

StorageWorks RAID Array 110 Subsystem

Service Guide

Order Number: EK-SZ200-SG. A01

Revision/Update Information: This is a new manual.

July 1993

First Printing, July 1993

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

Restricted Rights: Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of DFARS 252.227-7013, or in FAR 52.227-19 or FAR 52.227-14, Alt. III, as applicable.

© Digital Equipment Corporation 1993

All Rights Reserved.
Printed in the U.S.A.

FCC Notice: The equipment described in this manual 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 computing device pursuant to Part 15 of FCC Rules, which are designed to provide reasonable protection against such radio frequency interference when operated in a commercial environment. Operation of this equipment in a residential area may cause interference, in which case the user at his own expense may be required to take measures to correct the interference.

The postpaid READER'S COMMENTS form on the last page of this document requests the user's critical evaluation to assist in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

BA350, DEC, DECnet, DECnet-DOS, DECpc, DEC RAID, DECserver, DECstation, StorageWorks, DECwindows, MicroVAX, OpenVMS, VAX, VAXcluster, VAX DOCUMENT, VAXstation, VMS, VT, and the DIGITAL logo.

The following are third-party trademarks:

ASPI is a registered trademark of Adaptec, Inc.
DPT and SmartROM are registered trademarks of Distributed Processing Technology.
Intel is a registered trademark of Intel Corporation.
Motorola is a registered trademark of Motorola, Inc.
MS-DOS and Microsoft are registered trademarks of Microsoft Corporation.
Novell and NetWare are registered trademarks of Novell, Inc.
RAID Manager is a trademark of NCR Corporation.
SCO is a registered trademark of The Santa Cruz Operation, Inc.
UNIX is a registered trademark of UNIX System Laboratories, Inc.
Windows is a trademark of Microsoft Corporation.
X Window System is a trademark of the Massachusetts Institute of Technology.

All other trademarks and registered trademarks are the property of their respective holders.

This document was prepared using VAX DOCUMENT, Version 2.1.

Contents

Preface	xiii
1 Introduction	
1.1 Product Description	1-1
1.1.1 Overview	1-1
1.1.2 Product Highlights	1-3
1.1.2.1 Array Controller Features	1-3
1.1.2.2 Subsystem Features	1-4
1.1.3 Product Attributes	1-4
1.1.3.1 Data Reliability	1-4
1.1.3.2 Redundancy	1-4
1.1.3.3 Data Availability	1-5
1.1.3.4 Performance	1-5
1.1.3.5 Flexibility and Capacity	1-5
1.1.4 General Conclusions	1-6
1.2 StorageWorks RAID Array 110 Subsystem Component Descriptions	1-6
1.2.1 Description of the Subsystem	1-6
1.2.1.1 BA35X-VA Vertical Mounting Kit	1-7
1.2.1.2 BA350-EA Shelf	1-7
1.2.1.3 HSZ10-AA Controller (Disk Array Controller)	1-7
1.2.1.4 Features	1-8
1.2.1.5 Firmware	1-9
1.2.1.6 Software	1-9
1.2.2 System Building Blocks (SBBs)	1-9
1.2.3 Expansion Unit Description	1-10
1.2.4 StorageWorks RAID Array 110 Utilities	1-10
1.2.5 Enhanced SCSI Driver Support	1-10
1.2.6 SCSI Interconnects/Host Adapters	1-11
2 Technical Information	
2.1 RAID Overview	2-1
2.1.1 What Is a Disk Array?	2-1
2.1.1.1 Description of RAID 0	2-2
2.1.1.2 Description of RAID 1	2-3
2.1.1.3 Description of RAID 3	2-4
2.1.1.4 Description of RAID 5	2-5
2.1.2 Key Concepts	2-6
2.1.2.1 Array Channels	2-6
2.1.2.2 Logical Units (LUNs)	2-6
2.1.2.3 Drive Groups	2-7
2.1.2.4 Drive Ranks	2-8
2.1.2.5 Partitions	2-8

2.1.2.6	Reconstruction	2-9
2.1.2.7	Regeneration	2-9
2.2	StorageWorks RAID Array 110 Subsystem Specifications	2-10
2.2.1	General Specifications	2-10
2.2.2	Power Unit Specifications	2-10
2.2.3	Environmental Stabilization	2-11
2.2.4	Environmental Specifications	2-12
2.3	StorageWorks RAID Array 110 Subsystem Base Configuration	2-13
2.3.1	SCSI IDs for the StorageWorks RAID Array 110 Subsystem	2-13
2.3.2	HSZ10-AA Controller Location and SCSI Address	2-15
2.4	Physical Drive Map	2-16
2.5	Logical Drive Map	2-16
2.5.1	DWZZA-VA Bus Adapter	2-16

3 Installation

3.1	Basic Configuration Installation	3-1
3.1.1	Connecting to a Host	3-1
3.1.1.1	Terminating the SCSI Bus	3-2
3.1.1.2	Connecting to a 16-Bit Differential Host/Adapter	3-2
3.1.1.3	Connecting to an 8-Bit Differential Host/Adapter	3-2
3.1.1.4	Connecting to an 8-Bit Single-Ended Host/Adapter	3-3
3.1.1.5	Maintaining Bus Continuity	3-3
3.1.2	Verifying Cables and Connectors	3-6
3.1.3	Powering On the Subsystem	3-6
3.1.4	Functional Verification	3-7
3.2	StorageWorks RAID Array 110 Utility for Novell NetWare Installation	3-8
3.2.1	About Installation	3-8
3.2.1.1	StorageWorks RAID Array 110 Utility for Novell NetWare Files	3-9
3.2.2	When to Install the StorageWorks RAID Array 110 Utility for Novell NetWare	3-10
3.2.3	Before You Install the StorageWorks RAID Array 110 Utility for Novell NetWare	3-10
3.2.4	Running the Installation Program	3-10
3.2.4.1	Using an MS-DOS Partition	3-10
3.2.4.2	Without Using an MS-DOS Partition	3-11
3.2.5	Installation Error Messages	3-13
3.3	Installation of the StorageWorks RAID Array 110 Utility for MS-DOS	3-15
3.3.1	StorageWorks RAID Array 110 Utility for MS-DOS Kit	3-15
3.4	StorageWorks RAID Array 110 Utility for SCO UNIX Installation	3-15
3.4.1	About Installation	3-16
3.4.1.1	Installing the StorageWorks RAID Array 110 Utility for SCO UNIX Parity Check/Repair Utility	3-16
3.4.1.2	Installing the RAID Monitor Utility	3-17
3.4.1.3	RAID Status Utility	3-17
3.5	StorageWorks RAID Array 110 Utility for OpenVMS VAX Installation Procedure	3-18
3.5.1	Installation Requirements	3-18
3.5.2	Preparing Your System for the Installation	3-19
3.5.3	Using the InfoServer During the Installation	3-21
3.5.4	Using the InfoServer System to Update Your System	3-23

3.5.5	Matching Update Procedures to System Configurations	3-25
3.5.5.1	Updating VAXcluster Environments	3-25
3.5.5.2	Updating a Local Area VAXcluster System with One Boot Server and Two System Disks	3-26
3.5.6	Installing the StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit	3-26
3.5.7	Tasks to Perform After the Installation	3-28
3.6	Deinstalling the StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit	3-29
3.6.1	Post Deinstallation Procedure	3-31

4 Diagnostics and Software

4.1	StorageWorks RAID Array 110 Utilities	4-1
4.1.1	What are the StorageWorks RAID Array 110 Utilities?	4-1
4.1.2	Using the StorageWorks RAID Array 110 Utility for OpenVMS VAX	4-2
4.1.3	Using the StorageWorks RAID Array 110 Utility for MS-DOS	4-2
4.1.4	Using the StorageWorks RAID Array 110 Utility for Novell NetWare	4-2
4.2	Getting Started with the StorageWorks RAID Array 110 Utility for OpenVMS VAX	4-3
4.2.1	Before You Begin	4-3
4.2.2	Device Names	4-4
4.3	ACE Window Features	4-4
4.3.1	ACE Drive Window	4-5
4.3.2	ACE Logical Unit Window	4-5
4.3.3	Using ACE	4-5
4.3.3.1	Getting Help with ACE	4-6
4.3.3.2	Deleting a LUN Using ACE	4-7
4.3.3.3	Creating a LUN Using ACE	4-9
4.3.3.4	Exiting the ACE Window	4-12
4.3.4	Other "Getting Started" Tasks	4-12
4.3.4.1	Description of APC and APR	4-12
4.3.4.2	When You Should Use APC and APR	4-12
4.3.4.3	Using APC and APR	4-12
4.3.5	Starting the StorageWorks RAID Array 110 Utilities Under MS-DOS, SCO UNIX, and NetWare	4-12
4.3.5.1	Starting the StorageWorks RAID Array 110 Utility for MS-DOS	4-13
4.3.5.2	Starting the StorageWorks RAID Array 110 Utility for Novell NetWare	4-13
4.3.6	Moving Around in the StorageWorks RAID Array 110 Utilities Under MS-DOS, SCO UNIX, and NetWare	4-14
4.3.7	Utility Menus Under MS-DOS, SCO UNIX, and NetWare	4-14
4.3.7.1	Disk Array Services Window	4-14
4.3.7.2	Array Controller Configuration Window	4-16
4.3.7.3	StorageWorks RAID Array 110 Utilities Road Map	4-19
4.3.8	RAID Manager LUN and Drive Status	4-20
4.3.9	RAID Manager Device Names	4-22
4.3.9.1	StorageWorks RAID Array 110 Utility for MS-DOS Device Names	4-22
4.3.9.2	NetWare Device Name	4-22
4.4	Array Configuration Editor (ACE)	4-24

4.4.1	What is ACE?	4-24
4.4.2	Starting ACE	4-24
4.4.2.1	ACE Screen Window	4-25
4.4.2.2	Navigating in ACE	4-26
4.4.2.3	Quitting ACE	4-27
4.4.3	Drive Window	4-27
4.4.3.1	Drive Matrix	4-27
4.4.3.2	ACE Drive Status	4-28
4.4.3.3	Drive Window Options	4-28
4.4.4	Logical Unit Window	4-29
4.4.4.1	Logical Unit Display	4-29
4.4.4.2	ACE LUN Status	4-30
4.4.4.3	Logical Unit Window Options	4-31
4.5	Array Parity Check (APC) and Array Parity Repair (APR) Utilities for OpenVMS VAX	4-32
4.5.1	Overview	4-32
	APC	4-33
	APR	4-34
4.6	Array Status Monitor (ASM) Utility for OpenVMS VAX	4-35
4.6.1	Overview	4-35
	ASM	4-36

