21st MIRROR, STRIPE or RAIDset.

Delete unused RAID, STRIPE and MIRRORsets 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: Controller 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 controller's ID, set MAX_NODES first, then the controller'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 controller 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 controller's ID or the controller's ID was specified with a number greater than MAX_NODES—1.

If decreasing MAX_NODES, set the controller's ID first, then MAX_NODES.

Error 4090: Module has invalid serial number. This controller cannot be used Call field service

Explanation: This error is typically the result of faulty Non-Volatile memory. This error cannot be fixed in the field.

A replacement controller must be ordered. Contact Digital Multivendor Customer Services.

Error 4100: Unable to RESTART other controller.

Explanation: A communication error occurred when trying to restart the other controller.

Retry the RESTART command.

Error 4110: Unable to SHUTDOWN other controller.

Explanation: A communication error occurred when trying to shutdown the other controller.

Retry the SHUTDOWN command.

Error 4120: Unable to SELFTEST other controller.

Explanation: A communication error occurred when trying to self-test the other controller.

Retry the SELFTEST command.

Error 4130: Unable to setup controller restart.

Explanation: A communication error occurred when trying to restart or self-test the other controller.

Retry the RESTART or SELFTEST command.

Error 4140: Unable to lock the other controller's NV memory

Explanation: Most configuration commands such as ADD, DELETE, and SET require both controllers in a dual-redundant configuration to be running so configuration changes can be recorded in both controllers. If one controller is not running, the above message results when you attempt to change the configuration.

Restart the other controller and try the command again, or SET NOFAILOVER on the remaining controller.

Error 4150: Unable to rundown the following units on the other controller: <list of problem units>

Explanation: When attempting to shut-down, restart or selftest the other controller, 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 controller: <list of problem units>

Explanation: When attempting to SHUTDOWN, RESTART or SELFTEST this controller, 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 controller ID may be changed. All unit numbers must be in the range(s): <start> to <end>

Explanation: You attempted to change the controller 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 is no longer 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

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 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 controller.

Explanation: There was a communication problem with the other controller. This typically happens if the other controller is shutting down. If these messages happen often when the other controller is not shutting down, call Digital Multivendor Customer Services.

Error 6010: Other controller not present

Explanation: When asked to communicate with another controller (the result of any one of a number of commands), the other controller was found not to be running.

If the other controller is in the process of restarting, retry the command later. If the other controller is shut down or turned off, start it. If the other controller 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 controllers must communicate to set up a dual-redundant mode. This setup time is typically less than 1 minute. If commands that require controller-to-controller communication are entered during this setup time, error 6020 results.

Retry the command later.

Error 6030: Unable to communicate with the other controller to setup FAILOVER

Explanation: Could not setup FAILOVER due to communication problems between the controllers.

Retry the command later.

Error 6040: The write of the other controller's configuration information did not succeed; information may be in an inconsistent state. Before further use both controllers 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 controller while putting controllers 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 controller not configured for dual-redundancy

Explanation: A command was entered to a single controller configuration that requires two controllers to be in dual-redundant mode.

If two controllers 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 controller not currently in dual-redundant mode

Explanation: A command was entered to a dual-redundant-configured controller, but the other controller was not available for communication.

Restart the other controller and wait until it is communicating with this controller. If this controller 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 controllers when in a dual-redundant configuration. When in a single controller 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 controllers are 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: Controllers already configured for failover

Explanation: A SET FAILOVER cannot be issued on a controller already in failover.

Error 6130: RAID5 in use on this controller but not enabled on the other controller.

Explanation: When trying to SET FAILOVER, it was discovered that there were RAID5 configurations on this controller but the other controller did not have the RAID5 feature enabled.

If RAID5 is licensed on the other controller, enable it. If it is not licensed, either contact Digital Multivendor Customer Services for licensing information, or do not use the two controllers in dual-redundant mode, or do not use a RAID5 configuration.

Error 6140: Writeback cache in use on this controller but not enabled on the other controller.

Explanation: When trying to SET FAILOVER, it was discovered that there were write-back cache switches set on this controller but the other controller did not have the write-back cache feature enabled.

If write-back cache is licensed on the other controller, enable it. If it is not licensed, either contact Digital Multivendor Customer Services for licensing information, or do not use the two controllers in dual-redundant mode, or do not use the write-back cache switches.

Error 6150: RAID5 in use on other controller but not enabled on this controller

Explanation: When trying to SET FAILOVER, it was discovered that there were RAID5 configurations on the other controller but this controller did not have the RAID5 feature enabled.

If RAID5 is licensed on this controller, enable it. If it is not licensed, contact Digital Multivendor Customer Services for licensing information or do not use the two controllers in dual-redundant mode, or do not use a RAID5 configuration.

Error 6160: Writeback cache in use on other controller but not enabled on this controller.

Explanation: When trying to SET FAILOVER, it was discovered that there were write-back cache switches set on the other controller but this controller did not have the write-back cache feature enabled.

If write-back cache is licensed on this controller, enable it. If it is not licensed, either contact Digital Multivendor Customer Services for licensing information or do not use the two controllers in dual-redundant mode, or do not use the write-back cache switches.

Error 6170: An <controller type> and <controller type> cannot configured for failover

Explanation: Two different controllers cannot be configured for failover.

Replace the other controller with the same model as this one and reissue the command.

Error 6180: MIRRORING in use on this controller but not enabled on the other controller

Explanation: When trying to SET FAILOVER, it was discovered that there were MIRRORsets on this controller but the other controller did not have the MIRROR feature enabled.

If MIRROR is licensed on the other controller, enable it. If it is not licensed, either contact Digital Multivendor Customer Services for licensing information, or do not use the two controllers in dual-redundant mode, or do not use a MIRRORset configuration.

Error 6190: MIRRORING in use on the other controller but not enabled on this controller

Explanation: When trying to SET FAILOVER, it was discovered that there were MIRRORsets on the other controller but this controller did not have the MIRROR feature enabled.

If MIRROR is licensed on this controller, enable it. If it is not licensed, contact Digital Multivendor Customer Services for licensing information or do not use the two controllers in dual-redundant mode, or do not use a MIRRORset configuration.

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 controller 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 remains unavailable until the controller is restarted.

Error 7060: Cannot clear LOST_DATA on a unit without LOST_DATA

Explanation: If a unit does not have a LOST_DATA error, a CLEAR LOST_DATA <unit number> is an illegal command.

Check to assure the unit that you wished to CLEAR LOST_DATA on does exhibit a LOST_DATA error.

Error 7070: Cannot clear UNWRITEABLE_DATA on a unit without UNWRITEABLE_DATA

Explanation: If a unit does not have a UNWRITEABLE_DATA error, a CLEAR UNWRITEABLE_DATA <unit number> is an illegal command.

Check to assure the unit that you wished to CLEAR UNWRITEABLE_DATA on does exhibit an UNWRITEABLE_DATA error.

Error 7080: Cannot retry UNWRITEABLE_DATA on a unit without UNWRITEABLE_DATA

Explanation: If a unit does not have a UNWRITEABLE_DATA error, a RETRY UNWRITEABLE_DATA <unit number> is an illegal command.

Check to assure the unit that you wished to RETRY UNWRITEABLE_DATA on does exhibit a UNWRITEABLE_DATA error.

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.

Note

This error is commonly the result of failing to name a storageset. For example, when trying to create a stripeset out of three disks the user may type:

```
ADD STRIPESET DISK100 DISK200 DISK300
```

This results in a 9080 error since the controller attempted to create a stripeset named "DISK100" (since the user forgot to name the stripeset), which of course is a name already used by a disk drive.

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 CD–ROMS 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 nondisk 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 controller

Explanation: When trying to check for online drives on the other controller, there was a communication failure.

Retry the command.

Error 9220: You cannot specify NOREAD_CACHE while in WRITEBACK_ CACHE mode

Explanation: WRITEBACK_CACHE is specified for this unit. READ_ CACHE cannot be specified while WRITEBACK_CACAHE is set.

Reissue the SET command disabling both WRITEBACK_CACHE and READ_ CACHE.

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: Illegal command—this controller not configured for failover, but other controller present. Stop one controller or SET FAILOVER to put both in failover mode

Explanation: You may not change a non-dual-redundant controller's configuration if the other controller is running.

Stop the other controller (pop out the program card, for example) and retry the command.

Error 9280: Cannot rename the SPARESET or FAILEDSET

Explanation: The SPARESET and FAILEDSET cannot be renamed.

Error 9290: Communication failure with other controller, cannot check other controller's licensing

Explanation: Unable to communicate with the other controller to check licensing before creating a RAIDset or enabling write-back cache.

Check to make sure that both controllers are running. If one is broken, take this controller out of failover (SET NOFAILOVER) and reissue the command.

Warning 9300: Bad write cache or battery on <controller> controller

Explanation: When trying to set failover, it was discovered that one controller's cache had failed. For this reason the ability to place the controller's in a dual redundant mode has been blocked.

Repair the cache and retry the command.

Error 9310: No metadata found on container, unit not created. <reason for failure>

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 9330: NV memory write collision. Please try again

Explanation: Two processes were trying to modify the controller's configuration 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 controller.

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 controller 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 controller

Explanation: When trying to set failover a unit with unflushed data or a cache error was detected on this controller.

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 controller

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 controller

Explanation: When trying to set failover a unit with unflushed data or a cache error was detected on the other controller.

Delete the unit as requested and then retry the SET FAILOVER command.

Error 9450: Cannot mirror <disk name>

Explanation: Error when attempting to mirror a disk drive.

Check the configuration to assure it's the same as before. If so, retry the command. If this error persists, call Digital Multivendor Customer Services.

E.2.3 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 results in an unusable configuration; you 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.

E.2.4 CLI Warning Messages

Warning 1000: It is recommended that you read the controller product documentation to understand the significance of enabling WRITEBACK_ CACHE particularly for RAID Arrays

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 controller'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 works, it is strongly recommended that the storageset be deleted and reconfigured with one member per port.

Warning 3020: <storageset name> is configured with different sized containers. This results in a storageset of reduced size

Explanation: This warning results from an ADD *storageset-type* command. 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, the size of all devices that make up the storageset should be identical.

Warning 3030: Cannot determine state of disk <disk name>in mirrorset <mirrorset name>

Explanation: Before disk drives are removed from a MIRRORset, the MIRRORset is checked to assure that at least one NORMAL member is left after the REMOVE or REDUCE. If the state of a disk cannot be determined, the above message is printed.

Warning 3040: <storageset name> is disabled due to bad cache state on <controller> controller

Explanation: If the cache is in a state on this or the other controller such that the storageset cannot be used, it is created but the above message is printed so the user knows that the storageset is disabled.

Warning 4000: A restart of this controller is required before all the parameters modified take effect.

Explanation: This warning results from a SET THIS_CONTROLLER command. Some controller 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 controller is required before all the parameters modified take effect.

Explanation: This warning results from a SET OTHER_CONTROLLER command. Some controller parameters require a restart before they can take effect. If any of those parameters are changed, this warning is displayed.

Restart the controller and retry the command.

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

Explanation: This warning results from a SET THIS_CONTROLLER or a SET OTHER_CONTROLLER command. Some controller parameters require a restart of both controllers before they can take effect. If any of those parameters are changed, this warning is displayed. Restart both controllers and retry the command.

Warning 6000: Communication failure with the other controller while taking controllers out of dual-redundant mode. Issue a SET NOFAILOVER command on the other controller

Explanation: This warning results from a SET NOFAILOVER command. This controller was unable to communicate with the other controller to notify it that it is no longer in dual-redundant mode. Typically, this occurs when the other controller has already been removed prior to the SET NOFAILOVER command.

Enter a SET NOFAILOVER command on the other controller as soon as possible.

Warning 6010: Licensing different between the two controllers

Explanation: If the licensing is not identical on both controllers in a dual-redundant configuration, the above warning is displayed.

You should check the licensing on both controllers 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 controller

Explanation: When trying to clear LOST_DATA on the other controller, 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 controller

Explanation: When trying to clear UNWRITEABLE_DATA on the other controller, a communication error occurred.

Retry the command. If the failure persists, contact Digital Multivendor Customer Services.

Warning 9000: Drive has LOST_DATA

Explanation: During a check of the drive's metadata, it was detected that the drive had lost data.

Clear the lost data error on the drive.

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 nondisk 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 <controller> controller. 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 controller 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 <controller> writeback cache is not used.

Explanation: The battery is low or bad on the specified controller. The unit specified does not use write-back cache until the battery is charged or repaired.

Warning 9110: Bad or low battery or bad write cache on <controller> this unit cannot be used by <controller>

Explanation: The battery is low or bad on the specified controller. The unit specified requires the use of write-back cache, so its use has been disabled until the battery is charged or repaired.

E.3 Device Configuration Examples

The following examples cover the majority of configurations and the method of defining those configurations.

Example E–1 Initial Single Controller Configuration of an HS1CP Controller

```
CLI> SET THIS_CONTROLLER ID=5 SCS_NODENAME="HS1CP1"
CLI> SET THIS_CONTROLLER MSCP_ALLOCATION_CLASS=4 TMSCP_ALLOCATION_CLASS=4
CLI> RESTART THIS_CONTROLLER
[this controller restarts at this point]
CLI> SET THIS_CONTROLLER PATH
```

Example E–2 Setting the Terminal Speed and Parity

```
CLI> SET THIS_CONTROLLER TERMINAL_SPEED=19200 NOTERMINAL_PARITY
```

Note

Garbage appears on the terminal after setting the controller's terminal speed until you set the terminal's speed to match the controller's new terminal speed.

Example E–3 Creating a Unit from a Disk Device

```
CLI> ADD DISK DISK0 2 0 0
CLI> INITIALIZE DISK0
CLI> ADD UNIT D0 DISK0
```

Example E–4 Creating a Unit from a Tape Device

```
CLI> ADD TAPE TAPE0 3 0 0
CLI> ADD UNIT T0 TAPE0
```

Example E–5 Creating a Unit from a Four-Member Stripeset

```
CLI> ADD DISK DISK0 1 0 0
CLI> ADD DISK DISK1 2 0 0
CLI> ADD DISK DISK2 3 0 0
CLI> ADD DISK DISK3 1 1 0
CLI> 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
CLI> INITIALIZE STRIPE0
CLI> ADD UNIT D0 STRIPE0
```

Example E–6 Creating a Unit from a Five-Member RAIDset

```
CLI> ADD DISK DISK0 1 0 0
CLI> ADD DISK DISK1 2 0 0
CLI> ADD DISK DISK2 3 0 0
CLI> ADD DISK DISK3 1 1 0
CLI> ADD DISK DISK4 2 1 0
CLI> 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
CLI> INITIALIZE RAID9
CLI> ADD UNIT D0 RAID9
```

Example E–7 Creating a Unit From a Disk Device and Setting the Write Protection

```
CLI> ADD DISK DISK0 2 0 0
CLI> INITIALIZE DISK0
CLI> ADD UNIT D0 DISK0 WRITE_PROTECT
```

Example E–8 Setting the Write Protection for an Existing Unit

```
CLI> ADD DISK DISK0 2 0 0
CLI> INITIALIZE DISK0
CLI> ADD UNIT D0 DISK0
CLI> SET D0 WRITE_PROTECT
```

Example E–9 Renumbering Disk Unit 0 to Disk Unit 100

```
CLI> ADD DISK DISK0 2 0 0
CLI> INITIALIZE DISK0
CLI> ADD UNIT D0 DISK0
CLI> DELETE D0
CLI> ADD UNIT D100 DISK0
```

Note

No INITIALIZE is required because DISK0 has already been initialized.

Example E–10 Creating a Transportable Unit from a Disk Device

```
CLI> ADD DISK DISK0 2 0 0 TRANSPORTABLE
CLI> INITIALIZE DISK0
CLI> ADD UNIT D0 DISK0
```

[or]

```
CLI> ADD DISK DISK0 2 0 0
CLI> SET DISK0 TRANSPORTABLE
CLI> INITIALIZE DISK0
CLI> ADD UNIT D0 DISK0
```

Example E–11 Changing the Replacement Policy of a RAIDset

```
CLI> ADD DISK DISK0 1 0 0
CLI> ADD DISK DISK1 2 0 0
CLI> ADD DISK DISK2 3 0 0
CLI> ADD DISK DISK3 4 0 0
CLI> ADD DISK DISK4 5 0 0
CLI> ADD RAIDSET RAID9 DISK0 DISK1 DISK2 DISK3 DISK4
CLI> INITIALIZE RAID9
CLI> ADD UNIT D0 RAID9
CLI> SET RAID9 POLICY=BEST_FIT
```

Note

The replacement policy can be changed at any time.

Example E–12 Deleting the Unit, Stripeset, and All Disks Associated with a Stripeset

```
CLI> DELETE D0
CLI> DELETE STRIPE0
CLI> DELETE DISK0
CLI> DELETE DISK1
CLI> DELETE DISK2
CLI> DELETE DISK3
```

F

SWEAT

F.1 Overview

HS family subsystems log errors according to the format laid out in the *HS Family of Array Controllers Service Manual*. The OpenVMS VAX ERF utility[1] does not provide bit-to-text decoding of these error log entries until OpenVMS VAX Version 6.1. Furthermore, VAXsimPLUS™ does not provide analysis support for any HS family subsystem; this support is planned for a future release of DECevent on the VAX platform.