5 Operations

5.1	StorageWorks RAID Array 110 Subsystem Operations	5-1
5.1.1	StorageWorks RAID Array 110 Subsystem Monitoring Features	5-1
5.1.1.1	Monitoring Through the HSZ10-AA Controller	5-1
5.1.1.2	Monitoring Through the StorageWorks Shelf	5-4
5.1.2	User Monitoring Methods	5-4
5.1.2.1	Monitoring Operation Using the StorageWorks RAID Array 110 Utilities	5-4
5.1.2.2	Monitoring Operation Using LED Indicators	5-5
5.2	Adding LUNs to NetWare	5-11
5.2.1	Using the Install Command	5-11
5.2.2	Adding LUNs to Existing Volumes	5-11
5.2.3	Performance Notes	5-11
5.3	Configuring LUNs	5-12
5.3.1	Array Configuration Summary	5-12
5.3.2	LUN Parameters	5-15
5.3.2.1	RAID Level	5-16
5.3.2.2	Drive Map	5-17
5.3.2.3	Logical Block Size	5-18
5.3.2.4	Logical Unit Size	5-18
5.3.2.5	Segment Size	5-18
5.3.2.6	Segment Zero Size	5-19
5.3.2.7	Delay Interval	5-19
5.3.2.8	Blocks Per Delay Interval	5-19
5.3.3	Displaying and Modifying Logical Unit Parameters	5-20
5.3.4	Modify/Display Procedure	5-20
5.3.5	Creating a LUN from Spare Drives	5-22
5.3.6	Creating a LUN from an Existing Drive Group	5-25
5.3.7	Changing RAID Level/Drives on an Existing LUN	5-27

6 Troubleshooting

6.1	Before You Begin Troubleshooting	6-1
6.2	Using the Troubleshooting Table	6-1
6.3	If You Have Expanded Your StorageWorks RAID Array 110 Subsystem . .	6-3

7 Error Information and Symptoms Tables

7.1	System Message Tables	7-1
7.2	Symptoms and Status Tables	7-2
7.3	Error Reporting under OpenVMS VAX	7-3
7.3.1	Errors Reported by the StorageWorks RAID Array 110 Utility for OpenVMS VAX	7-3
7.3.2	Errors Reported by OpenVMS VAX	7-7
7.3.3	Request Sense Data Format	7-9

8 Removal and Replacement

8.1	How to Replace an HSZ10-AA Controller	8-1
8.2	How to Add an Additional Controller	8-7
8.3	When and How to Replace a Drive	8-8
8.4	When and How to Replace Power Supplies and Blowers	8-9
8.4.1	Replacing a Shelf Power Supply	8-9
8.4.2	Replacing a Blower	8-10
8.5	Deleting a LUN	8-11
8.6	Modifying and Displaying a Drive Status	8-14
8.6.1	Adding a Drive	8-14
8.6.2	Deleting a Drive	8-15
8.6.3	Failing a Drive	8-15
8.6.4	Reconstructing a Drive	8-16
8.7	Restoring a LUN	8-19
8.7.1	Restoring a LUN Task Summary	8-19
8.7.2	When to Replace a Drive	8-21
8.7.3	Drive Failures on RAID 0	8-23
8.7.3.1	Single Drive Failure on RAID 0	8-23
8.7.4	Drive Failures on RAID 1	8-24
8.7.4.1	Single Drive Failures on RAID 1	8-24
8.7.4.2	Multiple Drive Failures on RAID 1	8-24
8.7.5	Drive Failures on RAID 5	8-25
8.7.5.1	Single Drive Failure on RAID 5	8-25
8.7.5.2	Multiple Drive Failures on RAID 5	8-25
8.7.6	What is Reconstruction?	8-25
8.7.7	Reconstruction Rate	8-26
8.7.8	Restoring a RAID 0 LUN	8-27
8.7.9	Restoring a RAID 1 LUN	8-29
8.7.10	Restoring RAID 1 LUNs after Multiple Failures	8-31
8.7.11	Restoring a RAID 5 LUN	8-33
8.7.12	Restoring RAID 5 LUNs after Multiple Failures	8-35

9 Adjustments and Alignments

10 Configurations and Rules

10.1	Modifying Basic Configurations	10-1
10.2	Multiple Rank Configurations	10-2
10.3	Configuration Guidelines	10-2
10.3.1	Expansion Guidelines	10-3
10.4	Recommended Configurations	10-4
10.4.1	Expansion from a Single BA350-EA to a Single BA350-EA/Dual BA350-SA Configuration	10-5
10.4.1.1	Prereconfiguration	10-5
10.4.1.2	Shelf Reconfiguration	10-6
10.4.1.3	Postreconfiguration	10-10
10.4.2	Expansion from a Single BA350-EA to a Single BA350-EA/Quad BA350-SA Configuration	10-11
10.4.2.1	Prereconfiguration	10-11
10.4.2.2	Shelf Reconfiguration	10-12
10.4.2.3	Postreconfiguration	10-16
10.4.3	Expansion from a Single BA350-EA/Dual BA350-SA to a Single BA350-EA/Quad BA350-SA Configuration	10-17
10.4.3.1	Prereconfiguration	10-17
10.4.3.2	Shelf Reconfiguration	10-18
10.4.3.3	Postreconfiguration	10-22
10.5	Custom Expansions	10-22

11 Parts Information

11.1	Hardware	11-1
11.2	Firmware	11-6
11.3	Software	11-6
11.3.1	Hardware Requirements	11-6
11.3.2	Software Requirements	11-7
11.3.2.1	OpenVMS VAX V5.5-2	11-7
11.3.2.2	MS-DOS	11-7
11.3.2.3	NetWare	11-7
11.3.2.4	SCO UNIX	11-7
11.3.3	StorageWorks RAID Array 110 Utility Kits	11-8
11.3.3.1	StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit	11-8
11.3.3.2	StorageWorks RAID Array 110 Utility for MS-DOS Kit	11-8
11.3.3.3	StorageWorks RAID Array 110 Utility for Novell NetWare Kit ...	11-9
11.3.3.4	StorageWorks RAID Array 110 Utility for SCO UNIX Kit	11-10

12 Preventive and Proactive Maintenance

12.1	Upgrading Software	12-1
12.2	Checking/Repairing Array Parity for Novell NetWare	12-1
12.2.1	What is Parity Check/Repair?	12-2
12.2.1.1	When to Run Parity Check/Repair	12-2
12.2.2	How Automatic Parity Check/Repair Works	12-2
12.2.2.1	Automatic Parity Check/Repair Operation	12-2
12.2.3	How to Run Manual Parity Check/Repair	12-4
12.2.3.1	Manual Parity Check/Repair Operation	12-4
12.2.3.2	Running the Check	12-4

12.2.4	Changing Parity Check Scheduling	12-6
12.3	The Array Monitor Daemon for Novell NetWare	12-7
12.3.1	Array Monitor Daemon	12-7
12.3.1.1	How the AMD Works	12-7
12.3.1.2	What to do When an Error is Displayed	12-7
12.3.1.3	Error Log Example	12-8
12.3.2	Error Message Types	12-10
12.3.3	AMD Error Message Format	12-10
12.3.3.1	AMD Pop-Up Message Format	12-10
12.3.3.2	AMD Console Message Format	12-11
12.3.3.3	AMD Error Log Message Format	12-11
12.3.4	AMD Error Messages	12-12
12.3.4.1	Array Component Errors	12-12
12.3.4.2	Array Device Errors	12-14
12.3.4.3	Other Status Change Errors	12-16
12.3.4.4	RAID Manager Messages	12-16
12.4	Array Parity Check/Repair for SCO UNIX	12-19
12.4.1	Getting Started	12-19
12.4.1.1	Using the Shell Script	12-19
12.4.1.2	Invoking the Executable	12-19
12.4.2	Options	12-20
12.4.3	Output	12-21
12.4.4	Error Messages	12-21
12.5	RAID Monitor Utility for SCO UNIX	12-23
12.5.1	Getting Started	12-23
12.5.1.1	Using the Shell Script	12-23
12.5.1.2	Invoking the Executable	12-23
12.5.2	Options	12-24
12.5.3	Error Messages	12-24
12.6	RAID Status Utility for SCO UNIX	12-26
12.6.1	Getting Started	12-26
12.6.1.1	Using the Shell Script	12-26
12.6.1.2	Invoking the Executable	12-26
12.6.2	Options	12-27
12.6.3	Output	12-27
12.6.4	Error Messages	12-28

A SCSI-2 Error Codes

A.1	SCSI-2 Status Codes	A-1
A.2	SCSI Sense Keys and Additional Sense Codes	A-1

B KZESA Host Adapter Error Codes

C Total Call Concept (TCC)

D Supported Options

E Host SCSI Cables

Glossary

Index

Examples

7-1	Error Log Example	7-7
7-2	Request Sense Data Format	7-10

Figures

1-1	StorageWorks RAID Array 110 Subsystem	1-2
1-2	Logical View of the StorageWorks RAID Array 110 Subsystem	1-3
1-3	Subsystem Base Unit	1-7
1-4	Physical Layout of the HSZ10-AA Controller	1-8
2-1	Conceptual Diagram of the StorageWorks RAID Array 110 Subsystem	2-2
2-2	Diagram of RAID 0	2-3
2-3	Diagram of RAID 1	2-4
2-4	Diagram of RAID 3	2-5
2-5	Diagram of RAID 5	2-6
2-6	Diagram of a Regular LUN	2-7
2-7	Drive Partition Diagram	2-8
2-8	StorageWorks RAID Array 110 Subsystem SCSI IDs	2-14
2-9	Physical Drive Map	2-15
3-1	Typical Y Cable Connection	3-3
3-2	Connecting the HSZ10-AA Controller to an 8-Bit, Single-Ended Host or Adapter	3-4
3-3	Trilink Connector Midbus Connection	3-5
3-4	Trilink Connector End-Bus Connection	3-6
3-5	LED Indicators	3-7
4-1	Annotated ACE Window	4-4
4-2	ACE Window	4-6
4-3	Example of a HELP Screen	4-7
4-4	HELP Screen for Modifying and Displaying Parameters	4-7
4-5	Select LUN to be Deleted	4-8
4-6	Deletion Warning Message	4-8
4-7	After the Deletion of a LUN Screen	4-9
4-8	Exiting Without LUN 0 Warning	4-9
4-9	Predefined Configurations	4-10
4-10	Create LUN Screen	4-10
4-11	Modify/Display Parameters Screen	4-11
4-12	LUN Creation Completed	4-11

4-13	Disk Array Services Window	4-15
4-14	Array Controller Configuration Window	4-17
4-15	ACE Window	4-25
5-1	LUN State Diagram	5-2
5-2	LED Indicators	5-5
5-3	Power Supply Status LEDs	5-7
5-4	Shelf Status LEDs	5-9
7-1	Drive Groups	7-14
8-1	Removing the Front Bezel	8-2
8-2	HSZ10-AA in a BA350-EA Shelf	8-2
8-3	Removing the HSZ10-AA Controller	8-3
8-4	SW1—SCSI ID Settings	8-4
8-5	SCSI ID Setting Examples	8-5
8-6	LED Indicators	8-6
8-7	LED Indicators for the Fully Functional State	8-6
8-8	Diagram of the BA350-EA Shelf	8-7
8-9	Replacing a Drive	8-8
8-10	Replacing BA35X-MA Blowers	8-11
10-1	Single BA350-EA Shelf	10-6
10-2	Single BA350-EA to Single BA350-EA/Dual BA350-SA Shelf Reconfiguration	10-7
10-3	Single BA350-EA to Single BA350-EA/Dual BA350-SA Drive Reconfiguration	10-8
10-4	BA350-EA and BA350-SA Combination Storage Array	10-9
10-5	Single BA350-EA Shelf	10-12
10-6	Single BA350-EA to Single BA350-EA/Quad BA350-SA Shelf Reconfiguration	10-13
10-7	Single BA350-EA to Single BA350-EA/Quad BA350-SA Drive Reconfiguration	10-14
10-8	Single BA350-EA Shelf	10-18
10-9	Single BA350-EA/Dual BA350-SA Shelf to Single BA350-EA/Quad BA350-SA Shelf Reconfiguration	10-19
10-10	Single BA350-EA/Dual BA350-SA Shelf Configuration to Single BA350-EA/Quad BA350-SA Drive Reconfiguration	10-20

Tables

2-1	StorageWorks RAID Array 110 Subsystem General Specifications	2-10
2-2	StorageWorks Power Units	2-10
2-3	Thermal Stabilization Specifications	2-11
2-4	Environmental Specifications	2-12
3-1	Midbus Connections	3-4
3-2	End-Bus Connections	3-5
3-3	Files Copied to the File Server in NetWare	3-9
3-4	Installation Error Messages	3-13
3-5	StorageWorks RAID Array 110 Utility for MS-DOS Kit	3-15
3-6	StorageWorks RAID Array 110 Utility for SCO UNIX Kit	3-16
3-7	TCQ and SCSI Port Driver Support for VAX Systems	3-19

4-1	Moving within the RAID Manager Menus	4-14
4-2	LUN Status from the Disk Array Devices Window	4-20
4-3	Drive Status	4-21
4-4	Understanding the ACE Window	4-26
4-5	ACE Navigation Keys	4-26
4-6	LUN Status from the ACE Window	4-30
5-1	LUN States	5-2
5-2	Drive Status	5-3
5-3	Summary of HSZ10-AA Controller LED Codes	5-6
5-4	Shelf and Single Power Supply Status LEDs	5-8
5-5	Shelf and Dual Power Supply Status LEDs	5-8
5-6	Drive SBB Status LEDs	5-10
5-7	StorageWorks RAID Array 110 Utility for MS-DOS LUN Configuration Tasks	5-13
5-8	StorageWorks RAID Array 110 Utility for Novell NetWare LUN Configuration Tasks	5-13
5-9	StorageWorks RAID Array 110 Utility for OpenVMS VAX LUN Configuration Tasks	5-15
5-10	Logical Unit Parameters	5-16
5-11	RAID Level/Drive Selection	5-17
6-1	Troubleshooting Systems Problems	6-2
7-1	System Messages and How to Find Them	7-1
7-2	Symptoms and Status Codes and How to Find Them	7-2
7-3	StorageWorks RAID Array 110 Utility for OpenVMS VAX Error Codes	7-4
8-1	Blower Replacement	8-11
8-2	Procedure for Restoring LUNs	8-20
8-3	LUN Status	8-22
8-4	Reconstruction Rates	8-26
10-1	Expansion Paths	10-4
11-1	Cabinets	11-1
11-2	BA350 Shelves	11-2
11-3	SBBs	11-2
11-4	Accessories	11-3
11-5	Cables	11-4
11-6	StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit	11-8
11-7	StorageWorks RAID Array 110 Utility for MS-DOS Kit	11-9
11-8	StorageWorks RAID Array 110 Utility for Novell NetWare Kit	11-9
11-9	StorageWorks RAID Array 110 Utility for SCO UNIX Kit	11-10
12-1	Array Component Errors	12-12
12-2	Array Device Errors	12-14
12-3	Other Status Change Errors	12-16
12-4	RAID Manager Messages	12-16
12-5	Parity Check/Repair Utility Options	12-20
12-6	RAID Monitor Options	12-24
12-7	RAID Status Utility Options	12-27
A-1	SCSI-2 Status Codes for the Array Controller	A-1

A-2	SCSI Sense Keys	A-2
A-3	SCSI Error Codes	A-3
B-1	SCSI Status Codes for the KZESA Host Adapter	B-1
D-1	Disk Drives Supported	D-1
D-2	Adapters Supported	D-1
E-1	SCSI Cable Selections	E-1

Preface

Purpose

This guide introduces the StorageWorks RAID Array 110 Subsystem storage solution. It describes available RAID features, packaging options, and configuration and operating instructions. This guide serves as complimentary documentation to the StorageWorks® documentation set.

Important Note

The DEC RAID OpenVMS VAX Utility is renamed to the StorageWorks RAID Array 110 Utility for OpenVMS VAX. This guide refers to DEC RAID OpenVMS VAX Utility in code examples only. In all other cases, the new product name is used.

Intended Audience

This guide is intended for use by Digital Service Engineers responsible for installing and maintaining the StorageWorks RAID Array 110 Subsystem products.

It is presumed that readers of this guide have an understanding of SCSI technology. For readers installing and maintaining the StorageWorks RAID Array 110 utilities on systems with the OpenVMS VAX®, SCO® UNIX®, MS-DOS®, or NetWare® operating systems, familiarity with the specific operating system used is required.

Structure

This guide is organized as follows:

Chapter 1, Introduction – provides an overview of the StorageWorks RAID Array 110 Subsystem features and options and describes the StorageWorks RAID Array 110 Subsystem components including the following:

- Base unit
- Expansion units
- Utilities
- Driver support
- Interconnect
- Host adapters

Chapter 2, Technical Information – provides the following information:

- Presents background information on disk array levels and components

- Describes the physical, environmental, and performance specifications for the StorageWorks RAID Array 110 Subsystem product
- Describes the subsystem and its physical/logical mappings

Chapter 3, Installation – contains the following information:

- Provides installation instructions for the HSZ10-AA controller, host connections, and functional verification
- Describes the installation of the StorageWorks RAID Array 110 for Novell NetWare software
- Describes the installation of the StorageWorks RAID Array 110 Utility for SCO UNIX software

Chapter 4, Diagnostics and Software – describes the StorageWorks RAID Array 110 utilities and the StorageWorks RAID Array 110 Utility for OpenVMS VAX and discusses when and how to run a given utility. It also describes the RAID Manager™ menus.

Chapter 5, Operations – provides the following information:

- Describes modes of operation and methods of monitoring the StorageWorks RAID Array 110 Subsystem
- Describes how to add logical units (LUNs) for the StorageWorks RAID Array 110 Utility for Novell NetWare
- Describes how to initially set up your StorageWorks RAID Array 110 Subsystem
- Describes how to configure logical units (LUNs)

Chapter 6, Troubleshooting – provides error handling and troubleshooting guidelines.

Chapter 7, Error Information and Symptoms Tables – provides lists of system and error messages, symptom and status code tables, and error reporting under the OpenVMS VAX operating system.

Chapter 8, Removal and Replacement – contains the following information:

- Describes how to remove, replace, and add an HSZ10-AA controller
- Describes when and how to replace a drive, power supplies, and blowers
- Describes reconstruction and Parity Check/Repair, upgrading software
- Describes deleting, displaying, and modifying logical units
- Describes restoring logical units

Chapter 9, Adjustments and Alignments – informs the reader that there are no adjustments or alignments for this product.

Chapter 10, Configurations and Rules – discusses advanced configurations and expansion of the base unit.

Chapter 11, Parts Information – contains lists of part numbers and descriptions for the StorageWorks RAID Array 110 Subsystem and the StorageWorks RAID Array 110 utilities.

Chapter 12, Preventive and Proactive Maintenance – contains the following information:

- Describes Parity Check/Repair for the StorageWorks RAID Array 110 for Novell NetWare
- Describes the automatic and manual check and repair procedure for NetWare
- Describes Parity Check/Repair for the StorageWorks RAID Array 110 Utility for SCO UNIX
- Describes the RAID Monitor for the StorageWorks RAID Array 110 Utility for SCO UNIX and how it reports errors
- Describes the RAID Status Utility for the StorageWorks RAID Array 110 Utility for SCO UNIX and how it displays current status of logical units and drives

Appendix A, SCSI-2 Error Codes – describes SCSI status codes, SCSI sense keys, and additional sense codes.

Appendix B, KZESA Host Adapter Error Codes – lists the error codes for the KZESA Host Adapter.

Appendix C, Total Call Concept (TCC) – informs the reader that the *Total Call Concept* does not apply to this product.

Appendix D, Supported Options – includes information regarding the disk drives and adapters supported by the StorageWorks RAID Array 110 Subsystem.

Appendix E, Host SCSI Cables – lists the SCSI cables used with specific host computer systems and the StorageWorks shelves.

Glossary – provides an alphabetical listing of key terms.

Related Documents

The following is a list of documents that contain information related to this product:

Document Title	Order Number
<i>StorageWorks RAID Array 110 Subsystem Pocket Service Guide</i>	EK-SZ200-PS
<i>StorageWorks RAID Array 110 Subsystem User's Guide</i>	EK-SM2CA-UG
<i>DEC RAID Subsystem User's Guide</i>	EK-SZ200-UG
<i>HSZ10-AA Controller Site Preparation Guide</i>	EK-HSZ10-IN
<i>DEC RAID Utilities User's Guide</i>	EK-DECRA-UG
<i>DEC RAID OpenVMS VAX Utility User's Guide</i>	AA-PYZLA-TE

Document Title	Order Number
<i>DEC RAID OpenVMS VAX Utility Release Notes and Installation Guide</i>	AA-PYZUA-TE
<i>DEC SCSI Tagged Command Queuing (TCQ) Driver for OpenVMS VAX Release Notes and Installation Guide</i>	AA-PXKGA-TE
<i>StorageWorks RAID Array 110 Utility for Novell NetWare User's Guide</i>	AA-Q0N4A-TE
<i>StorageWorks RAID Array 110 Utility for SCO UNIX User's Guide</i>	AA-Q0N6A-TE
<i>StorageWorks RAID Array 110 Utility for MS-DOS User's Guide</i>	AA-Q0N5A-TE
<i>BA35X Vertical Mounting Kit User's Guide</i>	EK-BA350SV-UG
<i>StorageWorks Configuration Guide</i>	EK-BA350-CG
<i>BA350-EA Modular Storage Shelf User's Guide</i>	EK-350EA-UG
<i>StorageWorks Family User's Guide</i>	EK-350SA-UG

Documentation Conventions

The following conventions are used in this manual:

boldface type	Boldface type indicates the first instance of terms in the text that are defined in the glossary.
<i>italic type</i>	Italic type indicates emphasis and complete manual titles.
Note	A note calls the reader's attention to any item of information that may be of special importance.
Caution	A caution contains information essential to avoid damage to the system.
Warning	A warning contains information essential to the safety of personnel.

Introduction

This chapter contains the following information:

- Product description
- StorageWorks RAID Array 110 Subsystem components descriptions

1.1 Product Description

This section contains the following information:

- Overview
- Product highlights
- Product attributes
- General conclusions


1.1.1 Overview

The StorageWorks RAID Array 110 Subsystem is a modular, integrated, end-user **RAID** (Redundant Array of Independent Disks) solution based on the HSZ10-AA controller (disk array controller). Figure 1-1 shows the StorageWorks RAID Array 110 Subsystem.