By default, SWEAT produces output only for those error log entries that pertain to storage subsystems. It does **NOT** produce output for such error log entry types as system bugchecks or timestamps.

F.2 Features

SWEAT Version 2.6 has the following features:

- Nonprivileged code
- A summary report provided for nonverbose requirements
- Supports HSOF up to and including Version 2.5
- Supports HS1CP, HSJ30, HSJ40, HSD30, HSD05, HSD10, HSC/k.scsi, and SWXCR
- Supports the new error log format (including dynamic device recognition) as found in OpenVMS Alpha Version 6.1
- Supports Merge facility (similar to the function of VAXSIM/MERGE, but functions on Alpha and VAX error logs)
- Corrects the problem of new ERRLOG.SYS files being created if SWEAT uses the current error log file

F.3 Installing SWEAT

SWEAT is supplied as a VMS Backup save set. This save set consists of the following files:

- SWEAT.OBJ
- SWEAT_AXP.OBJ
- READ_ME.1ST
- SWEAT.PS
- SWEAT.CLD

[1] Invoked by the ANAL/ERR command.

- HSJ40.SYS
- KSCSI.SYS
- SWEAT_MESSAGES.OBJ
- SWEAT_MESSAGES_AXP.OBJ

Installation involves the following procedure:

1. Enter the following command to show your current directory:

```
$ SHOW DEFAULT
  SYS$SYSROOT:[SYSMAINT]
```

2. Create a directory for SWEAT:

```
$ CREATE/DIRECTORY [.SWEAT]
```

3. Set your default directory to the new SWEAT directory:

```
$ SET DEF [.SWEAT]
```

4. Unpack the save set to the SWEAT directory:

```
$ BACKUP MUA0:SWEAT_X26.BCK/SAV */LOG

%BACKUP-S-CREATED, created SYS$SYSROOT:[SYSMAINT.SWEAT]SWEAT.OBJ;5
%BACKUP-S-CREATED, created SYS$SYSROOT:[SYSMAINT.SWEAT]SWEAT_AXP.OBJ;1
%BACKUP-S-CREATED, created SYS$SYSROOT:[SYSMAINT.SWEAT]READ_ME.1ST;14
%BACKUP-S-CREATED, created SYS$SYSROOT:[SYSMAINT.SWEAT]SWEAT_MESSAGES.OBJ;1
%BACKUP-S-CREATED, created SYS$SYSROOT:[SYSMAINT.SWEAT]SWEAT_MESSAGES_AXP.OBJ;1
%BACKUP-S-CREATED, created SYS$SYSROOT:[SYSMAINT.SWEAT]SWEAT.CLD;6
%BACKUP-S-CREATED, created SYS$SYSROOT:[SYSMAINT.SWEAT]HSJ40.SYS;1
%BACKUP-S-CREATED, created SYS$SYSROOT:[SYSMAINT.SWEAT]KSCSI.SYS;1
```

5. Link SWEAT for either VAX or Alpha systems:

VAX Systems

Enter the following command to link the object file to produce the executable code for VAX systems:

```
$ LINK SWEAT,SWEAT_MESSAGES,SYS$INPUT:/OPT
  SYS$LIBRARY:VAXCRTL.EXE/SHARE ^z
```

The LINK command line informs the system that the *options file* will be supplied from the terminal (SYS$INPUT:). The system waits for the user to enter the following command:

```
SYS$LIBRARY:VAXCRTL.EXE/SHARE  ^z
```

Note that there is no prompt for this input.
Note also the Ctrl/Z (^Z) terminator; this is very important.

Alpha Systems

Enter the following command to link the object file to produce the executable code for Alpha systems:

```
$ LINK SWEAT_AXP,SWEAT_MESSAGES_AXP/EXE=SWEAT.EXE
```

SWEAT.CLD

Edit the SWEAT.CLD file to modify the verb's definition. Do this by entering the following command:

```
$ SHOW DEFAULT
  SYS$SYSROOT:[SYSMAINT.SWEAT] = SYS$SYSROOT:[SYSMAINT.SWEAT]
$ EDIT SWEAT.CLD
```

Make the command verb SWEAT and its qualifiers known to your process by entering the following command:

```
$ SET COMMAND SWEAT.CLD
```

The following shows the content of the SWEAT.CLD file and the modifications required when you enter the EDIT command:

```
define type dump_blocks
 keyword start value(type=$number)
 keyword end value(type=$number)
 keyword count value(type=$number)

define verb sweat
 Image DISK$USERS1:[LOANE.HSJ40]SWEAT.EXE
 Parameter p1, label=FILE, prompt="Errlog File:- ", value(list, type=$file)
 Qualifier Output, value(required, type=$file), nonnegatable
 Qualifier Config, value(required, type=$file), nonnegatable
 Qualifier Before, value(required, type=$datetime), nonnegatable
 Qualifier Since, value(required, type=$datetime), nonnegatable
 Qualifier Entry, value(required, list, type=dump_blocks), nonnegatable
 Qualifier Full,nonnegatable
 Qualifier All, nonnegatable
 Qualifier Summary,nonnegatable
 Qualifier Instance, value(required), nonnegatable
 Qualifier Last_fail, value(required), nonnegatable
 Qualifier Fsterr,nonnegatable
 Qualifier Merge, nonnegatable
 Qualifier Binary, nonnegatable, value(required, type=$file)
 Qualifier Ele_type, nonnegatable, value(required, list, type=$number)

 disallow entry.end and entry.count
 disallow instance and last_fail
 disallow merge and fsterr
 disallow merge and instance
 disallow merge and last_fail
 disallow merge and summary
 disallow merge and all
 disallow merge and full
 disallow merge and entry
 disallow merge and config
 disallow merge and output
```

In SWEAT.CLD, the line containing DISK$USERS1:[LOANE.HSJ40] must be changed to SYS$SYSROOT:[SYSMAINT.SWEAT] to match the previous example.

SWEAT is now ready for use.

F.4 Commands and Syntax

The following section provides descriptions of the command syntax used with SWEAT. The qualifiers are grouped into those that affect the filtering of error log entries and those that affect the output generated. Simple examples are given to demonstrate SWEAT's ease of use.

SWEAT

Invokes the SWEAT tool to analyze error log entries and break out the contents for further decoding.

Syntax

SWEAT input file-spec[, . . .]

Command Qualifiers	Defaults
/OUTPUT=file spec	None
/CONFIG=file spec	None
/BEFORE=time	None
/SINCE=time	None
/ENTRY=(START, END, COUNT)	None
/FULL	None
/ALL	None
/SUMMARY	None
/INSTANCE=instance_code	None
/LAST_FAIL=last_fail_code	None
/FSTERR	None
/MERGE	None
/ELE_TYPE=(entry_type,entry_type,..)	None
/BINARY=filespec	None

Restrictions

None.

Prompts

Errlog File:- Input File Spec[, . . .]

Parameters

Input file spec

It is recommended that this *not* be the system error log file, but should be a file previously merged[2] from all nodes in a cluster. This minimizes the processing required while holding the error log file open and obtains a more accurate view of the problem from a cluster point of view.

Description

The SWEAT command syntax is available by issuing the DCL SET COMMAND DISK:[DIRECTORY]SWEAT.CLD command, where DISK:[DIRECTORY] refers to the directory specification set up in the previous section. You can use SWEAT to selectively output decoded error log entries to the terminal or to a file.

[2] Using SWEAT/MERGE or VAXSIM/MERGE.

Filter Qualifiers

/BEFORE=time

Selects only those error log entries whose timestamps are before the given *time* criteria.

This time can be entered in VMS absolute time format (28–FEB–1996 14:36:21.05). Alternatively, the keywords TODAY, TOMORROW, and YESTERDAY can be used. There is no default for this qualifier.

/SINCE=time

Selects only those error log entries whose timestamps are after the given *time* criteria.

This time can be entered in VMS absolute time format (28–FEB–1996 14:36:21.05). Alternatively, the keywords TODAY, TOMORROW, and YESTERDAY can be used. There is no default for this qualifier.

/ENTRY=[START:nnn, END:mmm]

Selects only those error log entries whose entry numbers meet the criteria given.

The entry number is the relative error log record number within the errorlog file (the first ELE is Entry Number 1); it is *not* the sequence number recorded within the ELE itself.

Input Qualifiers

/CONFIG=file spec

Specifies the file that contains site-specific configuration information. This file is input to SWEAT and should be supplied by the local engineer. SWEAT uses the data to replace HEX Controller IDs or SCS Node IDs with more meaningful node names.

```
serial:000933700882,ci:0,node:HSJ00,id:420010003120
serial:000931900420,ci:1,node:HSJ01,id:420010013122
serial:000131600189,ci:1,node:HSJ01,id:420010013122
serial:000131800250,ci:2,node:HSJ02,id:420010023124
serial:000132100503,ci:5,node:HSJ05,id:42001005312a
node:HSC003,id:00000000FC03
node:HSC004,id:00000000FB04
node:SDC101,id:000000000734,ci:A
node:PEARL,id:000000000407,ci:C
node:BLACK,id:0000000004C8,ci:7
```

This is an example of a configuration file from a cluster with four HSJ40 controllers, three VAX systems, and two HSC controllers.

Note in the previous example that HSJ01 has two entries. This is because the HSJ01 controller was replaced (controller module serial# 000931900420 was replaced by controller module serial# 000131600189).

To get the correct information from the system to build the configuration file, use the following:

- **NODE**

  Obtain a list of node names by entering the SHOW CLUSTER command.

- **SERIAL**

This item is useful only for HSJ40 controllers. Obtain the serial number by entering the *SHOW CONTROLLER* command on the HSJ40's console. Then tie this up with the node name from the previous step.

- **ID**

  Obtain the SCS NodeID by entering the WRITE SYS$OUTPUT F$GETSYI("NODE_SYSTEMID", ",emphasis>(xxxxxx)") where *xxxxxx* relates to the node name in which you are interested.

- **CI**

  Obtain the CI node number for each node in the cluster by entering the following two commands:

```
$ SHOW CLUSTER/CONT
$ ADD RPORT
```

Output Qualifiers

/OUTPUT=file spec

Specifies the name of the file to which SWEAT writes the output of its analysis.

/ALL

Specifies that **ALL** error log entries get reported by SWEAT. Using this qualifier ensures that output is generated for all ELEs, not just those for which SWEAT adds value. For example, using this qualifier generates output for timestamp, system bugcheck, memory CRD, and I/O errors.

/FULL

This qualifier is not implemented for Version 2.6.

/SUMMARY

Specifies that output is restricted to a summary report of all ELEs within the error log file. Using this qualifier allows you to get a brief list of all events in the error log file being processed with corresponding Instance and/or Last Fail codes when appropriate, for example:

```
$ SWEAT/SUMM CUSTOMER.SYS
 Copyright Digital Equipment Corporation 1993,1994. All rights reserved.
 StorageWorks Errorlog Analyser Tool X2.6

  15 Logged Message Events
       3 Errs BB12VA$PAA0
       4 Errs BB12VB$PAA0
       6 Errs HSJ06$MUA1 TZ87
           6 Errs with Instance Code = 03644002
       2 Errs HSJ07$MUA4 TZ87
           2 Errs with Instance Code = 03644002
 153 ERL$LogStatus Events
     153 Errs HSJ06$MUA1
  10 Volume Mount Events
   9 Volume Dismount Events
```

/INSTANCE = instance code

This qualifier commands SWEAT to decode the given Instance code into meaningful English text, for example:

```
$ SWEAT/INSTANCE=03644002
 Copyright Digital Equipment Corporation 1993,1994. All rights reserved.
 StorageWorks Errorlog Analyser Tool V2.6

Decoded Instance Code is:
    An unrecoverable tape drive error was encountered while  performing
    work related to tape unit operations.

Repair action is:
    If the Sense Data FRU field is nonzero, follow repair  action  41.
    Otherwise,  replace the appropriate FRU associated with the device's
    SCSI interface or the entire device.
```

/LAST_FAIL = last fail code

This qualifier commands SWEAT to decode the given Last Fail code into meaningful English text, for example:

```
$ SWEAT/LAST_FAIL=01860080
 Copyright Digital Equipment Corporation 1993,1994. All rights reserved.
 StorageWorks Errorlog Analyser Tool V2.6

Decoded LastFail Code is:
    A processor interrupt was generated with an indication that the
    other controller in a dual controller configuration asserted the
    KILL line to disable this controller.

Repair action is:
    No action necessary.
```

/FSTERR

Specifies that SWEAT should report the contents of the specified error log file as one line per error log entry. The actual information output depends on the ELE type, but the intent is to provide a precis of the error log for further, detailed investigation. This switch would be useful *after* having used the /SUMMARY switch in homing in on device problems, for example:

```
$ SWEAT/FSTERR/CONF=CUSTOMER_CONF.TXT/ALL CUSTOMER.SYS
 Copyright Digital Equipment Corporation 1993,1994. All rights reserved.
 StorageWorks Errorlog Analyser Tool X2.6

Entry|   Date/Time    | Node |  ELE Type     |   Device    | Cntrlr ID |Instance|LastFail| Miscellaneous
=====|================|======|===============|=============|===========|========|========|==============================
    1 01SEP94 00:03:34 QUEST   Time Stamp
    2 01SEP94 00:06:02 VEEFER  Logged Msg      VEEFER$PAA0                                  Loc ID = 07, Rem ID = 0B
    3 01SEP94 00:06:02 QUEST   Logged Msg      QUEST$PAA0                                   Loc ID = 05, Rem ID = 0B
   .
   .
   .
   28 01SEP94 08:38:49 VEEFER  Logged MSCP Msg              THOR        4007640A          Src ID = 0B, Dst ID = 06
   29 01SEP94 08:38:50 QUEST   Logged MSCP Msg              THOR        4007640A          Src ID = 0B, Dst ID = 06
   30 01SEP94 09:16:02 VEEFER       Time Stamp
   .
   .
   .
  268 01SEP94 18:43:43 VEEFER  Logged MSCP Msg              ZEUS        01010302 01880080
  269 01SEP94 18:43:43 VEEFER  Logged MSCP Msg              THOR        01010302 01880080
```

Note that in the previous example, the /CONFIG qualifier has translated the "Cntrlr ID" field into a more meaningful name.

Examples

1. `$ SWEAT MEL.SYS/OUT=MEL.LOG/BEFORE/SINCE=YESTERDAY/CONFIG=CUSTOMER_CONFIG.TXT`

   This command opens the binary error log file MEL.SYS for input and decodes any error log entries timestamped yesterday generating the output file MEL.LOG as it goes. It also opens CUSTOMER_CONFIG.TXT for input as the configuration file, extracting node name information as it goes.

2. `$ SWEAT MEL.SYS/ENTRY=(START:10,END:19)`

   This command opens MEL.SYS for input, decodes error log entry numbers 10 through 19, and displays the information on the terminal.

Merge Qualifiers

/MERGE
This qualifier commands SWEAT to merge error log files in a similar manner to the VAXSIM/MERGE function.

Note

The /ELE_TYPE and /BINARY qualifiers are valid only when used in conjunction with the /MERGE qualifier\bold.)

The following example illustrates this qualifier's use:

```
$ SWEAT/MERGE
```

This command opens (for input) all files referenced by the logical *SWEAT$ERRLOGS*, opens (for output) the default file *SMEL.SYS*, and starts loading records chronologically from the input files into the output file. To mimic the function of VAXSIM/MERGE, define the logical SWEAT$ERRLOGS as SYS$SYSDEVICE:[SYS*.SYSERR]ERRLOG.SYS, such that the following applies:

```
$ SHOW LOGICAL SWEAT$ERRLOGS
"SWEAT$ERRLOGS" = "SYS$SYSDEVICE:[SYS*.SYSERR]ERRLOG.SYS"
```

Use the /BINARY qualifier to specify an output filespec *other* than SMEL.SYS.

Use the /BEFORE and /SINCE qualifiers to select a specific time frame of interest.

Use the /ELE_TYPE qualifier to specify (in decimal) specific ELE entry types to be processed.

/BINARY
The /BINARY qualifier used in conjunction with the /MERGE qualifier causes SWEAT to open for output the specified filespec (the default filespec is SMEL.SYS.), for example:

```
$ SWEAT/MERGE/BIN=MERGED_ERRLOG.SYS
```

This command merges all error log files referenced by SWEAT$ERRLOGS and creates the output file MERGED_ERRLOG.SYS. If SWEAT finds no ELEs that match the criteria of the MERGE command syntax, the output file (which is empty) is deleted.

/ELE_TYPE

This qualifier is used with the /MERGE qualifier to select specific ELE entry types for processing. The ELE_TYPES are entered as a comma-separated list of decimal entry types. These entry types are defined in the macro $EMBETDEF found in LIB.MLB.

```
$ SWEAT/MERGE/ELE_TYPE=(1,98,100)
```

This command selects only Device Error ELEs (1), Device Attention ELEs (98), and Logged Message ELEs (100) for processing to the default SMEL.SYS.

For details of the valid ELE_TYPE codes, see ELE_TYPE Codes.

F.5 Error Log Collation Techniques—Overview

In an OpenVMS VAX system, errors are logged in response to a device driver's call to ERL$ALLOCEMB or ERL$DEVICEATTN, [3] (requesting an error message buffer), followed by a call to ERL$RELEASEMB (to release the error message buffer). Once released, the error message buffer is written to the error log file by the ERRFMT process every 30 seconds (or when a request to ERL$ALLOCEMB fails due to a buffer being full).

Between the calls to ERL$ALLOCEMB (or ERL$DEVICEATTN) and ERL$RELEASEMB, the error message buffer is written by the driver with the data relating to the error.

The system maintains a unique counter (ERL$GL_SEQUENCE) called the sequence number. This is incremented for each error log prior to the attempt to allocate an EMB. This can be used to spot a high error rate (where the sequence numbers have "holes"). However, note that these sequence numbers are CPU node-specific.

Within a cluster, each node has its own error logging system and its own error log file. Attempting to build a picture of some failure event on a cluster can be hampered by the various mechanisms used logging errors. For example, the *SHOW ERROR* command on a particular Cluster node could show no significant errors, while on another node it could show thousands. This can be due to one of two reasons:

- Because the *SHOW ERROR* command reports the error count in a device's unit control block (UCB), if a node has been recently rebooted, few errors may have been logged.
- A single cluster node with a high error count may indicate high activity from that specific node. If this is the case, errors are logged back only to the node *consuming* the data.

Errors can be logged via the error logging datagram mechanism to *ALL* nodes simultaneously, and the picture of exactly what happened can become unclear.

F.5.1 Error Log Collation

The following process is helpful in quickly isolating problems from clusterwide error logs by enabling the engineer to build a time line of events spanning the period during which the symptoms were present:

1. Determine the time span over which an event has been evident.
2. Collate *ALL* error logs from the cluster for that time period into one binary file.
3. Filter this further by selectively discarding uninteresting events (like timestamps, volume (dis)mounts, and so forth).
4. Use *SWEAT* to provide readable text, either via a summary report or directly to a verbose (default) report.

[3] ERL$LOGSTATUS, ERL$LOGMESSAGE, ERL$LOG_DMSCP, and ERL$LOG_TMSCP also are used to log specific events.

F.5.2 Determining the Time Span of the Symptoms

Probably the easiest way to determine the time span of the symptoms is to use SYSMAN and have each cluster node generate a summary report, which should pinpoint the area of interest, for example:

```
$ MC SYSMAN
SYSMAN> SET ENVIRON/CLUSTER
%SYSMAN-I-ENV, current command environment:
            Clusterwide on local cluster
            Username FIELD will be used on nonlocal nodes

SYSMAN> DO DIRECTORY/DATE/SIZE SYS$ERRORLOG:ERRLOG.SYS
```

This command shows the creation dates of each cluster member's error log file and should be checked for validity.

```
SYSMAN>  DO ANALYZE/ERROR/SUMMARY/NOFULL SYS$ERRORLOG:ERRLOG.SYS
```

This command produces an output from each cluster member showing the span of errors and devices indicted, for example:

```
DEVICE SUMMARY LOGGED BY SID 12000003

                    ERRORS          TIMEOUTS        UCB ERROR   UCB OPERATION
                  [HARD]  [SOFT]   [HARD]  [SOFT]     COUNT         COUNT

_HSJ001$DUA100:
  ""
                     0.     633.      0.      0.         2014.        55522.
                 ------- -------  ------- -------  ------------ ------------
    TOTALS           0.     633.      0.      0.         2014.        55522.

_CORP01$PAC0:
  ""
                     0.      15.      0.      0.            0.           0.
                 ------- -------  ------- -------  ------------ ------------
    TOTALS           0.      15.      0.      0.            0.           0.

_HSJ002$DUA100:
  ""
                     0.     423.      0.      0.         2085.        55809.
                 ------- -------  ------- -------  ------------ ------------
    TOTALS           0.     423.      0.      0.         2085.        55809.

_HSJ002$DUA140:
  ""
                     0.       1.      0.      0.            0.           0.
                 ------- -------  ------- -------  ------------ ------------
    TOTALS           0.       1.      0.      0.            0.           0.

SUMMARY OF ALL ENTRIES LOGGED BY SID 12000003

      ERRLOG.SYS CREATED                    3.
      TIME-STAMP                            3.
      VOLUME MOUNT                         11.
      VOLUME DISMOUNT                      24.
      ERL$LOGSTATUS                       352.
      ERL$LOGMESSAGE                      720.
      ERL$LOGMSCP                          27.

      DATE OF EARLIEST ENTRY          7-FEB-1994 15:21:58.86
      DATE OF LATEST ENTRY            7-FEB-1994 15:53:01.72

PROCESSED ENTRIES HOUR-OF-DAY HISTOGRAM LOGGED BY SID 12000003
```

```
00:00      0.
01:00      0.
02:00      0.
03:00      0.
04:00      0.
05:00      0.
06:00      0.
07:00      0.
08:00      0.
09:00      0.
10:00      0.
11:00      0.
12:00      0.
13:00      0.
14:00      0.
15:00   1140. **************************************************
16:00      0.
17:00      0.
18:00      0.
19:00      0.
20:00      0.
21:00      0.
22:00      0.
23:00      0.
```

Errorlog collation around 15:00 appears appropriate.

F.5.3 Collating Error Logs to One File

You can collate error logs to one file by using the SWEAT/MERGE command (new to Version 2.6) or by using VAXSIM/MERGE. The command syntax for SWEAT/MERGE is described in Section F.4; the following describes the process when using VAXSIM/MERGE.

This process uses the merge error log[4] feature of VAXsimPLUS. Merging the error logs generates a single error log file from multiple input files and discards duplicate entries (clusterwide error logging datagrams).

Use the following process to determine whether VAXsimPLUS is installed, and, if so, to collate the error logs:

```
$ SHOW LOGICAL VAXSIM$ERRLOGS
"VAXSIM$ERRLOGS" = "SYS$SYSDEVICE:[SYS*.SYSERR]ERRLOG.SYS"
```

This command shows that VAXsimPLUS is installed and shows you what search criteria can be used for selecting error log files to merge.

```
$ DIRECTORY/SIZE/DATE VAXSIM$ERRLOGS:
```

Use this command to ensure from the output that the target error log files contain data pertinent to the time-span of symptoms as determined in Section F.5.2.

```
$ VAXSIM/MERGE/BEFORE=7-FEB-1994:16:30/SINCE=7-FEB-1994:14:30/BIN=FILTER.SYS
```

This command Merges the Errorlog files pointed to by the VAXSIM$ERRLOGS logical (the default if no "P1" parameter is specified), selecting only those ELEs whose timestamps match the selection, discarding duplicates, and outputting the binary data to FILTER.SYS.

[4] FMGR$MEL.EXE accessed via the VAXSIM/MERGE command.

F.5.4 Filtering the Error Log

After obtaining the merged error log file as described in Section F.5.3, you may need to further filter out unrelated error log information. Enter the following command to determine what the merged error log file contains, as well as what information is unrelated:

```
$ ANALYZE/ERROR/SUMM/NOFULL FILTER.SYS
```

This command produces a summary report (as in Section F.5.2) outlining which error log events have been captured and which devices are responsible. If you decided, for example, that you were not interested in timestamps, volume (dis)mounts, and magtape error log entries, use the following command to filter them out:

```
$ ANALYZE/ERROR/EXCLUDE=(CONT,VOL,TAPE)/BIN=MORE_FILTER.SYS FILTER.SYS
```

This command generates MORE_FILTER.SYS; use this as input to Section F.5.5.

F.5.5 Using SWEAT to Break Out the Error Log

Use SWEAT to break out the error log file to text and perform whatever *Recommended Action* is called for. To make the output more meaningful, remember to use a configuration file so that SWEAT recognizes and annotates the controllers to the output data.

ELE_TYPE Codes

The codes listed in Table F–1 are valid for the /ELE_TYPE qualifier, and are extracted from the $EMBETDEF Macro in LIB.MLB.

Table F–1 ELE_TYPE Values (In Decimal)

Code	Error Type	Code	Error Type	Code	Error Type	Code	Error Type
01	Device Error	02	Machine Check	03	Syndrome Error	04	Bus Error
05	SBI Alert	06	Soft ECC Error	07	Async Write Error	08	Hard ECC Error
09	Unibus Adapter Error	10	Fault thru SBI Vector	11	Unibus Error	12	Massbus Adapter Error
13	SBIA Error	14	CRD Log	15	Environ. Monitor	16	CPU Error Halt
17	Console Reboot	18	BI Adapter Error	19	BI Bus Error	20	NMI Fault
21	Console Timeout	22	NBW	23	Cache Error	24	CVAX Cache/Bus Error
25	Memory Error Found by Scanning	26	SCB 54 Error	27	SCB 60 Error	28	Adapter Error
29	Calypso Lastfail	30	Console Entry	31	Info Entry	32	Cold Start
33	Clock Module Error	34	SCAN Error	35	New File Created	36	Warm Start
37	Crash Re-start	38	Time Stamp Entry	39	System Service Message	40	System Bugcheck

(continued on next page)

Table F–1 (Cont.) ELE_TYPE Values (In Decimal)

Code	Error Type	Code	Error Type	Code	Error Type	Code	Error Type
41	Operator Message	42	Network Message	43	Configuration Entry	44	Polled Error Entry
64	Volume Mount	65	Volume Dismount	96	Device Timeout	97	Undefined Interrupt
98	Async Device Attention	99	Software Parameters	100	Logged Message	101	Logged MSCP Message
102	Laser Power Event	112	User Bugcheck				

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

Also called a *signal converter*. An adapter converts the protocol and hardware interface of one bus type into that of another without changing the function of the bus.

allocation class

A numerical value assigned to an integrated storage element to indicate which hosts on a cluster it will be served by.

American National Standards Institute

See **ANSI**.

ANSI

American National Standards Institute. ANSI is 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.

BA350–Mx controller shelf

The StorageWorks controller shelf used for HS controller modules, cache modules, and shelf power units.

BA350–Sx SBB shelf

A StorageWorks shelf used for only power units and SBBs.

bad block

A block containing a defect that:

- Exceeds the correction capability of the subsystem error correction scheme.
- Exceeds a drive-specified error threshold. Once a block exceeds this threshold, data integrity is not guaranteed.

- Imposes too great a strain on system performance. In this case, the subsystem still assures data integrity, but the extensive error correction required for each block access causes too great a strain on system performance.

bad block replacement

See **BBR**.

battery backup unit

See **BBU**.

BBR

Bad block replacement. BBR is the substitution of defect-free disk blocks for those found to have defects. BBR locates a replacement block, marks the bad block as replaced, and moves the data from the bad block to the replacement block.

BBU

Battery backup unit. A BBU extends power availability after the loss of primary ac power or a power supply to protect against the corruption or loss of data.

BIST

Built-in self-test. BIST is the internal self-test routine for the HS controller module microprocessor chip.

block

A stream of data transferred as a unit. Block is used interchangeably with the term "sector" for disk drives to represent 512 bytes (for 16- and 32-bit host architectures) or 576 bytes (for 36-bit architectures). A block is the smallest data unit addressable on a subunit. It occupies a specific physical position relative to the index and is available for reading or writing once per disk rotation. The five types of blocks follow:

1. Diagnostic block—Used for drive read or write diagnostics. The diagnostic block area is not visible to the host operating system. However, it is visible to the controller. Diagnostic block addresses are 28 bits wide and are called diagnostic block numbers (DBNs).
2. External block—Contains the format control tables. The external block area is not visible to the host operating system. However, it is visible to the controller. External block addresses are 28 bits wide and are called external block numbers (XBNs).
3. Logical block—Contains the host applications area and the Replacement Control Table. All logical blocks are visible to the host operating system. Logical block addresses are 28 bits wide and are called logical block numbers (LBNs).
4. Physical block—Contains all the blocks on a subunit. DBNs, LBNs, RBNs, and XBNs are subsets of the physical block area. Physical block addresses are 28 bits wide and are called physical block numbers (PBNs).
5. Replacement block—A reserved block used as a replacement for a bad block on a subunit. Replacement block addresses are 28 bits wide and are called replacement block numbers (RBNs).

blower

An airflow device mounted in a StorageWorks shelf.

built-in self-test

See **BIST**.

cable distribution unit

See **CDU**.

carrier

Sometimes called an SBB carrier. A carrier is a standard, StorageWorks shelf-compatible, plastic shell into which a device can be installed.

CDU

Cable distribution unit. The CDU is the power entry device for StorageWorks cabinets. The unit provides the connections necessary to distribute ac power to cabinet shelves and fans.

CER

CLI error report. The CER spontaneously posts controller error messages to a terminal, if one is connected.

CLI

Command line interpreter. Operator command line interface for the HS controller.

CLI error report

See **CER**.

cluster

A collection of processors, called nodes, attached to each other by a high-speed bus. These processors are independent and survivable. They may be general-purpose computers or special-purpose servers, such as the HS controller, providing a special set of services to the rest of the nodes.

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 disk

A special storage configuration for communication with generic SCSI devices. The command disk is a virtual unit (in controller memory) that acts as a conduit for SCSI communication between a device such as a *loader* and an MSCP-based host.

command line interpreter

See **CLI**.

container

Data storage consisting of either a single disk device, or group of disk devices linked as a storageset. Examples of containers include disk drives, stripesets, and RAIDsets.

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 functions to the host/device communication.

controller shelf

A StorageWorks shelf designed to contain controller and cache memory modules.

CRC

A checkword (polynomial checksum) generally appended to a disk data transfer. CRC is computed using data message bits as coefficients divided by a generating polynomial. The resulting remainder is the CRC. When a transmitter computes and transmits a CRC following a data transfer, the receiver can recompute and compare it with the received version to verify correct reception. EDC and ECC (both used by disks) are examples of CRC checkwords.

cyclic redundancy check

See **CRC**.

DAEMON

Diagnostic and execution monitor. DAEMON is a part of HS1CP controller self-testing that includes port and cache initialization and self-test routines.

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.

device driver

An operating system software module used to physically control an I/O device. In DSA, conventional device drivers are replaced by a single driver for an entire class of devices, such as disk drives, and a single port driver for the host-to-controller transport mechanism.

device shelf

A StorageWorks shelf designed to contain SBBs.

diagnostic and execution monitor

See **DAEMON**.

Diagnostics and Utilities Protocol

See **DUP**.

differential SCSI bus

A SCSI bus in which a signal's level is determined by the potential difference between two lines. A differential bus is more robust and less subject to electrical noise than is a single-ended bus.

digital audio tape

See **DAT**.

Digital Storage Architecture

See **DSA**.

DILX

Disk inline exerciser. DILX is 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. DSA is 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.

Digital Standard Disk Format

See **DSDF**.

DSDF

Digital Standard Disk Format. DSDF is the DSA standard for disk media format. DSDF specifies the mechanism for mapping a contiguous logical block address space into a possibly imperfect physical space, as well as defining diagnostic and factory areas. DSDF is transparent to the system.

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 data link

See **DDL**.

dual porting (or dual access)

The ability of a disk or tape drive to be accessed by two controllers. All DSA drives have a standard dual-port feature. DSA drives can be online to only one controller at a time. However, they are able to disconnect themselves from a failed controller (or be disconnected by a failing controller) and become available for continued service through the other controller.

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

Two controllers in one controller shelf, sharing access to each other's devices. This configuration provides for one controller to take over the work of a failed controller.

DUART

Dual universal asynchronous receiver transmitter. A DUART is an integrated circuit containing two serial, asynchronous transceiver circuits.

DUP

Diagnostic and Utilities Protocol. DUP is host application software that allows a host operator terminal to connect to the HSJ- and HSD-series controllers' command line interpreter. *See also* **virtual terminal**.

ECC

Error correction code. The ECC is 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. The EDC is 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.

EIP

Error information packet (or event information packet). The EIP includes bytes of data translated into information explaining error events.

electromagnetic interference

See **EMI**.

electrostatic discharge

See **ESD**.

EMI

Electromagnetic interference. EMI is the impairment of a signal by an electromagnetic disturbance.

error correction code

See **ECC**.

error detection code

See **EDC**.

error information packet

See **EIP**.

ESD

Electrostatic discharge. ESD is the discharge of a potentially harmful static electric voltage as a result of improper grounding.

event logging

See **EVL**.

EVL

Event logging. EVL is the spontaneous posting of EIP information to a connected terminal. The EVL display allows the user to read event log information without invoking a separate utility for translation.

EXEC

Firmware executive. EXEC is the kernel portion of HS controller firmware.

extended status

An additional set of status information maintained by the drive that is of interest to a host error log. Extended status is drive-type specific and is not utilized by the controller except as input to the host error log and diagnostic processes.

failedset

A disk drive or group of disk drives that has been removed from a RAIDset due to a failure or manual removal. Disk drives in the failedset should be considered defective and should be tested, then repaired or replaced.

failover

A process that takes place when one controller fails in a dual-redundant configuration, and the other controller takes over service to the devices of the failed controller.

fan

An airflow device mounted in a StorageWorks cabinet.

fast, differential SCSI

See **FD SCSI**.

fast, wide, differential SCSI

See **FWD SCSI**.

Fault Management Utility

See **FMU**.