Introduction

1.1 Product Description

Figure 1-1 StorageWorks RAID Array 110 Subsystem



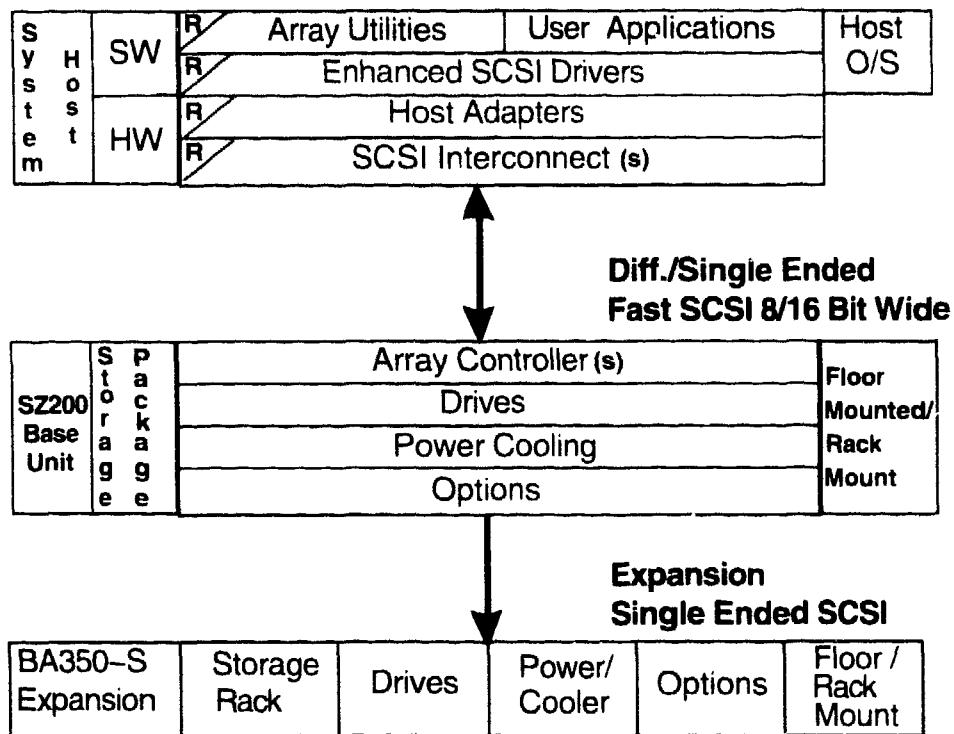
SHR-XR3000-VIF

The StorageWorks RAID Array 110 Subsystem provides you with a highly available and integrated storage solution through the use of RAID technology, optional redundancy of key components, value-added software utilities, and SCSI interconnect components. Flexibility and performance result from your ability to select a RAID level that best meets your application needs.

Logically, the subsystem looks like a large disk drive or multiple disk drives. All RAID functionality is provided by a controller and associated firmware and software. RAID management utilities are provided to allow you to configure and control the subsystem as illustrated in Figure 1-2.

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

Figure 1-2 Logical View of the StorageWorks RAID Array 110 Subsystem



SHR XR3002 GRA

Physically, the StorageWorks RAID Array 110 Subsystem is based on the StorageWorks family of products. The StorageWorks modular products provide configuration flexibility by offering various shelf options, cabinet kits, and system building blocks (SBBs). In addition, features such as **hot swapping** (meaning that the system does not have to be powered down to perform swapping) of disks and power options are achieved by means of carriers that interface directly with a rigid backplane.

This interface, along with optional redundant components, allows the removal and replacement of key subsystem components without having to render the system inoperable—a key attribute for systems that cannot tolerate down time.

1.1.2 Product Highlights

This section includes the following information:

- Array features
- Subsystem features

1.1.2.1 Array Controller Features

The following is a list of array features:

- RAID 0, 1, 0+1, 3, and 5
- Five drive channels
- Host array interface: differential 8- and 16-bit, 10/20 Mbytes per second, SCSI-2
- Array channel interface: single-ended 8-bit, 5 Mbytes per second, SCSI-2
- Array parity generation/checking and recovery with hardware assist

- Array parity/data reconstruction with hardware assist

1.1.2.2 Subsystem Features

The following is a list of subsystem features:

- Hot swapping drives and controller
- Drive fault light indicators
- Power supply and blower monitoring and indicators
- Support for up to 35 disk drives
- Redundant components
- Modular configuration and upgrade paths

1.1.3 Product Attributes

This section includes information on the following attributes:

- Data reliability
- Redundancy
- Data availability
- Performance
- Flexibility
- Capacity

1.1.3.1 Data Reliability

Unlike a single SCSI disk drive solution, the StorageWorks RAID Array 110 Subsystem provides dependable data reliability in the event of a disk failure. With a SCSI disk drive solution, the user must replace the product and reconstruct the data through backup media. This mechanism results in some user data loss, in the event of a failure, since backup of data is typically done at scheduled or infrequent intervals. Any activity that has occurred since that last user data backup is not saved.

In a StorageWorks RAID Array 110 Subsystem solution, there is no loss of data since redundant elements provide the capability of continuing I/O activity despite the loss of a drive and reconstruction of data back on the replaced drive. This means that data is updated up to the second of failure and throughout the reconstruction process.

1.1.3.2 Redundancy

Most RAID products provide the feature of disk drive redundancy. However, what is often neglected in a RAID solution is the redundancy of other key components.

This StorageWorks RAID Array 110 Subsystem provides optional redundancy of the following components:

- Power supplies
- RAID disk array controllers (HSZ10-AA controllers)
- Blowers

All components listed are removable and replaceable without needing to power down the subsystem. These features contribute to making the subsystem fully redundant.

1.1.3.3 Data Availability

One of the key attributes of the StorageWorks RAID Array 110 Subsystem is data availability. The different RAID levels offer data protection at the drive level. In addition, the redundancy designed into the StorageWorks packaging allows you to access data even through other component failures.

For users who depend upon minimal down time, these features are critical.

1.1.3.4 Performance

The performance of a disk array depends on the environment and the I/O workload. The following are general comments regarding array performance:

- An array of smaller capacity drives can perform better than a single, large-capacity drive. This is due to the availability of multiple actuators that allow execution of I/Os simultaneously.
- RAID 0 performs slightly better than a group of drives since it tends to achieve load balancing across the drives.
- The RAID 1 read performance is better than a group of drives since requests can be served by either drive in a **mirrored** pair of disks. However, the RAID 1 write performance is slightly less due to the need to write to both drives.
- RAID 0+1 is the combination of striping and mirroring, implemented by striping mirrored sets. RAID 0+1 offers the best I/O performance of any type of RAID by combining the performance advantages of RAID 0 and RAID 1. RAID 0+1 also provides data redundancy. Performance is equivalent to RAID 0 and the cost is the same as RAID 1.
- RAID 3 is suitable for environments with large, sequential transfers. It performs significantly better (four times better) than all the other alternatives.
- Like RAID 1, RAID 5 performs slightly better than a group of drives with read requests. However, write performance is slightly impacted due to the need to write parity information.

A description of these RAID levels is contained in Chapter 2.

1.1.3.5 Flexibility and Capacity

The disk array controller provides more capacity on a single-host adapter than other data storage alternatives. The array controller acts as a multiplexer and allows more SCSI disk drives to be connected to a single-host connection.

The disk array controller has the flexibility to support different RAID levels and classes of drives within the same array. This flexibility is important to meet the various price, performance, and capacity options needed for different business environments.

The combination of a modular packaging strategy along with an assortment of cabinet and mounting options provides for a flexible, highly available storage solution.

1.1.4 General Conclusions

The StorageWorks RAID Array 110 Subsystem provides a set of options from which users can choose. The options vary with the following features:

- Data availability
- Performance
- Capacity

The flexibility in choosing various options to satisfy these requirements allows you to meet your cost goals.

If your business cannot tolerate data loss or down time due to drive failures, then you need the protection that disk arrays (RAID 0+1, 1, 3, or 5) can provide.

Unlike multiple disks or RAID 0, the chance of losing data in a RAID 0+1, 1, 3, or 5 configuration is extremely small even with a large number of drives.

In many businesses, the initial price of disk arrays versus multiple disks is offset by the cost that would be incurred by a single-drive failure. This failure usually results in the loss of data and a disruption of business operations.

For the same supported capacity, RAID 3 and 5 provide reliable storage at a fraction of the drive price of RAID 1.

Disk arrays are the most cost-effective choice to meet the requirements of data reliability, availability, performance, and capacity demanded in most business environments.

1.2 StorageWorks RAID Array 110 Subsystem Component Descriptions

This section contains a description of the following StorageWorks RAID Array 110 Subsystem components:

- Subsystem base unit
- System building blocks
- StorageWorks RAID Array 110 utilities
- Enhanced SCSI driver support
- SCSI interconnects/host adapters

1.2.1 Description of the Subsystem

The subsystem base unit is the basic building block for the StorageWorks RAID Array 110 Subsystem. The subsystem base unit includes the following components:


- BA35X-VA vertical mounting kit
- BA350-EA shelf
- HSZ10-AA controller (disk array controller)
- System building blocks (SBBs)
 - Power
 - Disk drives
 - Adapters

1.2 StorageWorks RAID Array 110 Subsystem Component Descriptions

For StorageWorks RAID Array 110 Subsystem physical specifications, refer to Section 2.2. Figure 1-3 shows the subsystem base unit and identifies its components.

Figure 1-3 Subsystem Base Unit

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



SHR-XR3001-VIF

1.2.1.1 BA35X-VA Vertical Mounting Kit

The BA35X-VA vertical mounting kit is part of the StorageWorks family of cabinet options. For the subsystem base product, two BA35X-VA vertical mounting kits are used to house the BA350-EA shelf. Each BA35X-VA kit has an ac power controller that provides switch-controlled input voltages to the shelf power supplies. Refer to the *BA35X-VA Vertical Mounting Kit User's Guide* for more detailed information.

1.2.1.2 BA350-EA Shelf

The BA350-EA shelf is a double-width StorageWorks storage shelf option. The BA350-EA shelf houses the HSZ10-AA controller and SBBs (power, drive, and adapter options). All components are plugged into a common backplane.

The shelf provides six single-ended SCSI connectors. Five connectors are used for RAID storage expansion from the HSZ10-AA controller to the BA350-SA shelves. The sixth connector can be used for host interconnect adapters or storage devices.


1.2.1.3 HSZ10-AA Controller (Disk Array Controller)

The HSZ10-AA controller is an intelligent SCSI disk array controller based on the Motorola® MC68EC020 microprocessor running at 25 MHz. The HSZ10-AA controller supports RAID 0, 0+1, 1, 3, and 5. It supports expansion of up to 35 devices (seven ranks).

The host SCSI channel is a synchronous/asynchronous/fast, 8/16-bit differential. It supports SCSI-2 protocol. The host connection is made through a 68-pin, high-density connector (J1) as shown in Figure 1-4. A second connector (J2) is provided for daisy-chaining. SCSI term power is provided.