FD SCSI

Fast, differential SCSI. This differential SCSI bus has an 8-bit data transfer rate of 10 MB/s. *See also* **FWD SCSI** and **SCSI**.

field replaceable unit

See **FRU**.

Firmware Licensing System

See **FLS**.

filler panel

A sheet metal or plastic panel used to cover unused mounting areas in StorageWorks cabinets and shelves.

firmware executive

See **EXEC**.

FLS

Firmware Licensing System. The FLS allows a user to enter license keys and activate licensed features (such as RAIDsets) on the HS controller.

flush

To write cache data to storage media.

FMU

Fault Management Utility. The FMU controls certain spontaneous error displays to a maintenance terminal. FMU also provides for interactive review of last failure and memory system failure information.

FRU

Field replaceable unit. An FRU is the lowest-level hardware component that can be replaced by field personnel.

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."

FWD SCSI

Fast, wide, differential SCSI. This differential SCSI bus has a 16-bit data transfer rate of up to 20 MB/s. *See also* **FD SCSI** and **SCSI**.

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.

HBVS

Host-Based Volume Shadowing. HBVS also is known as Phase 2 Volume Shadowing.

HBVS assistance

RAID level 1a. The HS1CP controller performs HBVS assistance by independently directing shadow copy operations that were requested by the host between two units under the given controller.

Hierarchical Storage Controller

See **HSC**.

HIS

Host Interconnect Services. The firmware that communicates with the host in HS1CP controllers.

host

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

Host-Based Volume Shadowing

See **HBVS**.

Host Interconnect Services

See **HIS**.

host logical unit

See **logical unit**.

host terminal

See **virtual terminal**.

hot swap

A method of device replacement whereby the complete system remains online and active during device removal and reinstallation. The device being removed or reinstalled is the only device that cannot perform operations during this process.

HSC

Hierarchical Storage Controller. The HSC is an intelligent mass storage server used on the CI bus. Capable of supporting a total of eight disk and/or tape data channels, the HSC is part of the System Interconnect Architecture and Digital Storage Architecture. By performing as an I/O manager, the HSC can be classified as an I/O server, removing the burden of I/O management from the CPU.

IBR

Initial boot record. The IBR is a table of information placed in memory in accordance with i960 processor specifications. The i960 chip reads the IBR during initialization in order to configure itself, as well as position the program counter at the start of core MIST.

Initial boot record

See **IBR**.

Initiator

The SCSI bus member that requests an operation to be performed by another member (the target). When an HS controller interacts with its physical storage devices, it is the initiator. Another example is when the host CPU interacts with the HSZ-series controller, the host is the initiator.

Instance code

The four-byte value transmitted in the error log packet that is key to interpreting the error.

KILL line

The controller-to-controller disable signal used in a dual-redundant configuration.

last failure

Last failure refers to event/error information specifically generated upon the sudden termination of executing firmware. For example, last failure occurs when there is an abrupt controller failure (such as when power is removed).

last failure logging

See **LFL**.

least recently used

See **LRU**.

LFL

Last failure logging. LFL is the spontaneous posting of last failure information to a connected terminal. The LFL display allows the user to read last failure information without invoking a separate utility for translation.

loader

A robotic mechanism for selecting and loading media into associated drive devices. Such mechanisms are also known as "jukeboxes" or medium changers. A loader is a good application for a *command disk* because SCSI commands are needed to control its robotic movement.

logical unit

Also called a host logical unit or simply a unit. A logical unit is a device or group of devices addressable as one unit by the host. (A logical unit is not necessarily a *LUN*.)

logical unit number

See **LUN**.

LRU

Least recently used. LRU is the block replacement algorithm for the read cache.

LUN

Logical unit number. A LUN is a value from 0–7 that identifies a subset of a SCSI target to a SCSI initiator. LUNs use their target's bus connection to communicate on the SCSI bus.

maintenance terminal

The operator terminal used to identify a controller, to enable its host paths, to define its subsystem configuration, and to check its status. The maintenance terminal interface is designed to accept any terminal conforming to EIA–423. A maintenance terminal is only required to configure a storage subsystem and is not required for normal operations.

mass storage control protocol

See **MSCP**.

metadata

Special data written to a disk device and inaccessible to the host CPU. Metadata improve error detection and media defect management for the disk device.

MIST

Module integrity self-test. MIST tests controller functions upon initialization. *See also* **DAEMON**.

module integrity self-test

See **MIST**.

MSCP

Mass Storage Control Protocol. MSCP is the message-level protocol used by the HS1CP controllers to communicate with a host computer. The three types of MSCP communication are sequential messages, datagrams, and block data transfers.

Network Interconnect

See **NI**.

NI

One of two standard interconnects used in the System Interconnect Architecture (CI is the other). The NI (also known as the Ethernet) connects communications servers and computer servers, creating a local area network.

node

An intelligent entity in a distributed computing configuration. Nodes are independent but linked, as in a network or a cluster, becoming parts of a whole. In a cluster, HSJ-series controllers and host computers are cluster nodes.

nonredundant

A configuration in which there is no backup hardware in place for the hardware that is present.

nontransportable

An operator device assignment that indicates the device is MSCP compliant and contains metadata. Devices assigned as nontransportable can be moved amongst HS controller subsystems, but cannot be moved directly to non-HS controller systems. *See also* **transportable**.

nonvolatile

See **NV**.

nonvolatile memory

See **NVMEM**.

nonvolatile parameter memory

See **NVPM**.

NV

Nonvolatile. NV describes memory, the contents of which survive loss of power.

NVMEM

Nonvolatile memory. NVMEM is the battery backed-up SRAM on the controller module.

NVPM

Nonvolatile parameter memory. NVPM is a portion of NVMEM used to store controller configuration data.

OCP

Operator control panel. The OCP is the control/indicator panel associated with a device. The OCP is usually mounted on the device and is accessible to the operator.

offline

One of the possible status conditions of a mass storage device or server. When a device is offline, it is not capable of communicating with the controller. When the controller is offline, it is inaccessible to any node in the configuration.

operator control panel

See **OCP**.

passthrough

Part of a virtual unit structure associated with a *command disk*. Passthroughs do not consist of physical devices, but are part of a controller configuration allowing MSCP-based hosts to control generic SCSI devices.

PCMCIA

Personal Computer Memory Card Industry Association. PCMCIA is an organization that develops standards for ROM memory cards.

Personal Computer Memory Card Industry Association

See **PCMCIA**.

port

The hardware and software used to connect a controller to a communication bus, such as a Internal bus.

port/target/LUN

See **PTL**.

program card

The PCMCIA card containing the HS controller operating firmware.

PTL

Port/target/LUN. PTL is a three-number hierarchical value representing a device location to a SCSI initiator. For example, PTL 143 is a device on port 1 of the initiator, target 4 on port 1, and LUN 3 under target 4.

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. RAID is a set of storage techniques devised to increase the performance and availability of a storage subsystem.

RAIDset

Three to fourteen physical disks configured as a container. RAIDsets allow one physical disk to fail without loss of data.

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**.

reduced

A RAIDset that has one failed member is said to be running reduced. RAID functions allow the missing data from the failed member to be reconstructed and accessed on an as-needed basis, using the surviving members' data and parity data. A RAIDset will not run reduced (or at all) if more than one member has failed.

reconstruction

The process of regenerating the data from one failed RAIDset member. Reconstruction is possible through XORing the surviving members' data with the parity data (or recalculating parity data, if a parity member failed). When the missing data is reconstructed it is placed on a disk from the spareset, if one is available. If there is no spare disk, the RAIDset will go *reduced* and will not reconstruct.

replacement

The process of changing a *spareset* device into an active member of a RAIDset (replacing a failed RAIDset device). Spareset members are selected as replacements based on either "best fit" or "best performance." Data previously backed up on tape is retrieved for disk storage using the normal priority. Backup is used to preserve information in the event of a disk failure. Restore is used to recover the information.

restore

Data previously backed up on tape is retrieved for disk storage using the normal priority. Backup is used to preserve information in the event of a disk failure. Restore is used to recover the information.

Redundant Array of Independent Disks

See **RAID**.

replacement policy

The firmware-controlled method by which a spare disk (from the spareset) is selected to replace a disk that has failed in a RAIDset.

RFI

Radio frequency interference. RFI is the impairment of a signal by an unwanted radio signal or radio disturbance.

SBB

StorageWorks building block. An SBB is 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. This is a StorageWorks shelf, such as the BA350-Sx, designed to house plug-in SBB modules.

SCA

System Communication Architecture. The interface specifications and protocols defining the connection of independent computer systems into clusters.

SCS

System Communication Services. SCS is a delivery protocol for packets of information (commands or data) to or from the host.

SCSI

Small Computer System Interface. SCSI is 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.

SCSI device

A host computer adapter, a peripheral controller, or a peripheral device that can be attached to the SCSI bus.

SCSI device ID

Also referred to as a *target ID*. This ID is the physical address an initiator uses to connect with a target. Each target is assigned a unique target address.

SCSI–A cable

A 50-conductor, 25 twisted-pair cable used for single-ended, SCSI bus connections.

SCSI–P cable

A 68-conductor, 34 twisted-pair cable used for differential bus connections.

SHAC

Single Host Adapter Chip. The SHAC hardware is common to some DSSI host adapters compatible with HSD-series controlers.

shelf brackets

Sheet metal components designed to attach and position StorageWorks shelves in their associated enclosures.

signal converter

Also called an *adapter*. A signal converter converts the protocol and hardware interface of one bus type into that of another without changing the function of the bus.

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.

Single Host Adapter Chip

See **SHAC**.

skirt

A trim panel designed to mount around the base of a cabinet.

Small Computer System Interface

See **SCSI**.

software product description

See **SPD**.

spareset

A pool of disk drives from which a controller can draw to replace failed members of a RAIDset.

SPD

Software product description. The SPD describes the function of a program.

storageset

A grouping of disk drives that make up a new distinct container. Stripesets and RAIDsets are examples of storagesets.

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**.

stripeset

A virtual disk drive with its physical data spread across from 2–14 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.

surviving controller

The controller in a dual-redundant pair that assumes service to its companion's devices when the companion fails. *See also* **failover**.

System Communication Architecture

See **SCA**.

System Communications Services

See **SCS**.

tagged command queuing

A technique that allows a device to have multiple I/O requests outstanding to it at one time.

Tape Mass Storage Control Protocol

See **TMSCP**.

TMSCP

Tape Mass Storage Control Protocol. TMSCP is an applications protocol used by the HSJ- and HSD-series controllers to communicate with the host computer. TMSCP is tape-specific but overlaps and shares certain portions of MSCP.

Tape Inline Exerciser

See **TILX**.

target

A member of a SCSI bus that carries out operations requested by an initiator. Physical storage devices are targets of all HS1CP controllers, for example.

target ID

Also referred to as *SCSI device ID*. This ID is the physical address an initiator uses to connect with a target. Each target is assigned a unique target address.

template

The group or type to which an error event log to the host belongs. Template type determines how the EIP is formatted.

TILX

Tape Inline Exerciser. TILX is diagnostic firmware used to test the data transfer capabilities of tape drives in a way that simulates a high level of user activity.

transportable

An operator device assignment that indicates the device is not MSCP compliant and does not contain metadata. Transportable devices can be moved between HS controller subsystems and non-HS controller systems. However, such devices do not support forced error, and should not be set to transportable after correct installation in an HS controller subsystem. *See also* **nontransportable**.

unit

See **logical unit**.

unwritten cache data

Data suspended in write-back cache memory which has not been written to storage media yet, even though the host operation using the data has completed.

value-added firmware

The firmware that provides optional, licensed features such as logical block mapping, cache, RAID, and so on.

VAXcluster System Console

See **VCS**.

VCS

VAXcluster System Console. This terminal allows access to hosts (by networks). VCS is 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 (for example, using DUP) or via the maintenance port through an intermediary host (VCS). A virtual terminal is also sometimes called a host terminal.

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-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 have reached the cache.

write hole

Undetectable RAID level 5 data corruption. A write hole is caused by the successful writing of cached data but not the data parity, or vice versa. Write holes occur under conditions such as power outages, where the writing of these two elements can be abruptly interrupted. A battery backed-up cache design eliminates the write hole, because data is preserved and writes can be retried.

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**.

Index

3.5-Inch SBBs
  highest availability, D–29
  highest performance, D–29
  LEDs, 3–3, 3–5, 6–10, 6–12
  optimal performance, D–28
  status, 3–3, 6–10
5.25-Inch SBBs
  highest availability, D–29
  highest performance, D–29
  LEDs, 3–3, 3–4, 3–5, 6–10, 6–11, 6–12
  optimal performance, D–28
  status, 3–4, 6–11

A

Abort codes
  DILX, 7–31
  TILX, 7–51
AC power-up sequence, 2–10, 3–2
Acceptance testing, 4–13, 5–14
Accessing the server processor, 8–12
Acoustics, xxvi
Adapter board
  See device bus adapter
ADD CDROM command, E–2
ADD DISK command, E–3
ADD LOADER command, 5–21, E–5
ADD MIRRORSET command, E–6
ADD OPTICAL command, E–8
ADD PASSTHROUGH command, 5–21, E–10
ADD RAIDSET command, E–11
ADD SPARESET command, E–13
ADD STRIPESET command, E–14
ADD TAPE command, E–15
ADD UNIT command, E–16
Adding a device
  warm swap, 8–40
Adding devices, 5–13, 5–21, 7–71, 7–73
Adding SIMMs, 8–18
Allocation class, 5–10, 5–11
Amber LEDs, 6–5, 7–2, 7–3, A–12
  HS1CP front panel, 3–2, 6–4
arc command, 5–25
ARC interface, 5–24
  switching to SRM from, 5–25
Availability
  configuration, D–29
  highest, D–29

B

BA350–M Shelf replacement, 8–42
BA350–S Shelf replacement, 8–43
Basic function test
  DILX, 7–8
  TILX, 7–34
Battery
  charge time, write-back cache, A–16
  disposal, 8–57
  error messages, 6–21
  NVMEM, A–14
  replacement, 5–18, 8–58, A–16
  SRAM, A–14
  write-back cache, 6–2, 8–57, 8–58, A–15
  write-back cache, discharging, A–16
Beep codes, 2–1, 2–9, 2–10, 2–11
BIST, 7–2
Bit flags
  Connection State Codes
    0000, C–67
    0001, C–67
    0002, C–67
    0003, C–67
    0004, C–67
    0005, C–67
    0006, C–67
    0007, C–67
    0008, C–67
    0009, C–67
    000A, C–67
    000B, C–67
  Virtual Circuit State Codes
    01, C–66
    02, C–66
    03, C–66
    04, C–66
    05, C–66
Blower, 3–5, 6–12, A–21
Blowers
  replacing, 8–44

Boot
See initialization.
Boot diagnostic flow, 1–4
Built-in self-test
See BIST
Bus exchanger, A–15
Bus state problem
during bus quiesce, 8–41
shelf power failure, 8–41

C

Cabinet
FDDI Server, D–4
Cable distribution unit
description, A–3
See CDU., 8–36
Cables
host port, D–20
Cabling
internal bus cables, 8–69
Cache module, 7–4, 8–54
See also Read cache
See also Write-back cache
cache manager, 7–4, 7–5
configuration, D–27
DAEMON, 7–4
error messages, 6–21
failover, 6–2
operation, A–15
service consideration, 6–2
size restriction, D–27
testing of, 7–4
Card slot location, 5–32
cat el command, 2–6
CAT EL command, 4–5
CD–ROM drive
modifying SCSI ID, D–39
CD–ROM LEDs, 2–7
CER, 6–15
CFG files, 2–8, D–41
CFMENU, 7–73
devices, 7–74
example, 7–73
exiting, 7–84
failedset, 7–78
fields, 7–75, 7–76, 7–77, 7–79, 7–81
initializing containers, 7–80
invoking, 7–73
main menu, 7–73
messages, 7–82
passthrough containers, 7–79
RAIDset, 7–76
restrictions, 7–73
spareset, 7–78
storageset, 7–76, 7–78
stripeset, 7–76
CFMENU (cont'd)
terminal setup, 7–82
units, 7–81
Charge time, write-back cache battery, A–16
Chunksize default
for RAIDsets, E–38
for stripesets, E–38
CI node number, 5–9, 5–11
CLCP, A–21
CLEAR_ERRORS CLI command, E–25
CLEAR_ERRORS INVALID_CACHE command, E–26
CLEAR_ERRORS LOST_DATA command, E–27
CLEAR_ERRORS UNKNOWN command, E–28
CLEAR_ERRORS UNWRITEABLE_DATA command, E–29
CLI
command sets, 5–9
commands, E–1
error conventions, E–114
error messages, E–114
error messages, automatic, 6–16
error messages, interactive, 6–24
error report
See CER
firmware, A–20
warning conventions, E–134
warning messages, E–134
CLONE Utility
cloning a mirrorset, 7–111
cloning a single device unit, 7–106
cloning a striped mirrorset, 7–114
cloning a stripeset, 7–108
description of, 7–104
Cluster size, 5–15
Code patch, 7–3
Code Patch considerations, 7–85
Code Patch utility, 7–85
Code Patching, 7–85
Codes
CI Message Operation Codes
00, C–66
01, C–66
02, C–66
03, C–66
04, C–66
05, C–66
06, C–66
07, C–66
08, C–66
09, C–66
10, C–66
11, C–66
12, C–66
13, C–66
0A, C–66
0B, C–66
0C, C–66

Codes
CI Message Operation Codes (cont'd)
0D, C–66
0E, C–66
0F, C–66
Controller Restart Codes
0, C–138
1, C–138
2, C–138
Event Codes
0007, C–107, C–111
0014, C–102, C–111
0016, C–86, C–107, C–108, C–109, C–110, C–111, C–112, C–113
0037, C–112, C–113
0077, C–112
0097, C–112, C–113
0103, C–110, C–111
002A, C–102, C–103
006A, C–99, C–100, C–101, C–102, C–104
008A, C–103
012A, C–86, C–87, C–88, C–89, C–90, C–91, C–92, C–93, C–94, C–95, C–96, C–97, C–98, C–99
016A, C–102, C–104, C–107
020A, C–87
022A, C–85, C–86
040A, C–85
01AA, C–102, C–103, C–104
000B, C–109, C–111
002B, C–110, C–111, C–112
012B, C–110, C–111, C–112
014B, C–109, C–111
01CA, C–103
00CB, C–110, C–111
00E8, C–110, C–111
01EA, C–103, C–104
03EA, C–85
00EB, C–104, C–105, C–106, C–107, C–110, C–111, C–112
Event Notification/Recovery Threshold
Classification Value
01, C–139
02, C–139
64, C–139
0A, C–139
Firmware Component Identifier Codes
01, C–63
02, C–63
03, C–63
04, C–63
05, C–64
06, C–64
07, C–64
08, C–64
20, C–64
40, C–64
42, C–64

Codes
Firmware Component Identifier Codes (cont'd)
60, C–64
61, C–64
62, C–64
80, C–64
81, C–64
82, C–64
83, C–64
84, C–64
flashing HS1CP, 6–5
Format Codes
00, C–30, C–32, C–34, C–40, C–45, C–48, C–51
01, C–36, C–38, C–43
02, C–53
05, C–58
09, C–57
0A, C–62
Host Interconnect Services Status Codes
00000000, C–64
00000001, C–64
00000002, C–64
00000003, C–64
00000004, C–64
00000009, C–64
00000032, C–64
00000033, C–64
00000034, C–64
00000035, C–64
00000036, C–64
00000064, C–64
00000065, C–65
00010009, C–65
00020009, C–65
00030009, C–65
00040009, C–65
00050009, C–65
00060009, C–65
00070009, C–65
00080009, C–65
00090009, C–65
00100009, C–65
00110009, C–65
00120009, C–65
00130009, C–65
00140009, C–65
00150009, C–65
00160009, C–65
00170009, C–65
00180009, C–65
00190009, C–65
00200009, C–65
000A0009, C–65
001A0009, C–65
000B0009, C–65
001B0009, C–65
000C0009, C–65

Codes
Host Interconnect Services Status Codes (cont'd)
001C0009, C–65
000D0009, C–65
001D0009, C–65
000E0009, C–65
001E0009, C–65
000F0009, C–65
001F0009, C–65
HS1CP Controller Vendor Specific SCSI ASC/ASCQ Codes
82 01, C–84
84 04, C–84
85 05, C–84
89 00, C–84
93 00, C–84
8A 00, C–84
A0 00, C–84
A0 01, C–84
A0 02, C–84
A0 03, C–84
A0 04, C–84
A0 05, C–84
A0 06, C–84
A0 07, C–84
A1 00, C–84
A1 01, C–84
A1 02, C–84
A1 03, C–84
A1 04, C–84
B0 00, C–84
B0 01, C–84
8C 04, C–84
D0 01, C–84
D0 02, C–84
D0 03, C–84
D1 00, C–84
D1 02, C–84
D1 03, C–84
D1 04, C–84
D1 05, C–84
D1 07, C–84
D1 08, C–84
D1 09, C–85
D1 0A, C–85
D1 0B, C–85
D2 00, C–85
D2 03, C–85
D3 00, C–85
D4 00, C–85
D5 02, C–85
D7 00, C–85
E0 03, C–85
E0 06, C–85
E0 07, C–85
8F 00, C–84
3F 85, C–83
3F 87, C–83
Codes
HS1CP Controller Vendor Specific SCSI ASC/ASCQ Codes (cont'd)
3F 88, C–83
3F 90, C–83
3F C0, C–84
3F C2, C–84
3F D2, C–84
Instance Codes
01010302, C–85
01032002, C–86
01292602, C–88
01303702, C–88
01313702, C–88
01323702, C–88
01333702, C–88
01342802, C–88
01352902, C–88
01362902, C–88
01392802, C–89
01402902, C–89
01412902, C–89
01422902, C–89
01432902, C–89
01442902, C–89
01503702, C–90
01513702, C–90
01533702, C–90
01543702, C–90
01553702, C–90
01563702, C–90
01933702, C–94
01943702, C–94
01953702, C–94
01963702, C–94
01973702, C–94
01993002, C–94
02020064, C–111
02032001, C–87
02042001, C–87
02052301, C–87
02062301, C–87
02072201, C–98
02082201, C–98
02090064, C–107
02102401, C–99
02110064, C–111
02192401, C–99
02200064, C–107
02210064, C–107
02220064, C–108
02230064, C–108
02240064, C–108
02280064, C–108
02290064, C–108
02300064, C–109
02310064, C–109
03010101, C–102

Codes
Instance Codes (cont'd)
03022002, C–102
03034002, C–102
03044402, C–102
03052002, C–103
03062002, C–103
03070101, C–103
03080101, C–103
03094002, C–109
03104002, C–110
03134002, C–110
03144002, C–110
03154002, C–110
03164002, C–110
03170064, C–110
03180064, C–110
03194002, C–110
03204002, C–110
03214002, C–110
03224002, C–110
03234002, C–110
03244002, C–110
03254002, C–110
03270101, C–104
03644002, C–111
03674002, C–111
03694002, C–111
03704002, C–111
03714002, C–111
03720064, C–111
03730064, C–112
03744002, C–112
03754002, C–112
03760101, C–112
03774002, C–112
03784002, C–112
03794002, C–112
03804002, C–112
03820101, C–102
03832002, C–102
03844002, C–102
03854402, C–103
03862002, C–103
03872002, C–103
03880101, C–103
03890101, C–104
03964002, C–112
03994002, C–112
07050064, C–85
40440064, C–99
82012002, C–87
82042002, C–87
82052002, C–87
82062002, C–87
82072002, C–87
82082002, C–87
0102030A, C–85

Codes
Instance Codes (cont'd)
0121370A, C–88
0122330A, C–88
0217000A, C–86
0218000A, C–86
0227000A, C–86
0311430A, C–110
0312430A, C–110
0326450A, C–110
0328450A, C–109
0368000A, C–111
0381450A, C–112
4003640A, C–99
4004020A, C–99
4007640A, C–99
4015020A, C–100
4029010A, C–100
4051020A, C–100
4052020A, C–100
4053020A, C–100
4054020A, C–100
4055020A, C–100
4056020A, C–100
4057020A, C–100
4058020A, C–100
4059020A, C–100
4060020A, C–101
4061020A, C–101
4062020A, C–101
4063020A, C–101
4064020A, C–101
4065020A, C–101
4066020A, C–101
4067020A, C–101
039A000A, C–112
020A0064, C–111
021A0064, C–102
022A0064, C–108
038A0101, C–104
402A010A, C–100
01372A02, C–88
01382A02, C–88
01452A02, C–89
01462A02, C–89
01472A02, C–89
01482A02, C–89
01492A02, C–89
405A020A, C–100
01A02F02, C–94
03A04002, C–112
01A13002, C–94
03A14002, C–112
01A23002, C–94
03A24002, C–113
013A2802, C–89
014A2A02, C–89
015A2C02, C–90

Codes
Instance Codes (cont'd)
016A2D02, C–91
017A2D02, C–92
018A2E02, C–93
019A3002, C–94
01A33002, C–94
03A34002, C–113
012A3702, C–88
031A4002, C–110
036A4002, C–111
037A4002, C–112
03A40064, C–113
01A43002, C–94
03A50064, C–113
01A53002, C–94
01A63002, C–94
03A64002, C–113
01A73002, C–94
03A74002, C–113
03A80101, C–113
01A83002, C–95
01A93102, C–95
03A94002, C–113
01AA3102, C–95
03AA4002, C–113
01AB3102, C–95
03AB4002, C–113
01AC3102, C–95
03AC4002, C–113
01AD3102, C–95
03AD4002, C–113
01AE3102, C–95
03AE4002, C–113
01AF3102, C–95
03AF4002, C–113
021B0064, C–102
022B0064, C–108
023B0064, C–109
031B0101, C–110
402B010A, C–100
01522B02, C–90
405B020A, C–100
01B03102, C–95
03B04002, C–113
07030B0A, C–85
07040B0A, C–85
07080B0A, C–86
01B13702, C–95
03B14002, C–113
020B2201, C–98
01B23702, C–95
03B24002, C–113
013B2802, C–89
014B2A02, C–89
015B2C02, C–90
016B2D02, C–91
017B2E02, C–92

Codes
Instance Codes (cont'd)
018B2E02, C–93
019B3102, C–94
01B33702, C–95
03B3450A, C–113
012B3702, C–88
036B4002, C–111
037B4002, C–112
039B4002, C–112
03B40101, C–102
01B43702, C–95
038B450A, C–111
03B52002, C–102
01B53702, C–95
01B63202, C–95
03B64002, C–102
01B73202, C–95
03B74402, C–103
03B82002, C–103
01B83202, C–95
03B92002, C–103
01B93202, C–95
03BA0101, C–103
01BA3202, C–95
03BB0101, C–104
01BB3202, C–95
03BC0101, C–104
01BC3202, C–96
01BD3702, C–96
03BD450A, C–112
01BE3702, C–96
01BF3702, C–96
022C0064, C–108
023C0064, C–109
07060C01, C–86
07070C01, C–86
402C010A, C–100
01572C02, C–90
01582C02, C–90
01592C02, C–90
01602C02, C–90
01612C02, C–91
01622C02, C–91
405C020A, C–100
01C03702, C–96
01C13702, C–96
020C2201, C–98
01C22F02, C–96
013C2802, C–89
014C2A02, C–89
015C2C02, C–90
016C2D02, C–91
017C2E02, C–92
018C2E02, C–93
019C3102, C–94
01C32F02, C–96
012C3702, C–88

Codes
Instance Codes (cont'd)
030C4002, C–109
031C4002, C–110
037C4002, C–112
039C4002, C–112
01C42F02, C–96
036C430A, C–111
01C52F02, C–96
01C62F02, C–96
01C72F02, C–96
03C80101, C–102
01C82F02, C–96
03C92002, C–102
01C92F02, C–96
01CA3002, C–96
03CA4002, C–102
03CB0101, C–104
01CB3002, C–96
03CC0101, C–104
01CC3002, C–96
03CD2002, C–103
01CD3002, C–96
03CE2002, C–103
01CE3002, C–96
03CF0101, C–103
01CF3002, C–96
030D000A, C–110
021D0064, C–99
022D0064, C–109
023D0064, C–109
01632D02, C–91
01642D02, C–91
01652D02, C–91
01662D02, C–91
01672D02, C–91
01682D02, C–91
01692D02, C–91
01702D02, C–91
01712D02, C–91
01722D02, C–92
01732D02, C–92
01742D02, C–92
01752D02, C–92
01762D02, C–92
01772D02, C–92
01782D02, C–92
01792D02, C–92
403D020A, C–99
405D020A, C–100
01D03002, C–97
03D04002, C–104
01D13002, C–97
03D14002, C–104
01D23002, C–97
020D2401, C–98
03D24402, C–103
013D2902, C–89

Codes
Instance Codes (cont'd)
015D2C02, C–90
016D2D02, C–91
017D2E02, C–92
018D2E02, C–93
019D2F02, C–94
01D33002, C–97
03D3450A, C–105
012D3702, C–88
014D3702, C–89
031D4002, C–110
037D4002, C–112
039D4002, C–112
01D43002, C–97
036D430A, C–111
03D4450A, C–105
01D53002, C–97
03D5450A, C–105
01D63002, C–97
03D6450A, C–105
01D73002, C–97
03D7450A, C–105
01D83002, C–97
03D8450A, C–106
01D93002, C–97
03D9450A, C–106
01DA3102, C–97
03DA450A, C–106
01DB3102, C–97
03DB450A, C–106
01DC3102, C–97
03DC450A, C–106
01DD3102, C–97
03DD450A, C–106
01DE3102, C–97
03DE450A, C–106
01DF3102, C–97
03DF450A, C–106
021E0064, C–107
022E0064, C–109
01802E02, C–92
01812E02, C–92
01822E02, C–93
01832E02, C–93
01842E02, C–93
01852E02, C–93
01862E02, C–93
01872E02, C–93
01882E02, C–93
01892E02, C–93
01902E02, C–93
01912E02, C–93
01922E02, C–93
405E020A, C–101
01E03102, C–97
03E0450A, C–106
01E13102, C–97

Codes
Instance Codes (cont'd)
03E1450A, C–106
01E23102, C–97
020E2401, C–99
03E2450A, C–107
012E2702, C–88
013E2902, C–89
015E2C02, C–90
016E2D02, C–91
017E2E02, C–92
018E2E02, C–93
019E2F02, C–94
01E33102, C–97
014E3702, C–90
030E4002, C–110
031E4002, C–110
036E4002, C–111
037E4002, C–112
039E430A, C–112
01E43102, C–98
01E53102, C–98
01E63102, C–98
01E73102, C–98
01E83102, C–98
01E93102, C–98
01EA3702, C–98
01EB3702, C–98
01EC3702, C–98
01ED3702, C–98
01EE3702, C–98
03F00402, C–104
021F0064, C–107
022F0064, C–109
01982F02, C–94
405F020A, C–101
03F10502, C–105
03F20064, C–105
020F2401, C–99
013F2902, C–89
015F2C02, C–90
016F2D02, C–91
017F2E02, C–92
018F2E02, C–93
019F2F02, C–94
03F30064, C–105
012F3702, C–88
014F3702, C–90
030F4002, C–110
031F4002, C–110
036F4002, C–111
037F4002, C–112
03F40064, C–104
039F430A, C–112
03F50402, C–107
400F640A, C–99
Last Failure Codes
firmware
Codes
Last Failure Codes
firmware (cont'd)
01000100, C–113
01010100, C–113
01020100, C–113
01030100, C–113
01040100, C–114
01050104, C–114
01070100, C–114
01082004, C–114
01090105, C–114
02000100, C–115
02010100, C–115
02040100, C–115
02050100, C–115
02080100, C–115
02090100, C–115
02100100, C–116
02170100, C–116
02180100, C–116
02210100, C–116
02220100, C–116
02270104, C–116
02360101, C–116
02370102, C–116
02440100, C–117
02530102, C–117
02560102, C–117
02570102, C–117
02620102, C–118
02690102, C–118
02730100, C–118
02790102, C–118
02800100, C–118
02820100, C–118
02830100, C–118
02840100, C–118
02850100, C–118
02860100, C–118
02880100, C–118
02890100, C–118
02900100, C–118
02910100, C–119
02920100, C–119
02930100, C–119
02940100, C–119
02950100, C–119
02960100, C–119
02970100, C–119
02990100, C–119
03020101, C–120
03030101, C–120
03040101, C–120
03050101, C–120
03060101, C–120
03070101, C–120
03080101, C–120

Codes
Last Failure Codes
firmware (cont'd)
03150100, C–121
03280100, C–121
03290100, C–121
03320101, C–121
03370108, C–123
03390108, C–124
03410101, C–125
03470100, C–126
03480100, C–126
03490100, C–126
03500100, C–126
03510100, C–126
03520100, C–126
04010101, C–126
04020102, C–127
04030102, C–127
04040103, C–127
04050100, C–127
04060100, C–127
04070103, C–127
04080102, C–127
04090100, C–127
04100101, C–128
04110101, C–128
04120101, C–128
05010100, C–128
05030100, C–128
06010100, C–128
06020100, C–128
06030100, C–128
06040100, C–128
06050100, C–128
06060100, C–128
07010100, C–128
07020100, C–128
07030100, C–128
07040100, C–129
07050100, C–129
07060100, C–129
07070100, C–129
08010101, C–129
08020100, C–129
08030101, C–129
08040101, C–129
08050100, C–129
08060100, C–129
08070100, C–129
08080000, C–129
08090010, C–129
08100101, C–129
08110101, C–129
08120100, C–130
08130100, C–130
08140100, C–130
08150100, C–130
08160100, C–130
08170100, C–130
08180100, C–130
08190100, C–130
20010100, C–130
20020100, C–130
20030100, C–130
20070100, C–130
20080000, C–130
20090010, C–130
20100101, C–131
20110100, C–131
20120101, C–131
20130101, C–131
40000101, C–131
40150100, C–131
40280100, C–131
40290100, C–131
40300100, C–131
40510100, C–131
40520100, C–131
40530100, C–131
40560100, C–131
40900100, C–132
40930100, C–132
40950100, C–132
40960100, C–132
40970100, C–132
40980100, C–132
42000100, C–132
42020100, C–132
42030100, C–132
42060100, C–132
42120100, C–132
42340100, C–132
42590001, C–133
60050100, C–133
60070100, C–133
60090100, C–133
60120100, C–133
60130100, C–133
60140100, C–133
60150100, C–133
60160100, C–133
60170100, C–133
60180100, C–134
60190100, C–134
60250100, C–134
60260100, C–134
60270100, C–134
60280100, C–134
60290100, C–134
60400100, C–134
60410100, C–134
60420100, C–134