There are five array SCSI channels or interfaces that are made through connectors J3, J4, and J5 as shown in Figure 1-4. Each interface is an 8-bit, single-ended synchronous/asynchronous SCSI-2. Active SCSI termination and SCSI termination power are provided on the controller.

Figure 1-4 Physical Layout of the HSZ10-AA Controller



Legend			
J1	68 pin host connector	J7	drive fault signals (reserved)
J2	68 pin host connector	J10	RS232 port (reserved)
J3	2 array SCSI channels	J11	diagnostic port (reserved)
J4	2 array SCSI channels	J17	processor test port (reserved)
J5	1 array SCSI channel	J18	processor test port (reserved)
J6	subsystem signals		
	SW1	SCSI switches	
	SW2	manufacturing switches (reserved)	

SHR-XR3003-GRA

1.2.1.4 Features

This section discusses the following features:

- Controller features
- Array features

HSZ10-AA Controller Features

- 68EC020 microprocessor running at 25 MHz
- DRAM, 1 Mbyte for program and structures
- EPROM, 256 Kbytes (expandable to 512 Kbytes) for diagnostics and start of day code
- EEPROM, 8 Kbytes for configuration and logging
- 68901 multifunction peripheral controller for timers, serial communications, and interrupt control

1.2 StorageWorks RAID Array 110 Subsystem Component Descriptions

- Host SCSI target ID switch (4-bit)
- Status LEDs (8-bit)
- Self-test capability

Array Controller Features

- SCSI host interface, 8/16-bit, synchronous, differential SCSI-2, 10/20 Mbytes per second
- SCSI array channels, 8-bit, synchronous, single-ended SCSI-2, 5 Mbytes per second, active termination
- Hot swapping controller and drive interfaces
- RAID 0, 0+1, 1, 3, and 5
- Hardware-assisted array parity generation/checking, recovery and reconstruction
- 64 Kbyte read/modify/write buffer
- Serial drive fault interface on each array channel
- Blower and power supply monitoring

1.2.1.5 Firmware

The HSZ10-AA controller firmware is the part of the code that resides with the controller's EPROM. This code is responsible for controller diagnostics, initialization, response to a limited set of SCSI commands, controller software downloading/uploading, and passing control to the HSZ10-AA controller software.

1.2.1.6 Software

The HSZ10-AA controller software is the part of the code that is stored on the array drives and must be uploaded into the controller's DRAM prior to operation. This code is responsible for all non-ROM operation of the board, including all RAID algorithms and configurations, read and write operations, and advanced SCSI functionality.

The software attempts to tolerate failures in the array that would cause loss of data access. By maintaining multiple copies of the array software on multiple drives, the controller software has as high availability as the user data. The software also maintains configuration information in both EEPROM and on the disk in order to tolerate the loss of one of these components.

The software is factory installed on the drives. Future upgrades to the software are accomplished by a downline load across the SCSI bus from the host that is saved on the drives. A utility is provided to upgrade the software, which is described in Section 4.3.7.2.

1.2.2 System Building Blocks (SBBs)

All power, storage, and adapter options are mounted inside 3-1/2 inch modular carriers that plug into slots in the BA350-EA shelf. A key feature is that each SBB can be removed and replaced without removing of power to the system. In addition, each SBB has visual status indicators so that it is easy to determine whether devices are functioning properly.

Note

The SBB is a field replaceable unit (FRU) and should not be opened. No attempt should be made to replace the parts associated with an SBB.

Power SBBs

The subsystem base unit uses two BA35X-HA universal ac input power supplies. Additional power supplies can be added for power redundancy.

For more information on the StorageWorks power options, refer to the *StorageWorks Configuration Guide*.

Disk SBBs

The subsystem base unit is configured with five 3-1/2 inch disk drives. For supported disk options, refer to Appendix D.

Adapter SBBs

The subsystem unit provides a slot for a bus adapter option that connects to different host types. For example, the DWZZA-VA bus adapter, which is a single-ended to differential SCSI converter, fits into the subsystem unit.

Refer to Appendix D for supported adapter options.

1.2.3 Expansion Unit Description

Since the StorageWorks family of products is based on a building block concept, expansion of the subsystem base unit can be customized to allow for expansion.

Expansion of the BA350-EA shelf is done by using the BA350-SA shelf. The BA350-SA shelf is a single-width storage module that holds up to seven 3-1/2 inch devices. The shelf can be configured as either one or two SCSI buses.

Depending on the needed expansion, additional cabinet options are available for installation of the shelf modules. For information on the available cabinet options, refer to the *StorageWorks Configuration Guide*.

1.2.4 StorageWorks RAID Array 110 Utilities

StorageWorks RAID Array 110 utilities are either applications or standalone or operating system-specific utilities that execute on the host system. These are used to configure, monitor, diagnose, and maintain the array subsystem.

Details on these utilities are operating system specific and are described in Chapter 4.

1.2.5 Enhanced SCSI Driver Support

The enhanced SCSI drivers are part of the host adapter support. Enhanced SCSI features are available in the SCSI drivers to better utilize the disk array. These features include the following:

- SCSI-2 tagged command queuing support
- SCSI-2 logical unit support
- Enhanced SCSI error handling capability

1.2 StorageWorks RAID Array 110 Subsystem Component Descriptions

1.2.6 SCSI Interconnects/Host Adapters

The subsystem base unit can connect to a variety of qualified host adapters through various SCSI interconnect schemes in order to meet the needs of SCSI users. Host adapters currently supported are listed in Appendix D. Connecting to a host and terminating the SCSI bus are described in Chapter 3.

The purpose of the SCSI interconnect is to provide fast/wide differential (FWD) SCSI as a standard connection. This provides a direct interconnect into the HSZ10-AA controller and uses the standard SCSI-3 P cable interconnect.

Other cable options provide resolution for the following:

- Connections to single-ended devices
- Connections to 8-bit devices initiator/targets with correct bus termination

Technical Information

This chapter contains information on the following:

- RAID overview
- Physical specifications
- StorageWorks RAID Array 110 Subsystem Base Configuration
- Physical drive map
- Logical drive map
- Key concepts

2.1 RAID Overview

This section contains information on RAID (Redundant Array of Independent Disks) levels, including the following:

- What is an array?
- Description of RAID 0
- Description of RAID 1
- Description of RAID 3
- Description of RAID 5

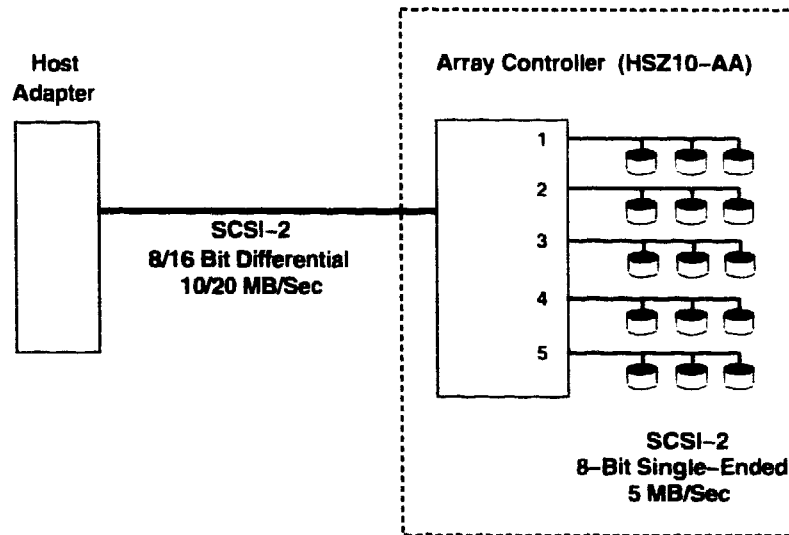
2.1.1 What Is a Disk Array?

This section gives a brief definition of disk arrays and a description of the four RAID levels currently supported by the StorageWorks RAID Array 110 Subsystem.

An **array** is a set of multiple disk drives and a specialized controller (an array controller), which keeps track of how the data is distributed across the disk drives.

Figure 2-1 shows a conceptual model of the StorageWorks RAID Array 110 Subsystem.

Figure 2-1 Conceptual Diagram of the StorageWorks RAID Array 110 Subsystem



SHR-XR3004-GRA

Data for a given file is segmented by dividing and storing the data to the different drives in the array, rather than writing the data to a single drive. A **segment** is a group of continuous data blocks that can be stored on a disk drive. By using multiple drives, the array can provide higher data transfer rates and higher I/O rates when compared to a single large drive.

Arrays can also provide data redundancy, so that no data is lost if a single drive in the array fails. Depending on the RAID level, data is either **mirrored** or **striped**. The RAID levels currently supported are RAID 0, 0+1, 1, 3, and 5. RAID 0+1, 1, 3, and 5 offer data redundancy.

Arrays are contained within an array subsystem or within an integrated array. Depending on the model and on how you configure it, an array subsystem can contain one or more arrays (referred to as **logical units numbers (LUNs)** (sometime referred to as logical units) – see the following sections for a description of arrays and LUNs). Each LUN has its own characteristics (RAID level, logical block size, logical unit size, and so on).

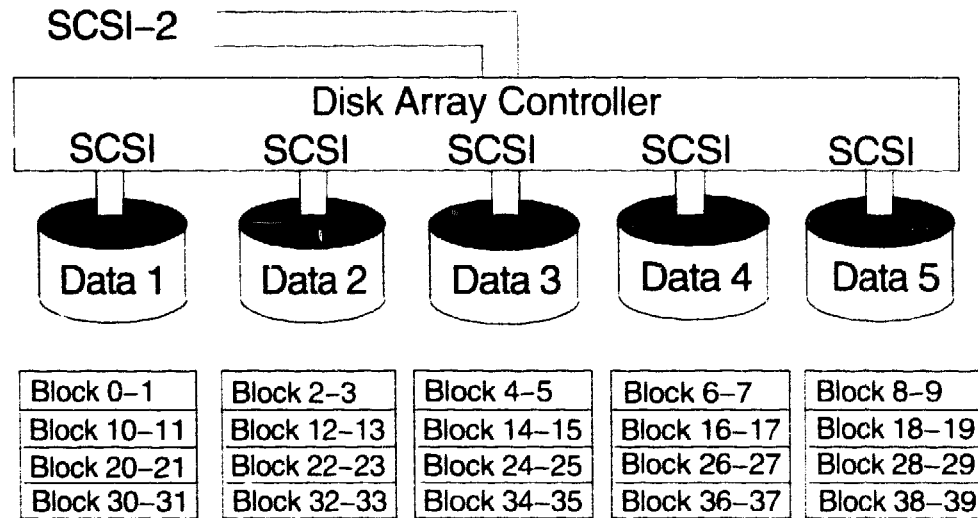
2.1.1.1 Description of RAID 0

RAID 0 stores data across the drives in the array one segment at a time, as shown in Figure 2-2.

In Figure 2-2, a segment is defined as two blocks of 512 bytes each. Blocks 0 and 1 are written to drive 1, blocks 2 and 3 are written to drive 2, and so on until each drive contains a single segment. Then blocks 10 and 11 are written to drive 1, blocks 12 and 13 to drive 2, and so on. (Note that segment zero is 0 blocks in this figure, and therefore is not shown.)