Codes
Last Failure Codes
firmware (cont'd)
60440100, C–134
60450100, C–134
60460100, C–135
60480100, C–135
60490100, C–135
60500100, C–135
60550100, C–135
60560100, C–135
60570100, C–135
60580100, C–135
60610100, C–135
60620100, C–135
60640100, C–135
60650100, C–135
60660100, C–135
60680100, C–135
61010000, C–136
61020100, C–136
61090100, C–136
62000100, C–136
62020100, C–136
62030100, C–136
80010100, C–136
80020100, C–136
80030100, C–136
80040100, C–136
80050100, C–136
80060100, C–136
80070100, C–136
80080100, C–136
80090100, C–137
80100100, C–137
80120100, C–137
80130100, C–137
80140100, C–137
81010100, C–137
81020100, C–137
81030100, C–137
81040100, C–137
81050100, C–137
81060100, C–137
81070100, C–137
81080100, C–137
81090100, C–137
81100100, C–137
81110100, C–138
81120100, C–138
81130100, C–138
81140100, C–138
83020100, C–138
83030100, C–138
83050100, C–138
83060100, C–138
84010100, C–138
080A0000, C–129

Codes
Last Failure Codes
firmware (cont'd)
200A0000, C–130
425A0001, C–133
02A00100, C–119
020A0100, C–115
028A0100, C–118
029A0100, C–119
030A0100, C–120
032A0100, C–121
034A0100, C–126
081A0100, C–130
402A0100, C–131
602A0100, C–134
604A0100, C–135
800A0100, C–137
810A0100, C–137
811A0100, C–138
025A0102, C–117
02A10100, C–119
02A20100, C–119
02A30100, C–119
02A40100, C–119
02A50100, C–119
02A60100, C–119
02A70100, C–119
02A80000, C–119
02A90100, C–119
02AA0100, C–119
02AB0100, C–119
02AC0100, C–119
02AE0100, C–120
02AF0102, C–120
424B0001, C–132
02B00102, C–120
020B0100, C–115
028B0100, C–118
032B0100, C–121
034B0100, C–126
080B0100, C–129
200B0100, C–130
402B0100, C–131
407B0100, C–131
420B0100, C–132
601B0100, C–134
602B0100, C–134
603B0100, C–134
604B0100, C–135
606B0100, C–135
610B0100, C–136
800B0100, C–137
810B0100, C–137
811B0100, C–138
081B0101, C–130
025B0102, C–117
027B0102, C–118
02B10100, C–120

Codes
Last Failure Codes
firmware (cont'd)
40B40101, C–132
02BC0100, C–120
02BD0100, C–120
424C0001, C–133
020C0100, C–115
028C0100, C–118
034C0100, C–126
080C0100, C–129
200C0100, C–130
402C0100, C–131
407C0100, C–131
600C0100, C–133
601C0100, C–134
602C0100, C–134
603C0100, C–134
606C0100, C–135
610C0100, C–136
800C0100, C–137
810C0100, C–137
811C0100, C–138
033C0101, C–124
081C0101, C–130
025C0102, C–117
021D0100, C–116
027D0100, C–118
029D0100, C–119
034D0100, C–126
040D0100, C–127
080D0100, C–129
409D0100, C–132
600D0100, C–133
601D0100, C–134
602D0100, C–134
604D0100, C–135
605D0100, C–135
606D0100, C–136
800D0100, C–137
810D0100, C–137
200D0101, C–130
402D0101, C–131
010D0110, C–114
021E0100, C–116
027E0100, C–118
028E0100, C–118
031E0100, C–121
408E0100, C–131
601E0100, C–134
602E0100, C–134
605E0100, C–135
606E0100, C–136
800E0100, C–137
810E0100, C–137
080E0101, C–129
200E0101, C–131
402E0101, C–131

Codes
Last Failure Codes
firmware (cont'd)
022E0102, C–116
033E0108, C–125
010E0110, C–115
021F0100, C–116
027F0100, C–118
028F0100, C–118
029F0100, C–119
031F0100, C–121
402F0100, C–131
408F0100, C–132
424F0100, C–133
604F0100, C–135
605F0100, C–135
810F0100, C–137
080F0101, C–129
200F0101, C–131
040F0102, C–128
033F0108, C–125
010F0110, C–115
hardware
02392084, C–116
03330188, C–122
03350188, C–122
03360188, C–123
03380188, C–124
03420188, C–126
42332080, C–132
42382080, C–132
42392080, C–132
42442080, C–132
42452080, C–132
42472080, C–132
42482080, C–132
42602080, C–133
42612080, C–133
42622080, C–133
42632080, C–133
018000A0, C–115
018600A0, C–115
018700A0, C–115
018800A0, C–115
018900A0, C–115
010A2080, C–114
423A2080, C–132
023A2084, C–117
02AD0180, C–120
030B0188, C–121
423B2080, C–132
425B2080, C–133
010B2380, C–114
018B2580, C–115
423C2080, C–132
425C2080, C–133
010C2380, C–114
018C2580, C–115

Codes
Last Failure Codes
hardware (cont'd)
423D2080, C–132
424D2080, C–133
425D2080, C–133
034E2080, C–126
423E2080, C–132
424E2080, C–133
425E2080, C–133
423F2080, C–132
425F2080, C–133
Port/Port Driver Message Operation Codes
0000, C–66
0001, C–66
0002, C–66
0003, C–66
0004, C–66
0005, C–67
0006, C–67
Recommended Repair Action Codes
00, C–140
01, C–140
02, C–140
03, C–140
04, C–140
05, C–140
06, C–140
07, C–140
08, C–140
09, C–140
20, C–141
22, C–141
23, C–141
24, C–141
25, C–141
26, C–141
27, C–141
28, C–142
29, C–142
30, C–143
31, C–144
32, C–144
33, C–144
34, C–144
35, C–144
36, C–144
37, C–144
40, C–144
41, C–144
43, C–144
44, C–144
45, C–144
60, C–145
61, C–145
63, C–145
64, C–145
0A, C–140
Codes
Recommended Repair Action Codes (cont'd)
2A, C–142
0B, C–140
2B, C–142
0C, C–141
2C, C–142
2D, C–143
2E, C–143
2F, C–143
SCSI ASC/ASCQ Codes
00 00, C–71, C–75, C–78, C–81
00 01, C–75
00 02, C–75
00 03, C–75
00 04, C–75
00 05, C–75
00 06, C–72, C–75, C–78, C–81
00 11, C–79
00 12, C–79
00 13, C–79
00 14, C–79
00 15, C–79
01 00, C–72
02 00, C–72, C–79, C–81
03 00, C–72, C–75
03 01, C–75
03 02, C–75
04 00, C–72, C–75, C–79, C–81
04 01, C–72, C–75, C–79, C–81
04 02, C–72, C–76, C–79, C–81
04 03, C–72, C–76, C–79, C–81
04 04, C–72, C–76
06 00, C–72, C–79, C–82
07 00, C–72, C–76, C–79, C–82
08 00, C–72, C–76, C–79, C–82
08 01, C–72, C–76, C–79, C–82
08 02, C–72, C–76, C–79, C–82
09 00, C–72, C–76, C–79
09 01, C–79
09 02, C–79
09 03, C–79
10 00, C–72
11 00, C–72, C–76, C–79
11 01, C–72, C–76
11 02, C–72, C–76
11 03, C–72, C–76
11 04, C–72
11 05, C–79
11 06, C–79
11 08, C–76
11 09, C–76
12 00, C–72
13 00, C–72
14 00, C–72, C–76, C–79
14 01, C–72, C–76, C–79
14 02, C–76
14 03, C–76

Codes
SCSI ASC/ASCQ Codes (cont'd)
14 04, C–76
15 00, C–72, C–76, C–79, C–82
15 01, C–72, C–76, C–79, C–82
15 02, C–72, C–76, C–79
16 00, C–72
17 00, C–72, C–76, C–79
17 01, C–73, C–76, C–79
17 02, C–73, C–76, C–79
17 03, C–73, C–76, C–79
17 04, C–79
17 05, C–73, C–79
17 06, C–73
17 07, C–73
17 08, C–73
18 00, C–73, C–76, C–79
18 01, C–73, C–79
18 02, C–73, C–79
18 03, C–80
18 04, C–80
18 05, C–73, C–80
18 06, C–73, C–80
19 00, C–73
19 01, C–73
19 02, C–73
19 03, C–73
20 00, C–73, C–76, C–80, C–82
21 00, C–73, C–76, C–80, C–82
21 01, C–82
22 00, C–73
24 00, C–73, C–76, C–80, C–82
25 00, C–73, C–76, C–80, C–82
26 00, C–73, C–77, C–80, C–82
26 01, C–73, C–77, C–80, C–82
26 02, C–73, C–77, C–80, C–82
26 03, C–73, C–77, C–80, C–82
27 00, C–73, C–77
28 00, C–73, C–77, C–80, C–82
28 01, C–82
29 00, C–73, C–77, C–80, C–82
29 01, C–73, C–77, C–80, C–82
29 02, C–74, C–77, C–80, C–82
29 03, C–74, C–77, C–80, C–82
30 00, C–74, C–77, C–80, C–82
30 01, C–74, C–77, C–80
30 02, C–74, C–77, C–80
30 03, C–74, C–77
31 00, C–74, C–77
31 01, C–74
32 00, C–74
32 01, C–74
33 00, C–77
37 00, C–74, C–77, C–80, C–82
39 00, C–74, C–77, C–80, C–82
40 00, C–74
41 00, C–74
42 00, C–74

Codes
SCSI ASC/ASCQ Codes (cont'd)
43 00, C–74, C–78, C–81, C–83
44 00, C–74, C–78, C–81, C–83
45 00, C–74, C–78, C–81, C–83
46 00, C–74, C–78, C–81, C–83
47 00, C–74, C–78, C–81, C–83
48 00, C–75, C–78, C–81, C–83
49 00, C–75, C–78, C–81, C–83
50 00, C–78
50 01, C–78
50 02, C–78
51 00, C–78
52 00, C–78
53 00, C–75, C–78, C–81, C–83
53 01, C–78
53 02, C–75, C–78, C–81, C–83
57 00, C–81
63 00, C–81
64 00, C–81
11 0A, C–72, C–76
0A 00, C–72, C–76, C–79, C–82
1A 00, C–73, C–76, C–80, C–82
2A 00, C–74, C–77, C–80, C–82
3A 00, C–74, C–77, C–80, C–82
4A 00, C–75, C–78, C–81, C–83
5A 00, C–75, C–78, C–81, C–83
2A 01, C–74, C–77, C–80, C–82
5A 01, C–75, C–78, C–81, C–83
2A 02, C–74, C–77, C–80, C–82
5A 02, C–75, C–78
5A 03, C–75, C–78
11 0B, C–72
1B 00, C–73, C–76, C–80, C–82
2B 00, C–74, C–77, C–80
3B 00, C–77
4B 00, C–75, C–78, C–81, C–83
5B 00, C–75, C–78, C–81, C–83
3B 01, C–77
5B 01, C–75, C–78, C–81, C–83
3B 02, C–77
5B 02, C–75, C–78, C–81, C–83
5B 03, C–75, C–78, C–81, C–83
3B 08, C–77
3B 0D, C–82
3B 0E, C–82
11 0C, C–72
0C 00, C–76
1C 00, C–73
2C 00, C–74, C–77, C–80, C–82
4C 00, C–75, C–78, C–81, C–83
5C 00, C–75
0C 01, C–72
1C 01, C–73
5C 01, C–75
0C 02, C–72
1C 02, C–73
5C 02, C–75

Codes
SCSI ASC/ASCQ Codes (cont'd)
1D 00, C–73
2D 00, C–77
3D 00, C–74, C–77, C–80, C–83
1E 00, C–73
3E 00, C–74, C–77, C–80, C–83
4E 00, C–75, C–78, C–81, C–83
2F 00, C–74, C–77, C–80, C–82
3F 00, C–74, C–77, C–80, C–83
3F 01, C–74, C–77, C–80, C–83
3F 02, C–74, C–77, C–80, C–83
3F 03, C–74, C–77, C–80, C–83
40 nn, C–74, C–78, C–81, C–83
SCSI Buffered Modes Codes
0, C–70
1, C–70
2, C–70
3, C–70
4, C–70
5, C–70
6, C–70
7, C–70
SCSI Command Operation Codes
00, C–68
01, C–68
03, C–68
04, C–68
05, C–68
07, C–68
08, C–68
10, C–68
11, C–68
12, C–68
13, C–68
14, C–68
15, C–68
16, C–68
17, C–68
18, C–68
19, C–68
25, C–68
28, C–68
30, C–69
31, C–69
32, C–69
33, C–69
34, C–69
35, C–69
36, C–69
37, C–69
39, C–69
40, C–69
41, C–69
42, C–69
43, C–69
44, C–69
45, C–69

Codes
SCSI Command Operation Codes (cont'd)
47, C–69
48 C–69
49, C–69
55, C–69
0A, C–68
1A, C–68
2A, C–69
3A, C–69
5A, C–69
A5, C–69
A6, C–70
A8, C–70
A9, C–70
AF, C–70
0B, C–68
1B, C–68
2B, C–69
3B, C–69
4B, C–69
B0, C–70
B1, C–70
B2, C–70
B3, C–70
B5, C–70
B6, C–70
B8, C–70
1C, C–68
3C, C–69
4C, C–69
1D, C–68
4D, C–69
1E, C–68
2E, C–69
3E, C–69
0F, C–68
2F, C–69
3F, C–69
SCSI Device Type Codes
00, C–67
01, C–68
05, C–68
08, C–68
SCSI Sense Key Codes
0, C–70
1, C–70
2, C–71
3, C–71
4, C–71
5, C–71
6, C–71
7, C–71
8, C–71
9, C–71
A, C–71
B, C–71
C, C–71

Codes
SCSI Sense Key Codes (cont'd)
D, C–71
E, C–71
F, C–71
Solid HS1CP front panel, 6–5
System Communication Services Message Operation Codes
0000, C–67
0001, C–67
0002, C–67
0003, C–67
0004, C–67
0005, C–67
0006, C–67
0007, C–67
0008, C–67
0009, C–67
000A, C–67
000B, C–67
Template Codes
01, C–28
05, C–32
06, C–34
11, C–36
12, C–38, C–63
13, C–40
14, C–42
32, C–45
33, C–47
41, C–50
51, C–53
57, C–55
61, C–58
71, C–60
COM2 and parallel port loopback tests, 4–3
Command disk, 5–21
Command line interpreter
See CLI
Commands
ADD CDROM, E–2
ADD DISK, E–3
ADD LOADER, 5–21, E–5
ADD MIRRORSET, E–6
ADD OPTICAL, E–8
ADD PASSTHROUGH, 5–21
ADD PASSTRHOUGH, E–10
ADD RAIDSET, E–11
ADD SPARESET, E–13
ADD STRIPESET, E–14
ADD TAPE, E–15
ADD UNIT, E–16
CLEAR_ERRORS CLI, E–25
CLEAR_ERRORS INVALID_CACHE, E–26
CLEAR_ERRORS LOST_DATA, E–27
CLEAR_ERRORS UNKNOWN, E–28
Commands (cont'd)
CLEAR_ERRORS UNWRITEABLE_DATA, E–29
DELETE *container-name*, E–30
DELETE FAILEDSET, E–31
DELETE SPARESET, E–32
DELETE *unit-number*, E–33
diagnostic, summarized, 4–1
DIRECTORY, E–34
disk-device-name, E–113
EXIT, E–35
firmware console, functions of, 1–7
HELP, E–36
INITIALIZE, E–37
LOCATE, E–39
LOCATE CANCEL, E–39
LOCATE DISKS, E–39
LOCATE entity, E–40
LOCATE LOADERS, E–39
LOCATE OPTICALS, E–39
LOCATE PTL SCSI-location, E–40
LOCATE TAPES, E–39
LOCATE UNITS, E–40
MEMBERSHIP=*number-of-members*, E–59
MIRROR *disk-device-name1 container-name*, E–41
REDUCE *disk-device-name1 [disk-device-nameN]*, E–43
REDUCED, E–12
REMOVE=*disk-container-name*, E–67
REMOVE=*disk-device-name*, E–60
RENAME, E–45
REPLACE=*disk-container-name*, E–67
REPLACE=*disk-device-name*, E–59
RESTART OTHER_CONTROLLER, E–46
RESTART THIS_CONTROLLER, E–48
RETRY_ERRORS UNWRITEABLE_DATA, E–50
RUN, E–51
SELFTEST OTHER_CONTROLLER, E–52
SELFTEST THIS_CONTROLLER, E–54
SET *disk-container-name*, E–56
SET FAILOVER, E–57
SET *mirrorset-container-name*, E–59
SET NOFAILOVER, E–62
SET OTHER_CONTROLLER, E–63
SET *raidset-container-name*, E–66
SET THIS_CONTROLLER, E–68
SET *unit-number*, E–70
SHOW *cdrom-container-name*, E–80
SHOW CDROMS, E–79
SHOW DEVICES, E–81
SHOW *disk-container-name*, E–83
SHOW DISKS, E–82
SHOW FAILEDSET, E–84
SHOW *loader-container-name*, E–86
SHOW LOADERS, E–85
SHOW *mirrorset-container-name*, E–89