The host system treats a RAID 0 array like a standard hard drive. RAID 0 errors are reported in the same way as normal drive errors, and the recovery procedures are the same as those used on a standard hard drive. For example, if a drive

Figure 2-2 Diagram of RAID 0



SHR-XR3005-GRA

fails, the array controller returns the same errors that occur during a read or write retry failure.

RAID 0 offers a high I/O rate, but is a nonredundant configuration, meaning that there is no array parity information generated for reconstructing data in the event of a drive failure. Therefore, there is no error recovery over and above what is normally provided on a single drive. All data in the array must be backed up regularly to protect against data loss.

2.1.1.2 Description of RAID 1

RAID 1 transparently mirrors data by striping segments across data drives and mirrored data drives. Any time data is written to a drive, it is also written to a mirrored drive without the system directly controlling the location of the mirrored data as shown in Figure 2-3.

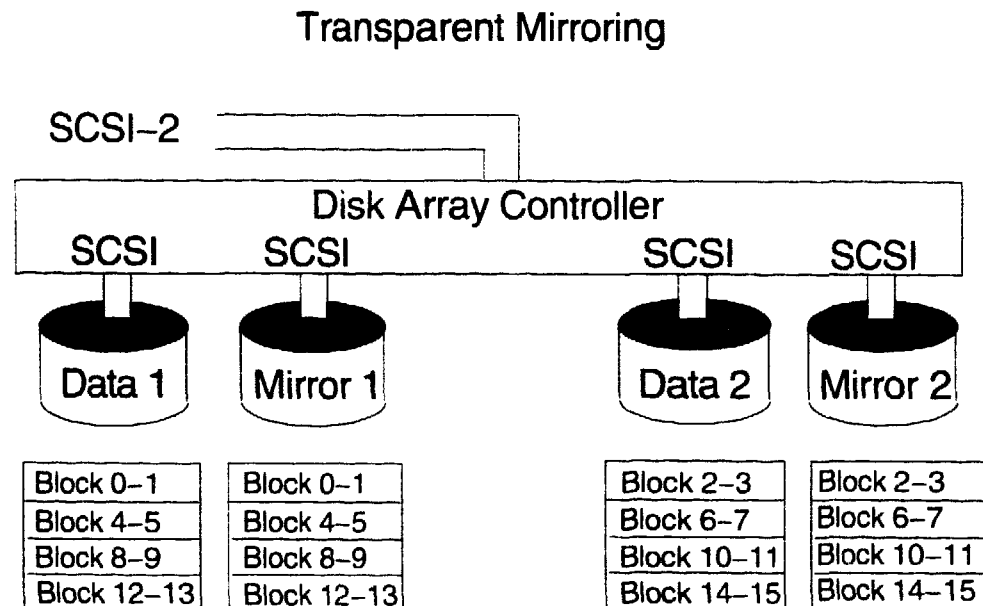
Traditionally, RAID 1 has been used for critical fault-tolerant transaction processing. There is high reliability provided with the mirrored data and there is high I/O rate (with small block size). However, RAID 1 is a more costly RAID solution because it requires a mirrored data drive for every data drive in the array.

In Figure 2-3, a segment is defined as two blocks of 512 bytes. Blocks 0 and 1 are written to data drive 1, while the mirrored data blocks 0 and 1 are written to the mirrored data drive 1. Then data blocks 2 and 3 are written to data drive 2, while the mirrored data blocks 2 and 3 are written to the mirrored data drive 2, and so on.

If a drive fails in RAID 1 array, you can continue to use the array normally since data from its mirrored drive is retrieved.

A RAID 1 array is unique in that you can have more than one failed drive and still use the data as long as there is only one failed drive in each mirrored pair. For example, if you have a six-drive RAID 1 LUN, you have three pairs. If there is a drive failure in each pair, you can still access your data. It is still

Figure 2-3 Diagram of RAID 1



SHR-XR3006-GRA

recommended that you replace the failed drives as soon as possible because if a second drive fails in a pair, all data on the LUN is lost.

2.1.1.3 Description of RAID 3

RAID 3 transfers data in parallel to data drives and to one parity drive. The result is an array that works best with large block transfers, but does not work well for transaction processing.

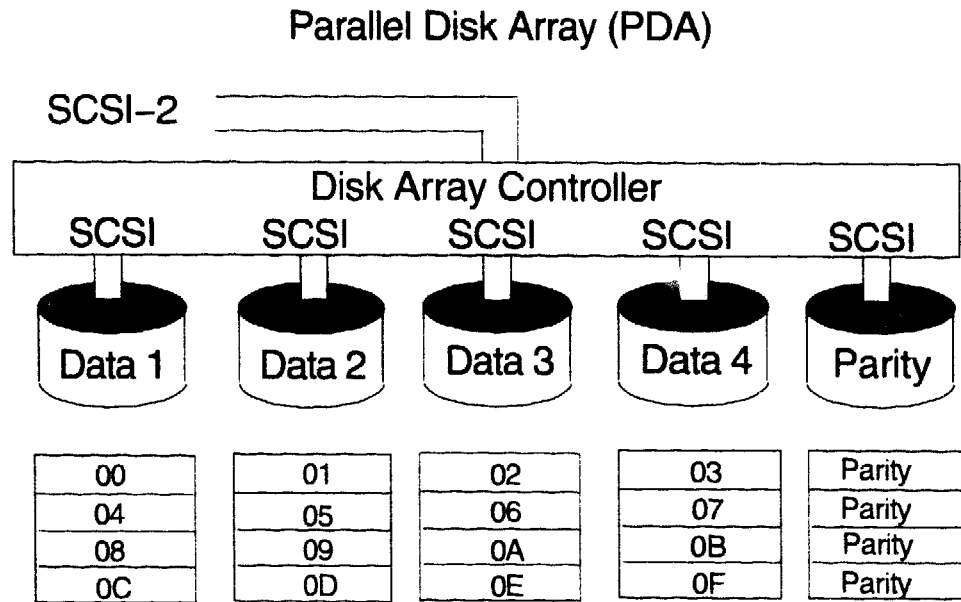
In Figure 2-4, the segment is defined as 1 byte of data. The host transfers the block of data, and the array controller stripes the data to the drives one byte at a time. Byte 00 is written to drive 1, byte 01 is written to drive 2, byte 02 is written to drive 3, byte 03 is written to drive 4, and the parity byte for bytes 00 through 03 is written to the parity drive.

Note

The logical block size for a RAID 3 LUN varies, and is determined by the operating system's I/O and the number of data drives in the LUN. If the I/O block is 512 bytes, then the array controller can process two times (for a three-drive LUN) or four times (for a five-drive LUN) the I/O size of 512 bytes at a time. Remember that a RAID 3 LUN has only two data drives in a three-drive LUN, and four data drives in a five-drive LUN. This means that the logical block size of a three-drive LUN is 1024 bytes, and 2048 bytes for a five-drive LUN.

If a drive fails in RAID 3 array, you can continue to use the array normally since the array controller recalculates the data and parity blocks from the data and parity on the operational drives.

Figure 2-4 Diagram of RAID 3



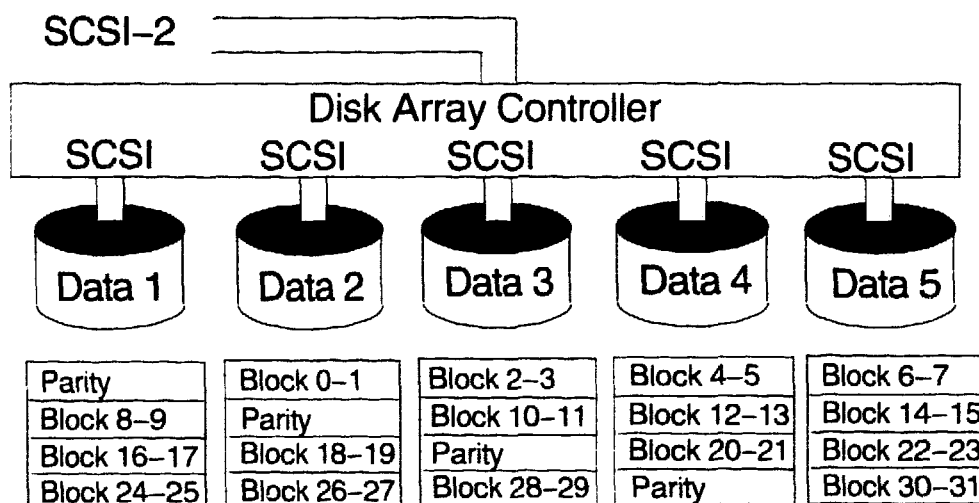
SHR-XR3007-GRA

For example, if drive 1 in Figure 2-4 fails, the array controller retrieves the data by using the parity information for that segment and the data from drives 2, 3, and 4 to reconstruct the data for drive 1. This process is repeated to reconstruct each segment on a failed drive as needed, so you can continue to operate the RAID 3 array.

2.1.1.4 Description of RAID 5

RAID 5 stores data across the drives in the array one segment at a time (a segment can contain multiple blocks). It also writes array parity data, but the parity data is spread across all the drives. The result is a transfer rate equal to that of a single drive but with a high overall I/O rate as shown in Figure 2-5.

Figure 2-5 Diagram of RAID 5



SHR-XR3008-GRA

In Figure 2-5, a segment is defined as two blocks of 512 bytes. Blocks 0 and 1 are written in the first position on drive 2, blocks 2 and 3 are written on drive 3, blocks 4 and 5 are written on drive 4, blocks 6 and 7 are written on drive 5, and the parity data for the blocks in data segments 1, 2, 3, and 4 is written on drive 1, and so on.

If a drive fails in the RAID 5 array, you can continue to use the array normally since the array controller recalculates the data on the failed drive using data and parity blocks on the operational drives.

For example, to recalculate data in data segment 4 in Figure 2-5 (the first position on drive 5), the array controller would use the parity information from drive 1 and the data from drives 2, 3, and 4 (data segments 1, 2, and 3) to reconstruct the data. This process is repeated to reconstruct each block of the failed drive as needed, so you can continue to operate the RAID 5 array.

2.1.2 Key Concepts

The following sections define some of the terminology used to describe arrays.


2.1.2.1 Array Channels

Array channels are the SCSI-2 compliant buses on which the disk drives are located. Each array channel is an independent SCSI bus.

2.1.2.2 Logical Units (LUNs)

A LUN is a grouping of drives. A LUN has its own device SCSI ID and LUN number. Each LUN has its own array parameters (RAID level, segment size, and so on). For most purposes, a LUN is equivalent to an array. See Figure 2-6.

Figure 2-6 Diagram of a Regular LUN



A LUN is treated as a single disk drive by most operating systems, so there are usually no special requirements. Depending on the operating system, you can create partitions or volumes in the same way you would treat an internal or external hard disk.

A LUN can be in various modes of operation, depending upon the RAID level configuration.