Commands (cont'd)
SHOW MIRRORSETS, E–87
SHOW *optical-container-name*, E–91
SHOW OPTICALS, E–90
SHOW OTHER_CONTROLLER, E–92
SHOW PASSTHROUGH, E–93
SHOW *passthrough-container-name*, E–94
SHOW *raidset-container-name*, E–97
SHOW RAIDSETS, E–95
SHOW SPARESET, E–98
SHOW STORAGESETS, E–99
SHOW *stripeset-container-name*, E–103
SHOW STRIPESETS, E–101
SHOW *tape-container-name*, E–105
SHOW TAPES, E–104
SHOW THIS_CONTROLLER, E–106
SHOW *unit-number*, E–108
SHOW UNITS, E–107
SHUTDOWN OTHER_CONTROLLER, E–109
SHUTDOWN THIS_CONTROLLER, E–111
to examine system configuration, 5–25
Component Access
cabinet preparation, 8–2
CONFIG, 1–5
CONFIG command, 7–71
CONFIG utility, 7–71
Configuration
See also ECU
5.25–Inch SBB restrictions, D–12
3.5-inch SBBs, D–11
atypical, D–14
cache module, D–27
command disk, 5–21
console port, 5–37
devices, 5–13, D–9
dual-redundant, 5–11
restrictions, D–27
EISA boards, 5–37, D–41
error messages, automatic, 6–16
FDDI Server cabinet, D–4
for 5.25-inch SBBs, D–13
hardware, D–1
highest availability, D–29
highest performance, D–29
initial steps, 8–63
intermixing SBB sizes, D–13
loaders, 5–21
mismatch, 3–2, 6–5
mixing disk and tape, D–9
nonredundant, 5–9, D–27
of environment variables, 5–28
optimal performance, D–28
power supply, A–6
read cache, D–27
server processor
device bus adapter, D–41
shelf, D–10
software and firmware, 5–1
Configuration (cont'd)
storage
automatic, 7–71, 7–73
manual, 5–13
typical
recommended, D–10
verifying, OpenVMS, 5–25
write-back cache, D–27
Configuration menu
See CFMENU
Configuration parameters
Are they preset?, D–31
Configuration rules
for FDDI Server cabinets, D–4
Console
diagnostic flow, 1–2
firmware commands, 1–7
Console commands, 1–7
CAT EL, 4–5
diagnostic and related, summarized, 4–1
KILL, 4–11
KILL_DIAGS, 4–11
memory, 4–6
MORE EL, 4–5
NET -IC, 4–10
net -s, 4–9
NETEW, 4–7
network, 4–8
set bootdef_dev, 5–28
set boot_osflags, 5–28
set envar, 5–28
show auto_action, 5–28
show config, 5–25
show device, 5–26
show envar, 5–28
show memory, 5–27
SHOW_STATUS, 4–12
test, 4–3
Console event log, 2–6
Console firmware
diagnostics, 2–11, 3–2
Console interfaces
switching between, 5–25
Console output, 5–38
Console port configurations, 5–37
Controller
See also HS1CP
Controller storage
explained, A–22
Core MIST, 7–2
hardware tests, 7–3
IBR, 7–2
program card validation, 7–2
CPU daughterboard, 5–33
Crash dumps, 1–7

C_SWAP
definition, 8–48

D

DAEMON, 7–4, 7–5
manually running, 7–5
manually stopping, 7–6
Data test patterns
DILX, 7–22
TILX, 7–46
DC power-up sequence, 2–10, 3–2
DEC VET, 1–7, 4–13
DECevent, 1–6, 6–26, C–6
example, C–9
invoking, 6–26
Defaults
DILX, 7–11
DELETE *container-name* command, E–30
DELETE FAILEDSET command, E–31
DELETE SPARESET command, E–32
DELETE *unit-number* command, E–33
Deleting a patch, 7–87
DESCRIBE command
FMU, 7–98
Deskew procedure, error log, C–146
Device bus adapter
configuring, D–41
Device channel processor
I/O demands, A–11
Device identifier
See SCSI ID, D–36
Device LEDs, 3–3, 6–10
SBB active LED, 3–3, 6–10
SBB fault LED, 3–3, 6–10
storage SBB faults, 3–3, 6–10
Device naming convention
SRM, 5–26
Device ports, A–15
running on fewer, 7–4
testing, 7–4
Device services
firmware, A–21
Device shelf status
power supply faults, 3–5, 6–12
power supply LEDs, 3–5, 6–12
shelf faults, 3–5, 6–12
single power supply
power supply faults, 3–6, 6–13
shelf faults, 3–6, 6–13
Devices
adding, 5–13
CFMENU, 7–74
configurations, D–9
configuring, automatic, 7–71, 7–73
configuring, manually, 5–13
initializing, 5–19
maximum
Devices
maximum (cont'd)
per port, D–29
moving between HS1CPs, 5–19
moving under the same controller, 5–20
nontransportable, 5–19
preserving device assignments, failover, 5–18
transportable, 5–19
Diagnostic
and execution monitor
See DAEMON
error messages, 6–18
registers, 7–3, A–12
tests, 5–6, 7–1
utility protocol
See DUP
Diagnostic flows
boot problems, 1–4
console, 1–2
errors reported by operating system, 1–4
power, 1–2
problems reported by console, 1–3
Diagnostics
command summary, 4–1
command to terminate, 4–11
console firmware-based, 2–11, 3–2
firmware power-up, 2–10, 3–2
power-up, 2–1, 3–1
power-up display, 2–1, 3–1
related commands, summarized, 4–1
ROM-based, 1–6, 4–1
serial ROM, 2–10
showing status of, 4–12
Digital Assisted Services (DAS), 1–8
DILX, 1–5, A–20
abort codes, 7–31
basic function test, 7–8
configuring all units, 7–27
data test patterns, 7–22
defaults, 7–11
defined, 7–7
end message display, 7–19
error codes, 7–31
error information packets, 7–20
example, 7–23
interrupting, 7–8
output messages, 7–15
performance summary, 7–29
running from maintenance terminal, 7–7
running from VCS, 7–7
running from virtual terminal, 7–8
test definition questions, 7–10
tests available, 7–8
user-defined test, 7–9
using all defaults, 7–23
using all functions, 7–24

DIRECTORY command, E–34
Disk bracket
  See also drive bracket
Disk drive
  warm swap, 8–38, 8–40
Disk drive removal
  warm swap, 8–38
Disk drive replacement
  warm swap, 8–40
Disk in-line exerciser
  See DILX
Disposal, write-back cache batteries, 8–57
DRAB
  See Shared memory
DRAM
  See Shared memory
Drive bracket
  See also metal housing
DSSI host interconnection
  supported protocols, A–20
DSSI node number, 5–9, 5–11
Dual-redundant HS1CP
  downtime, 6–1
  failover, 5–16, 6–1, A–14, A–22
  initialization, 5–7
  on separate hosts, 5–12
  service consideration, 6–1
DUP, 5–14, A–20

E

ECC, A–15
ECU, D–41
  ecu command, 5–25, 5–36
  how to run, D–41
  invoking console firmware, 5–35, D–41
  procedure for running, 5–35, D–41
  procedures, 5–37
  starting up, 5–35
EDC, 7–2, 7–3, A–15
EIA–423 terminal port, 5–7
EISA boards
  configuring, 5–37, D–41
EISA bus
  features of, 5–35
  problems at power-up, 2–7
  troubleshooting, 2–7
  troubleshooting tips, 2–8
EISA configuration utility
  EISA, D–41
EISA option card, 8–20
EISA slots on system board, 8–20
Electrostatic discharge
  grounding procedure, 8–2, D–2
Electrostatic discharge protection
  See ESD, 8–1, D–2
End message display
  DILX, 7–19
  TILX, 7–44
Entering a patch, 7–86
Environment variables
  configuring, 5–28
  setting and examining, 5–28
Environment variables set during system configuration, 5–28
ERF, C–1
  example, C–2, C–7
  invoking, 6–26
ERF/uerf, 1–6
Error
  handling, 1–6
  logging, 1–6
  report formatter (ERF), 1–6
Error correction code
  See ECC
Error detection code
  See EDC
Error logging, 1–5, 6–26
  controller model, 6–27
  DECevent, 6–26
  deskewing log, C–146
  ERF, 6–26
  firmware, A–21
  FMU, 1–5, 6–26
  HS1CP, C–1
  HS1CP-series, C–6
  translations, 6–26
Error messages, 6–15
  automatic, 6–15
  cache module, 6–21
  CER, 6–15
  CLI, automatic, 6–16
  CLI, interactive, 6–24
  configuration, automatic, 6–16
  diagnostics, 6–18
  EVL, 6–25
  failover, 6–18
  failure indications, 6–23
  initialization, 6–18
  LFL, 6–24
  NVPM, 6–19
  self-test, 6–18
  shelf, 6–21
  system information page, 6–24
  system information structure, 6–24
Error reporting, 6–3
  types, 6–3
Errorlog Report Formatter
  See ERF
ESD, 8–1, D–2
  grounding procedure, 8–2, D–2

Ethernet
external loopback, 4–3
Event logging
See EVL
Event logs, 1–6
EVL, 6–25, 7–94
Example
CFMENU, 7–73
changing the replacement policy of a RAIDset
See example., E–139
creating a transportable unit from a disk device
See example, E–139
creating a unit from a disk device
See example, E–138
creating a unit from a disk device and setting it write protected
See example, E–138
creating a unit from a five-member RAIDset
See example, E–138
creating a unit from a four-member stripeset.
See example, E–138
creating a unit from a tape device
See example, E–138
DECevent, C–9
deleting the unit, stripeset and disks associated with a stripeset
See example, E–140
DILX, 7–23
ERF, C–2, C–7
FLS, 7–91
FMU, 7–99
renumbering disk unit 0 to disk unit 100
See example, E–139
TILX, 7–47
write protecting an existing unit
See example, E–139
EXEC, 7–3, A–20
Exercisers
See DILX
See TILX
Exit
CFMENU, 7–84
DILX, 7–8
failover configuration, 5–17
FMU, 7–99
TILX, 7–34
VTDPY, 7–53
EXIT command, E–35
FMU, 7–99
Exiting Code Patch, 7–89

F

Fail-safe loader, 2–8
activating, 2–8
power-up using, 2–8
Failedset, 5–9
CFMENU, 7–78
Failover, 5–16, A–14
cache, 6–2
copying configuration, 5–11
correcting mismatch, 5–18
error messages, 6–18
exiting, 5–17
firmware, A–22
initialization, 5–18
mismatch, 5–18
preferred paths, 5–18
preserving device assignments, 5–18
RAIDset, 6–2
reviving failed HS1CP, 5–17
setup for, 5–16
shared commands, 5–16
testing for, 5–17
time required for, 5–17
Failure indications
error messages, 6–23
Fan assembly, A–6
Fault management
DECevent, 6–26, C–9
ERF, 6–26, C–2, C–7
firmware, A–21
FMU, 6–26
SWEAT, 6–26
Fault Management Utility
See FMU
FDDI Adapter
Physical Description, A–11
FDDI adapter–DEFEA or DEFPA module, A–11
FDDI Server
configuration, 5–1
FDDI Server cabinet
configuration, D–4
configuration rules, D–4
FDDI Server Component Access, 8–2
Firmware
CLI, A–20
console commands, 1–7
controller storage addressing, A–22
description, A–19
levice services, A–21
diagnostics, 4–1
downloading of, 7–3
DUP, A–20
error logging, A–21
event reporting, A–21
See Also Error Logging
EXEC functions, A–20
failover, A–22
fault management, A–21
host interconnect functions, A–20
host protocol, A–20
host storage addressing, A–24

Firmware (cont'd)
init functions, A–19
local programs, A–20
operator interface, A–20
power-up diagnostics, 2–10, 3–2
self-test, A–19
value-added, A–22
version restriction, D–27
Firmware executive
See EXEC
Firmware Licensing System
See FLS
Firmware version
verifying, 8–58
Flashing codes, HS1CP, 6–5
Floppy drive
LEDs, 2–7
FLS, 1–5, 7–90, A–21
example, 7–91
invoking, 7–90
license key, 7–90
main menu, 7–90
messages, 7–92
options, 7–90
FMU, 1–5, 7–94, A–21
DESCRIBE command, 7–98
example, 7–99
EXIT command, 7–99
invoking, 7–94
options, 7–94
restrictions, 7–94
SET command, 7–94
SHOW command, 7–97
Front panel
fault notification, 7–2, 7–3
HS1CP, A–12

G

General information, 1–4
Green LED, 5–6, 6–5, 7–1
Grounding
ESD, 8–2, D–2

H

H8571–J adapter, 8–5, D–3
Hardware
HS1CP
bus exchangers, A–15
device ports, A–15
diagnostic registers, A–12
dual controller port, A–14
front panel, A–12
host interface, A–17
LEDs, A–12
nonvolatile memory, A–14
policy processor, A–12
Hardware
HS1CP (cont'd)
program card, A–12
RAID, A–15
reset button, A–13
shared memory, A–15
value-added, A–15
write-back cache module, A–15
server processor
maintenance terminal, A–12
OCP, 2–11, A–4
storage devices, A–3
Hardware, HS1CP
front panel, 3–2, 6–3
maintenance terminal, 5–7
HELP command, E–36
Host cable, D–20
length, D–20
Host interface, A–17
DSSI, A–17
HS1CP to internal bus, D–20
SCSI, A–18
testing, 7–4
Host logical units, A–24
Host port path, 5–10, 5–12
Host protocol, A–20
Host storage
all HS controllers, A–24
HS operating firmware
See Firmware
HS110 description, 5–6
HS111 upgrade
BA350–M shelf
SCSI cabling, 8–67
SCSI cable installation, 8–67
setting initial parameters, 8–63
HS1CP
codes, 6–5
configuring, D–27
device support, D–24
error logging, 6–27
flashing codes, 6–5
front panel, 1–5
green LED, A–13
hardware support, D–24
host protocol, A–20
LEDs, 3–7, A–12
operating system, D–24
reset button, A–13
warm swap, 8–48
write-back cache module, A–15
HS1CP front panel, 3–2, 6–3, A–12
amber LEDs, 3–2, 6–4
codes, 3–3
fault notification, 3–2, 6–5
green LED, 3–2, 6–4
normal operation, 3–2, 6–4
reset button, 3–2, 6–4

HS1CP front panel (cont'd)
solid codes, 6–5
HS1CP ID, 5–9, 5–11, 5–18
HS1CP program card
removing, 5–6
HS1CP warm swap, 8–48
HS1xx maintenance
accessing cabinet interior, 8–2
cabinet preparation, 8–2
opening cabinet, 8–2
remove cabinet power, 8–2
HS211/111 conceptual model, 5–6
HS211/111 description, 5–5
HS221–AX upgrade
verifying firmware version, 8–58
HS221/121 conceptual model, 5–5
HS221/121 description, 5–4
HS241 conceptual model, 5–2
HS241 description, 5–2
HS280 conceptual model, 5–3
HS280 description, 5–3

I

I/D cache, 7–3, 7–4
I/O bus, EISA features, 5–35
i960, 7–2
IBR, 7–2
ILX error codes
DILX, 7–31
TILX, 7–52
ILX error information packets
DILX, 7–20
TILX, 7–44
Information resources, 1–7
Initial boot record
See IBR
Initialization, 4–13
BIST, 7–2
causes of, 5–6, 7–1
CFMENU, 7–80
command, 5–13
described, 7–1
device port, 7–4
dual-redundant HS1CP, 5–7, 5–18
error messages, 6–18
failover, 5–18
host port, 7–4
nontransportable devices, 5–19
RAID hardware, 7–5
subsystem, 5–7
tests performed, 7–1
time required, 7–1
transportable devices, 5–19
value-added functions, 7–5
INITIALIZE command, E–37
Installation
batteries, write-back cache, 8–58
EISA/PCI option cards, 8–20
read cache
write-back cache, 8–56
setting parameters, 8–63
write-back cache batteries, 8–58
Installing components
preparation for, 8–11
required equipment, 8–11
Installing server processors
server processor installation, 8–35
Installing SIMMs, 8–16, 8–18
configuration, 8–18
rules, 8–18
Installing storage devices
CD–ROM drive, 8–30
Interfaces
switching between, 5–25
Interrupt lines
and EISA, 2–8
Invoking the Code Patch Utility, 7–85
IRQs
and EISA, 2–8

J

Jumper
device shelves, D–8

K

Key, licensed features, 7–90
KFESA option card, 8–20
KILL command, 4–11
KILL_DIAGS command, 4–11