LUN modes of operation for nonredundant RAID 0:

- Optimal (All drives in the LUN are functional.)
- Dead (One or more drives in the LUN are not functional.)

LUN modes of operation for redundant RAID 0+1, 1, 3, and 5:

- Optimal (All drives in the LUN are functional.)
- Degraded (One drive in the LUN is not functional.)
- Dead (Two or more drives in the LUN are not functional.)

2.1.2.3 Drive Groups

A **drive group** is a set of from 1 to 10 drives that have been configured into one or more LUNs. A LUN can be contained in only one drive group, and all the LUNs in a drive group must have the same RAID level, drive type, and capacity. Operating systems treat the LUN (not the drive group) as a single disk drive, sending I/O to it and retrieving I/O from it.

Most of the time, you should create only one LUN in a drive group, so the use of the terms *logical unit* and *drive group* are synonymous.

Technical Information

2.1 RAID Overview

2.1.2.4 Drive Ranks

Drive ranks represent a numbering scheme providing information on the maximum number of drives on every array channel. A one-rank system indicates that there is a maximum of one drive per disk channel. A two-rank array indicates that there is a maximum of two drives per disk array channel. However, any array channel can have zero for its maximum number.

If you visualize the array channels as providing the user with a physical limitation on the vertical number of planes in the array, then ranks provide the limitation in a horizontal direction.

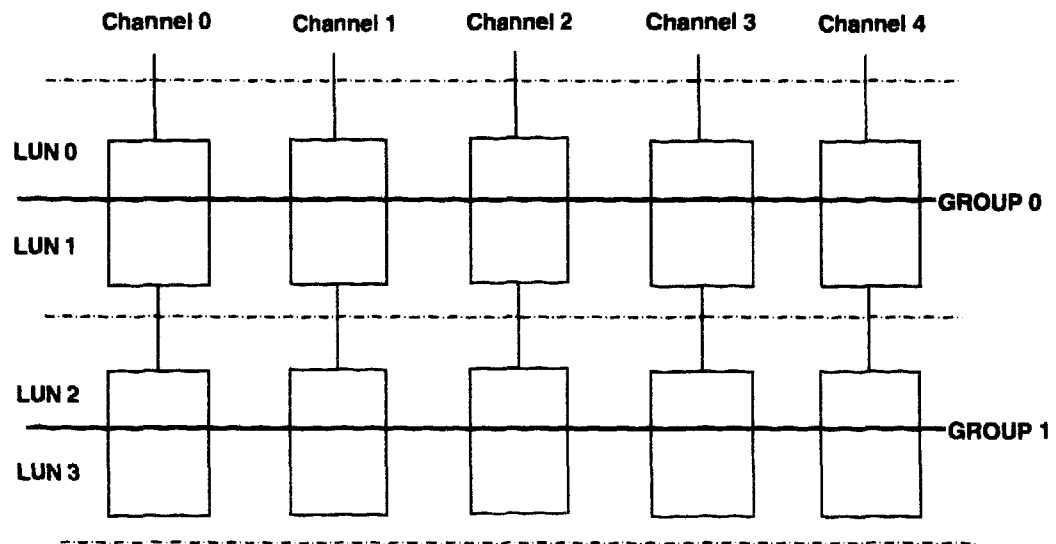
A rank is a physical grouping of drives that cannot be changed (unlike a drive group, which is an arbitrary grouping of drives that can be changed by reconfiguring). One rank of drives can contain up to five drives in the same horizontal plane (in the multirank array) or vertical plane (in the single-rank array). Drives in an array are identified by channel number and SCSI ID. Channel numbers go from 1 to 5 within a rank. All drives in a rank must be on different array channels. The drives are not required to have the same SCSI ID.

2.1.2.5 Partitions

A **partition** is defined as a portion of each drive in a rank. The partition allows multiple LUNs on one rank of drives. This is useful if the host has a capacity limit per LUN.

Figure 2-7 is a conceptual diagram of the drive partition.

Figure 2-7 Drive Partition Diagram



SHR-XR3010-GRA

2.1.2.6 Reconstruction

Reconstruction is done after a failed drive has been replaced in a RAID 0+1, 1, 3, or 5 configuration. It is the process of rebuilding the data on the new drive using data and parity from the other drives or the mirrored drive.

2.1.2.7 Regeneration

Regeneration occurs when data cannot be accessed from a drive in a RAID 0+1, 1, 3, or 5 configuration. The inability to access data can be due to a drive error or failure. Regeneration of data is done from the mirrored pair or from a combination of the other drives and parity.

Regeneration is done on a request-by-request basis and is needed only if the I/O cannot be satisfied. This happens only when a drive is in a *warning* or a *failed* state. For detailed drive status descriptions, refer to Chapter 5.

2.2 StorageWorks RAID Array 110 Subsystem Specifications

This section describes the physical, environmental, and performance specifications for the StorageWorks RAID Array 110 Subsystem product. For additional StorageWorks component specifications, refer to the *StorageWorks Configuration Guide*.

2.2.1 General Specifications

Table 2-1 StorageWorks RAID Array 110 Subsystem General Specifications

Dimensions (height × width × depth)	610 mm × 508 mm × 368 mm 24 in × 20 in × 14.5 in
Power Supply	Standard - 262 W (2 BA35X-HA)
Operating Temperature	+10°C to +40°C (+50°F to +104°F)

The standard power option for the StorageWorks RAID Array 110 Subsystem is two power supplies. The power supply SBB is rated at 131 watts and sufficient power for a fully loaded shelf.

2.2.2 Power Unit Specifications

Each StorageWorks shelf requires either a primary ac or dc power unit. The power unit type is determined by the enclosure ac power distribution unit or dc power controller.

All shelves can have a redundant power unit to ensure that a power unit failure does not disable the shelf. In most cases, battery backup units can be combined with the primary power unit to ensure that in the event of a power unit failure the integrity of the SBB data is maintained while the devices power down. See Table 2-2 and *StorageWorks Shelf Building Block Subsystem User's Guide* for more information about the power units.

Table 2-2 StorageWorks Power Units

Specifications	BA35X-HA	BA35X-HB	BA35X-HC
Power unit type	ac input	dc input	Battery Backup
Input voltage range	90–264 V ac	36–72 V dc	N/A
Nominal input voltage	110 V ac	48 V dc	12 V dc
Autoranging feature	Yes	Yes	N/A
Output voltages	12 V dc 5 V dc	12 V dc 5 V dc	N/A
Output power [†]	131 W	131 W	N/A

[†] Sequential device spin-up at 9-second interval is mandatory.

2.2.3 Environmental Stabilization

To ensure proper operation of Digital storage devices, the SBB temperature must be within 18°C to 29°C (65°F to 85°F). Table 2-3 specifies the time required to thermally stabilize SBBs based on the ambient shipping temperature.

CAUTION

Always stabilize storage devices in the operating environment prior to installation or operation. Otherwise, the media or associated electronics may be damaged when power is applied to the unit.

If condensation *is visible* on the outside of the storage device:

Stabilize the device and the SBB in the operating environment for 6 hours or until the condensation is no longer visible, whichever is longer. *Do not* insert the storage device into the shelf until it is fully stabilized.

If condensation *is not visible* on the outside of the storage device:

Thermally stabilize the device for the amount of time specified in Table 2-3.

Table 2-3 Thermal Stabilization Specifications

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

Technical Information

2.2 StorageWorks RAID Array 110 Subsystem Specifications

2.2.4 Environmental Specifications

The StorageWorks product line environmental specifications listed in Table 2–4 are the same as for other Digital storage devices.

Table 2–4 Environmental Specifications

Condition	Specification
Optimum Operating Environment	
Temperature	+18°C to +24°C (+65°F to +75°F)
Rate of change	3°C (5.4°F)
Step change	3°C (5.4°F)
Relative humidity	40% to 60% (noncondensing) with a step change of 10% or less (noncondensing)
Altitude	From sea level to 2400 m (8000 ft)
Air quality	Maximum particle count .5 micron or larger, not to exceed 500,000 particles per cubic ft of air
Inlet air volume	.026 cubic m per second (50 cubic ft per minute)
Maximum Operating Environment (Range)	
Temperature	+10°C to +35°C (+50°C to +95°F) Derate 1.8°C for each 1000 m (1.0°F for each 1000 ft) of altitude Maximum temperature gradient 11°C/hr (20°F/hr) ±2°C/hr (4°F/hr)
Relative humidity	10% to 90% (noncondensing) Maximum wet bulb temperature: 28°C (82°F) Minimum dew point: 2°C (36°F)
Maximum Nonoperating or Storage Environment (Range)	
Temperature	
Nonoperating	+18°C to +29°C (+65°F to +85°F)
Storage	–40°C to +66°C (–40°F to +151°F)
Relative humidity	
Nonoperating	10% to 90% (noncondensing)
Storage	8% to 95% in original shipping container (noncondensing); otherwise, 50% (noncondensing)
Altitude	From –300 m (–1000 ft) to +3600 m (+12,000 ft) MSL

2.3 StorageWorks RAID Array 110 Subsystem Base Configuration

The StorageWorks RAID Array 110 Subsystem is preconfigured at the factory and ready for use. Preconfigured means that the disk array software has been installed, and the default RAID level and settings have been configured. The five drives have been preconfigured to look as one LUN.

The basic StorageWorks RAID Array 110 Subsystem configuration is as follows:

- One HSZ10-AA controller
- RAID 5
- Five drives—one drive per array channel
- One LUN
- Segment size of 512 blocks
- DWZZA-VA bus adapter (optional)

If this configuration does not meet your needs, it can be modified using the StorageWorks RAID Array 110 utilities. Refer to the appropriate *DEC RAID Utilities User's Guide* for your operating system. If you need to expand your configuration to include additional storage, refer to Chapter 10.


2.3.1 SCSI IDs for the StorageWorks RAID Array 110 Subsystem

Each device connected to the StorageWorks RAID Array 110 Subsystem is a small computer system interface (SCSI) device. A SCSI device is a host computer adapter, or a peripheral controller, or an intelligent peripheral that can be attached to the SCSI bus. SCSI IDs are based upon bit-significant representation of the SCSI address that refers to one of the signal lines (data buses) DB 7-0. In the StorageWorks RAID Array 110 Subsystem, the SCSI IDs are assigned on several different levels as shown in Figure 2-8.

Technical Information

2.3 StorageWorks RAID Array 110 Subsystem Base Configuration

Figure 2–8 StorageWorks RAID Array 110 Subsystem SCSI IDs



SHR-XR3030-GRA

On the host adapter side of the configuration, two connectors provide the SCSI interconnect to the host system either for local termination or for daisy chaining to another host.

The HSZ10-AA controller has two switchpacks for setting the SCSI ID of the controller itself. The HSZ10-AA controller is preset in the factory to the SCSI ID 0 position. The *HSZ10-AA Controller Site Preparation Guide* describes setting SW1 to another SCSI ID.

Three connectors on the HSZ10-AA controller (J3, J4, and J5), shown in Figure 2–8, connect to the five channels of SCSI drive buses. The three connectors plug directly into the backplane of the shelf.