L

Lamp test, E–39
Last failure logging
See LFL
LEDs
CD–ROM drive, 2–7
device, 3–3, 3–4, 6–10, 6–11
floppy drive, 2–7
front panel
HS1CP, A–12
HS1CP, 3–7, 6–5
power supply, 3–5, 6–12
SBB, 3–3, 3–4, 6–10, 6–11
storage device, 2–6, 3–3
StorageWorks, 2–6
LFL, 6–24, 7–94
Licensed features, 7–90
Listing Patches, 7–87

Loaders, 5–21
Local programs, A–20
LOCATE CANCEL command, E–39
LOCATE command, E–39
LOCATE DISKS command, E–39
LOCATE entity, E–40
LOCATE LOADERS command, E–39
LOCATE OPTICALS command, E–39
LOCATE PTL SCSI-location command, E–40
LOCATE TAPES command, E–39
LOCATE UNITS command, E–40
Logical Unit Number
  See LUN
Logs
  event, 1–6
Loopback tests, 1–7
  COM2 and parallel ports, 4–3
  command summary, 4–2
LUN
  controller perspective, A–23

M

Maintenance port, 5–7
Maintenance strategy, 1–1
  service tools and utilities, 1–6
Maintenance terminal, 1–5, 5–7
  server processor, A–12
    modified modular jack
      See MMJ.
Mating guide
  See also trilink connector, 5–10
Maximum devices
  per port, D–29
Media tray
  3.5-Inch floppy drive, A–5
  CD–ROM drive, A–5
MEMBERSHIP=*number-of-members*, E–59
Memory capacity
  and maximum, A–9
  memory, A–9
Memory command, 4–6
Memory module
  configuration, 5–33
  displaying information for, 5–27
  minimum and maximum, 5–33
  See also SIMM, 8–16
Memory tests, 2–3
Memory, main
  exercising, 4–6
Messages, 7–88
  cache module, 6–21
  CFMENU, 7–82
  CI I, automatic, 6–16
  CLI, interactive, 6–24
  configuration, automatic, 6–16
  diagnostics, 6–18
Messages (cont'd)
  DILX, 7–15
  EVL, 6–25
  failover, 6–18
  failure indications, 6–23
  FLS, 7–92
  initialization, 6–18
  LFL, 6–24
  NVPM, 6–19
  self-test, 6–18
  shelf, 6–21
  system information page, 6–24
  system information structure, 6–24
  TILX, 7–40
Metadata, 5–13, 5–19
  defined, E–37
Metal housing
  See also drive bracket
MIRROR *disk-device-name1 container-name*
    command, E–41
Mirrorset, 5–9
MIST, 7–2, 7–4
  See also Core MIST
  See also DAEMON
Mixing disk and tape, D–9
MMJ adapter, 8–5, D–3
Module integrity self-test
  See MIST
Modules
  CPU, 5–33
  memory, 5–33
  motherboard, 5–34
MORE EL command, 4–5
Motherboard, 5–34
  memory banks, A–9
  SIMM connectors, A–9
Moving devices
  between HS1CPs, 5–19
  under the same controller, 5–20
MSCP, 5–10, 5–11
MSCP timeout, 5–15

N

NET -IC command, 4–10
NET -S command, 4–9
NETEW command, 4–7
network command, 4–8
Nonredundant
  configuration, D–27
Nonredundant controller
  downtime, 6–1
  service consideration, 6–1
Nontransportable devices, 5–19
Nonvolatile memory
  See NVMEM.

Nonvolatile parameter memory
See NVPM
NOTRANSPORTABLE qualifier, 5–13
NVMEM, A–14
battery, A–14
NVPM, 6–19
error messages, 6–19

O

OCP, 2–11, A–4
server processor, 2–11, A–4
OCP controls and indicators
description, 2–12, A–5
OpenVMS
cluster size, 5–15
MSCP timeout, 5–15
polling parameters, 5–15
shadow member timeout, 5–15
shadowsets, 5–15
storageset size, 5–15
support, 5–15
TMSCP timeout, 5–15
Operating system
boot failures, reporting, 1–4
crash dumps, 1–7
exercisers, 1–7
support, 5–15
Operating system, controller
See EXEC
Operator control panel, 2–12, A–5
See OCP
See OCP.
server processor, 2–11, A–4
Operator interface
firmware, A–20
maintenance terminal, 5–7
server processor
maintenance terminal, A–12
virtual terminal, 5–14
Operator interfaces, switching between, 5–25
Option cards
configuring, 8–22
Options
system bus, 5–31

P

Parameters
setting, 8–63
Parameters, initial HS1CP, 5–9, 5–11
Passthrough, 5–21
CFMENU, 7–79
Path, host port, 5–10, 5–12
PCI option card, 8–20
Performance
configuration, D–28, D–29
Performance summary
DILX, 7–29
TILX, 7–50
Policy processor, A–12
Polling parameters, 5–15
Port/target/LUN
See PTL
Power configurations
dual cabinet power configuration, D–19
dual shelf power configuration, D–17
full power redundancy, D–18
nonredundant power, D–15
partially redundant power, D–16
Power controller
See CDU., 8–36
Power problems
diagnostic flow, 1–2
Power supply
configuration, A–6
failure, 8–41
redundant, configuring, A–6
Power-on tests, 2–9, 3–2
Power-up
diagnostics, 2–10, 3–2
displays, interpreting, 2–1, 3–1
screen, 2–5
sequence, 2–9, 3–2
AC, 2–10, 3–2
DC, 2–10, 3–2
Preserving device assignments, failover, 5–18
Program card, A–12
contents, 7–2
removing, 5–6, 5–17, 7–1
self-test, 7–2
validation, 7–2
version restriction, D–27
PTL
controller perspective, A–23

Q

Quiesce
definition, 8–38, 8–49
troubleshooting, 8–41

R

RAID
FLS, 7–90
hardware, A–15
RAIDset, 5–9
CFMENU, 7–76
failover, 6–2, A–16
size, 5–15, E–11

Read cache, 8–54
  cache manager, 7–5
  configuration, D–27
  installing, 8–56
  removing, 8–54
  replacing, 8–56
  testing, 7–5
Read only test
  TILX, 7–35
Rear panel ports and slots
  Description, A–8
REDUCE *disk-device-name1 [disk-device-nameN]* command, E–43
REDUCED, E–12
Redundant Array of Inexpensive Disks
  See RAID
Removal
  batteries, write-back cache, 8–57
  HS1CP program card, 5–6
  program card, 5–6, 5–17, 7–1
  read cache
    write-back cache, 8–54
  write-back cache batteries, 8–57
Removal and replacement
  Cache Memory Battery, 8–46
  Cache Memory Board, 8–46
  CPU daughterboard, 8–25
  fans, 8–30
  FDDI adapter, 8–25
  hard drive, 8–22
  HS1CP board, 8–46
  motherboard, 8–25
  OCP, 8–32
  power supplies, 8–32
  speaker, 8–24
Removal and Replacement
  BA350–M Shelf, 8–42
  BA350–S Shelf, 8–43
REMOVE=*disk-container-name*, E–67
REMOVE=*disk-device-name*, E–60
Removing a device
  warm swap, 8–38
Removing components
  preparation for, 8–11
  required equipment, 8–11
Removing SIMMs, 8–16
Removing the cabinet panels, 8–4
Removing the cover panels, 8–14
Removing the front cover panel, 8–13
Removing the rear top cover panel, 8–13
Removing the server processor, 8–12
RENAME command, E–45
REPLACE=*disk-container-name*, E–67
REPLACE=*disk-device-name*, E–59
Replacement
  batteries, write-back cache, 8–58
  read cache
    write-back cache, 8–56
  shelf, 8–43
  write-back cache batteries, 8–58
Replacement procedure
  internal bus cables
    trilink connector block
      terminators, 8–69
Replacing
  blowers, 8–44
  shelf power supplies, 8–45
  shelves
    installing, 8–43
    removing, 8–42, 8–43
Replacing a device
  warm swap, 8–40
Replacing power SBBs
  cold swap method, 8–45
  warm swap method, 8–45
Replacing SIMMs, 8–16, 8–18
Reset button, 5–6, 5–17, 7–1, 7–6
RESTART OTHER_CONTROLLER command, E–46
RESTART THIS_CONTROLLER command, E–48
Restrictions
  cache size, D–27
  cache type, D–27
  CFMENU, 7–73
  dual-redundant controller, D–27
  firmware version, D–27
  FLS, 7–90
  FMU, 7–94
  hardware, D–1
  tape devices, A–22
  warm swap, A–22
RETRY_ERRORS UNWRITEABLE_DATA command, E–50
ROM-based diagnostics (RBDs), 1–6
  running, 4–1
  utilities, 4–1
RUN command, E–51
RZ*xx*
  LED indicators, A–10

S

Saving storage configuration, 8–60
SCS node name, 5–11
SCSI hosts
  supported protocols, A–20
SCSI ID, D–36
  modifying, D–36
SCSI identifier
  See SCSI ID
Self-test, 1–4, 7–5, A–19
  See also DAEMON
  error messages, 6–18

Self-test (cont'd)
running, 7–5
stopping, 7–6
SELFTEST OTHER_CONTROLLER command, 7–5, E–52
SELFTEST THIS_CONTROLLER command, 7–5, E–54
Serial ports, 5–38
Serial ROM diagnostics, 2–10
Server processor
device identifier
See SCSI ID.
hard drive
See system disk
shutdown, 8–11
terminal port
connecting to, 8–4, D–3
Server processor cover panels
front cover, removing, 8–13
rear top cover, removing, 8–13
Server processor front view—exploded, A–4
Server processor OCP, 2–11, A–4
Server processor rear view—exploded, A–6
Service
tools and utilities, 1–6
SET command
FMU, 7–94
set command (SRM), 5–28
SET *disk-container-name* command, E–56
SET FAILOVER command, E–57
SET *mirrorset-container-name* command, E–59
SET NOFAILOVER command, E–62
SET OTHER_CONTROLLER command, E–63
SET *RAIDset-container-name* command, E–66
SET THIS_CONTROLLER command, E–68
SET *unit-number* command, E–70
Shadow member timeout, 5–15
Shadowsets, 5–15
Shared memory, 7–3, A–15
testing, 7–3
Shelf
configurations, D–10
error messages, 6–21
replacement, 8–43
Shelf blowers
replacing, 8–44
SHOW *cdrom-container-name* command, E–80
SHOW CDROMS command, E–79
SHOW command
FMU, 7–97
show command (SRM), 5–28
show configuration command (SRM), 5–25
show device command (SRM), 5–26
SHOW DEVICES command, E–81
SHOW *disk-container-name* command, E–83
SHOW DISKS command, E–82
SHOW FAILEDSET command, E–84
SHOW *loader-container-name* command, E–86
SHOW LOADERS command, E–85
Show memory
verifying memory, 8–19
show memory command (SRM), 5–27
SHOW *mirrorset-container-name* command, E–89
SHOW MIRRORSETS command, E–87
SHOW *optical-container-name* command, E–91
SHOW OPTICALS command, E–90
SHOW OTHER_CONTROLLER command, E–92
SHOW PASSTHROUGH command, E–93
SHOW *passthrough-container-name* command, E–94
SHOW *raidset-container-name* command, E–97
SHOW RAIDSETS command, E–95
SHOW SPARESET command, E–98
SHOW STORAGESETS command, E–99
SHOW *stripeset-container-name* command, E–103
SHOW STRIPESETS command, E–101
SHOW *tape-container-name* command, E–105
SHOW TAPES command, E–104
SHOW THIS_CONTROLLER command, E–106
SHOW *unit-number* command, E–108
SHOW UNITS command, E–107
SHOW_STATUS command, 4–12
SHUTDOWN OTHER_CONTROLLER command, E–109
SHUTDOWN THIS_CONTROLLER command, E–111
SIMM connectors, A–9
SIMMs, 5–33, A–9
installation rules, 8–18
installing, 8–16
removing, 8–16
replacing, 8–18
show memory, 8–19
SHOW MEMORY command, 8–19
troubleshooting, 2–3
verify memory, 8–19
Sliding the server processor out of the cabinet, 8–12
Software, HS controller
See Firmware
Solid codes, HS1CP front panel, 6–5
Spareset, 5–9
CFMENU, 7–78
defined, E–13
SRM interface
switching to ARC from, 5–25
SROM memory tests, 2–3
Storage
controller perspective, A–22
controller PTL, A–23
host perspective, A–24
how addressed, A–22

Storage configuration
save, 8–60
Storage device LEDs, 2–6, 3–3
Storageset
CFMENU, 7–76, 7–78
defined, E–99
size, 5–15
StorageWorks
LEDs, 2–6
StorageWorks FDDI Server
power, configuring, D–15
Stripeset, 5–9
CFMENU, 7–76
size, 5–15
Subsystem, initialization, 5–7
SWAP_L interrupt signal, 8–41
SWEAT, 6–26
System
options, 5–31
power-up display, interpreting, 2–1, 3–1
troubleshooting categories, 1–1
System bus
location, 5–32
System disk, 8–22, D–36
modifying SCSI ID, D–36
SCSI ID
modifying
System information page
error messages, 6–24
System information structure
error messages, 6–24
System shutdown procedure, 8–11

T

Tape drives
warm swap, 8–41
Tape in-line exerciser
See TILX
Target
controller perspective, A–23
Terminal port
connecting to, 8–4, D–3
Terminal setup
CFMENU, 7–82
maintenance terminal, 5–7
virtual terminal, HS1CP, 5–14
Test
acceptance test, 5–14
after installation, 5–14
BIST, 7–2
cache module, 7–4
command disk, 5–24
core MIST, 7–2
DAEMON, 7–4
disk device, 5–14, 7–7, A–20
error messages, self-test, 6–18
Test (cont'd)
failover, 5–17
RAID, 7–5
read cache, 7–5
self-test, 1–4, 7–5, A–19
tape device, 5–14, 7–33, A–20
write-back cache, 7–4
TEST command, 4–3
Test definition questions
DILX, 7–10
TILX, 7–36
Testing
acceptance, 4–13
memory, 4–6
with DEC VET, 4–13
TILX, 1–5, A–20
abort codes, 7–51
basic function test, 7–34
data test patterns, 7–46
defined, 7–33
end message display, 7–44
error codes, 7–52
error information packets, 7–44
example, 7–47
interrupting, 7–34
output messages, 7–40
performance summary, 7–50
read only test, 7–35
running from maintenance terminal, 7–33
running from VCS, 7–33
running from virtual terminal, 7–33
test definition questions, 7–36
tests available, 7–34
user-defined test, 7–35
using all defaults, 7–47
using all functions, 7–48
TMSCP, 5–10, 5–11
TMSCP timeout, 5–15
Tools, 1–6
console commands, 1–6, 1–7
crash dumps, 1–7
DEC VET, 1–7
DECevent, 1–6
ERF/uerf, 1–6
error handling, 1–6
log files, 1–6
loopback tests, 1–7
RBDs, 1–6
Training, 1–8
Transportable devices, 5–19
TRANSPORTABLE qualifier, 5–13
Trilink connector, 5–10
Trilink connector block
installation, 8–69
Troubleshooting, 1–1, 6–15
See also Diagnostics; diagnostic flow
actions before beginning, 1–1
boot problems, 1–4

Troubleshooting (cont'd)
categories of system problems, 1–1
crash dumps, 1–7
diagnostic flows, 1–2, 1–3, 1–4
EISA problems, 2–7
error logs, 6–2
error messages, 6–15
error report formatter, 1–6
errors reported by operating system, 1–4
interpreting error beep codes, 2–1
power problems, 1–2
problem categories, 1–1
problems getting to console, 1–2
problems reported by the console, 1–3
SIMMs, 2–3
visual indicators, 3–3, 6–2, 6–10
with DEC VET, 1–7
with loopback tests, 1–7
with operating system exercisers, 1–7
with ROM-based diagnostics, 1–6
Troubleshooting Basics, 1–6
Troubleshooting HS1CP
using front panel, 3–2, 6–3

U

Units
adding, E–16
CFMENU, 7–81
UNMIRROR *disk-device-name* command, E–113
User-defined test
DILX, 7–9
TILX, 7–35

V

Value-added functions
See also RAID
firmware, A–22
hardware, A–15
initialization, 7–5
VAXcluster console system
See VCS
VCS, 5–14, 7–7, 7–8, 7–33
Verifying memory
SHOW MEMORY command, 8–19
Virtual terminal, 1–5, 5–14
DUP, 5–14
FYDRIVER, 5–14
HS1CP, 5–14
VTDPY, 1–5, 7–53, A–20
help, 7–70

W

Warm swap
controller, 1–5, A–20
definition, 8–38
disk and tape SBB replacement
device addition, 8–40
for the HS1CP, 8–48
for cache modules, 8–48
of devices, 8–38
tape drives, 8–41
Write hole, A–15, A–16
Write-back cache
battery, 5–18, 6–2, A–15
battery, discharging, A–16
cache manager, 7–4
charge time, battery, A–16
configuration, D–27
disposing of batteries, 8–57
error messages, 6–21
FLS, 7–90
installing, 8–56
installing batteries, 8–58
power failure, A–15
removing, 8–54
removing batteries, 8–57
replacing, 8–56
replacing batteries, 8–58
service consideration, 6–2
testing, 7–4