The position of the HSZ10-AA controller in its slot determines its SCSI ID on those buses, which can be either SCSI ID 7 or 6. If the controller is in the BA35X-VA vertical mounting kit, then the left position is SCSI ID 7 and the right (or middle) position is SCSI ID 6. Refer to Controller 1 and Controller 2 in Figure 2–9. Controller 1 has a SCSI ID of 7 and Controller 2 has a SCSI ID of 6.

The SCSI ID of the drives is determined by the backplane of the BA350-EA shelf, and is dependent upon its location in a slot. If the BA35X-EA shelf is in an upright position (as it is when it is contained in the BA35X-VA vertical mounting kit), then the bottom slot has a SCSI ID of 0. The top slot has a SCSI ID of 1. The two slots at the bottom are linked, and 5 and 0 share the connection in the backplane. Refer to Figure 2–9.

2.3 StorageWorks RAID Array 110 Subsystem Base Configuration

2.3.2 HSZ10-AA Controller Location and SCSI Address



The HSZ10-AA controller must be mounted in the slot marked "Controller 1" in Figure 2-9. The position of this controller is important since its SCSI ID on the array channels is set by the backplane. In the Controller 1 slot position, the controller is set to SCSI ID 7.

A second or redundant controller can be placed in the position marked "Controller 2."

Note

The redundant controller option can only be used by a host adapter/operating system environment that can accommodate a redundant controller configuration. Such a configuration is not currently supported.

Figure 2-9 Physical Drive Map



CXO-3589A-MC

2.4 Physical Drive Map

The drive map defines the drives included in the LUN. Individual drives are identified by the channel number and SCSI ID. The channel number and SCSI ID are determined by the physical location of the drive in the array subsystem. In the BA350-EA shelf, the slot number and SCSI ID are the same and are used interchangeably.

The following table shows the physical arrangement of the StorageWorks RAID Array 110 Subsystem base configuration:

SLOT/SCSI ID	CHANNEL	DRIVE
0	-	- *
1	5	Drive 1
2	4	Drive 2
3	3	Drive 3
4	2	Drive 4
5	1	Drive 5

SLOT - Defines the slot position in the BA350-EA shelf
CHANNEL - SCSI bus for the drives on the Array controller
SCSI ID - Drive's SCSI ID on the Array channel

* NOTE: Reserved for host adapters; for example, the DWZZA-VA.

2.5 Logical Drive Map

The logical mapping of the drives is not the same as the physical arrangement. The logical map is a grouping of drives. The default configuration is preconfigured for one LUN and includes the following five drives:

LOGICAL UNIT	(CHANNEL, SCSI ID)
0	(1,5) (2,4) (3,3) (4,3) (5,1)

2.5.1 DWZZA-VA Bus Adapter

An 8-bit differential SCSI system cannot be connected directly to a single-ended 8-bit SCSI system. However, with the proper adapter, a differential SCSI system operating in the 8-bit mode is compatible with a single-ended SCSI bus. The StorageWorks modular storage shelf subsystem uses the DWZZA-VA bus adapter for compatibility between these two SCSI bus types.

The major features of the DWZZA-VA bus adapter are as follows:

- No SCSI device modification is required.
- Supports data transfers at rates up to 10 Mbytes per second.
- Does not use a SCSI device address.
- It converts two physical buses into one logical bus with a total of eight device addresses (0 through 7).
- It can terminate both buses in the end-bus configuration.
- It can terminate either bus in the midbus configuration.
- DWZZA-VA operation is transparent to both buses.

Installation

This chapter contains information on the following:

- Basic configuration installation
- StorageWorks RAID Array 110 Utility for Novell NetWare installation
- StorageWorks RAID Array 110 Utility for MS-DOS installation
- StorageWorks RAID Array 110 Utility for SCO UNIX installation
- StorageWorks RAID Array 110 Utility for OpenVMS VAX installation
- StorageWorks RAID Array 110 Utility for OpenVMS VAX deinstallation

Important Note

The DEC RAID OpenVMS VAX Utility is renamed to the StorageWorks RAID Array 110 Utility for OpenVMS VAX. This guide refers to DEC RAID OpenVMS VAX Utility in code examples only. In all other cases, the new product name is used.

3.1 Basic Configuration Installation

This section discusses installation topics that are specific to the StorageWorks RAID Array 110 Subsystem. Since cabinet installation and service instructions are specific to the cabinet option selected, they are not discussed in this section but can be found in the StorageWorks documentation set.

This section discusses the following:

- Connecting to a host
- Verifying cables and connectors
- Powering on the subsystem
- Functional verification

3.1.1 Connecting to a Host

The HSZ10-AA controller has a fast, wide, differential host interface. This SCSI-3 interface on the controller is provided through two 68-pin, high-density female connectors. The specific cable required to connect to the HSZ10-AA controller (SCSI-3 "P" cable) depends on the type of host/adaptor used. Refer to the sections below for a description of connecting to 16-bit differential, 8-bit differential, and 8-bit single-ended hosts.

3.1 Basic Configuration Installation

3.1.1.1 Terminating the SCSI Bus

Every SCSI bus must be terminated at both ends to operate reliably. How bus termination is accomplished with an HSZ10-AA controller depends upon whether the controller is located at the middle of the bus or at the end of the bus. Refer to the following list of examples:

- If an HSZ10-AA controller is located at the end of the bus, terminate the bus by placing an H879-AA bus terminator (68-pin) on one of the two HSZ10-AA controller 68-pin connectors. Either connector can contain the terminator. Connect a SCSI-3 "P" cable to the other HSZ10-AA controller 68-pin connector.
- If an HSZ10-AA controller occupies a midbus position, attach a SCSI-3 "P" cable to both of the HSZ10-AA controller 68-pin connectors. The bus is then terminated on whatever device is at the end of the bus.

3.1.1.2 Connecting to a 16-Bit Differential Host/Adapter

To connect directly to a 16-bit (wide) differential host, a SCSI-3 "P" cable is needed. The BN21K-series cable provides a male, 68-pin, high-density, right-angle connector to the HSZ10-AA controller, and a male, 68-pin, high-density straight connector to the host. If a right-angle, 68-pin, high-density connection is required on both ends, then the BN21L-series cable should be used.

If the StorageWorks RAID Array 110 Subsystem is at the end of the bus, an H879-AA bus terminator is needed. If the StorageWorks RAID Array 110 Subsystem is in the middle of the bus, use another SCSI-3 "P" cable to daisy-chain to another device.

3.1.1.3 Connecting to an 8-Bit Differential Host/Adapter


Connecting to an 8-bit differential host or adapter varies depending on the connector and termination on that host or adapter.

In all cases, the SCSI-3 "P" cable is used on the HSZ10-AA controller. In the case where the host or adapter uses the SCSI-2 50-pin low-density connector, an additional "Y" transition cable is required.

Specifically, connectivity to an EISA-based host system that uses the KZESA EISA-to-SCSI host adapter requires the following cables:

- The 68-pin, high-density, right-angle connector on the BN21K or BN21L cable connects to the HSZ10-AA controller. The other connector on the HSZ10-AA controller contains the terminator (H879-AA).
- The other end of the cable connects to one of the 68-pin, female connectors on the "Y" transition cable (BN21P-0B). This "Y" cable then connects to the KZESA with the male, 50-pin, low-density straight connector.
- An H879-AA terminator connects to the other female, 68-pin, high-density straight connector, correctly terminating the bus. See Figure 3-1.

Figure 3-1 Typical Y Cable Connection



CXO-3587A-MC

3.1.1.4 Connecting to an 8-Bit Single-Ended Host/Adapter

Connecting the HSZ10-AA controller to an 8-bit, single-ended host or adapter requires the use of the DWZZA-VA single-ended to differential converter. The DWZZA-VA is mounted in a standard 3-1/2 inch system building block (SBB) located in slot 0 of the StorageWorks shelf.

The DWZZA-VA contains a 68-pin, female differential connector on the front of the SBB. The connection from the HSZ10-AA controller to the DWZZA-VA is accomplished using a BN21L cable, which contains a male, 68-pin, right-angle, high-density connector on each end. The HSZ10-AA must also be terminated using the H869-AA terminator in the second 68-pin, female connector on the controller.

The single-ended signals to the DWZZA-VA are provided using a BC09D cable. This cable has a 50-pin, low-density straight connector for the host connection, and a 50-pin, high-density straight connector for insertion into the backplane of the BA350-EA shelf. The section of the BA350-EA shelf that contains the HSZ10-AA controller also contains a column of 50-pin, high-density connectors on the right-hand side. The top connector is connected to the DWZZA-VA in slot 0 through the backplane. This is where the high-density connector of the BC09D cable is inserted. Refer to Figure 3-2.

3.1.1.5 Maintaining Bus Continuity


A trilink connector block (H885-AA) allows the removal of devices from the bus without breaking the continuity of the SCSI bus.

The trilink connector block (H885-AA), shown in Figure 3-3, is a small, metal block with two 68-pin, high-density, female connectors with 2-56 jack standoffs on one side and a single 68-pin, high-density, male connector with 2-56 jackscrews on the other side. Using the trilink connector, the HSZ10-AA controller can occupy either an end-bus or a midbus position, yet remain removable for service.

Installation

3.1 Basic Configuration Installation

Figure 3-2 Connecting the HSZ10-AA Controller to an 8-Bit, Single-Ended Host or Adapter



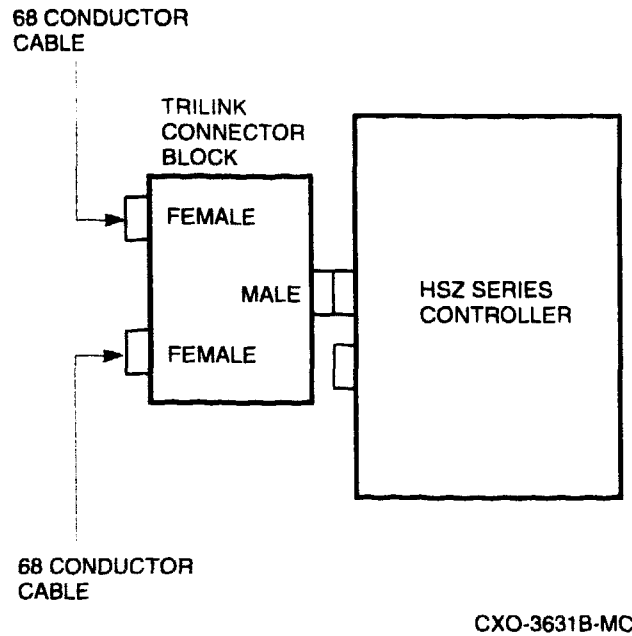
SHR-XR3029-GRA

For midbus connections, follow the steps listed in Table 3-1 and refer to Figure 3-3.

Table 3-1 Midbus Connections

Step	Procedure
1.	Leave one of the two HSZ10-AA controller connectors open.
2.	Connect the second HSZ10-AA controller connector to the male connector on the H885-AA trilink connector block.
3.	Attach two BN21L cables to the other side of the trilink connector block.

Figure 3-3 Trilink Connector Midbus Connection



For end-bus connections, follow the steps listed in Table 3-2 and refer to Figure 3-4.

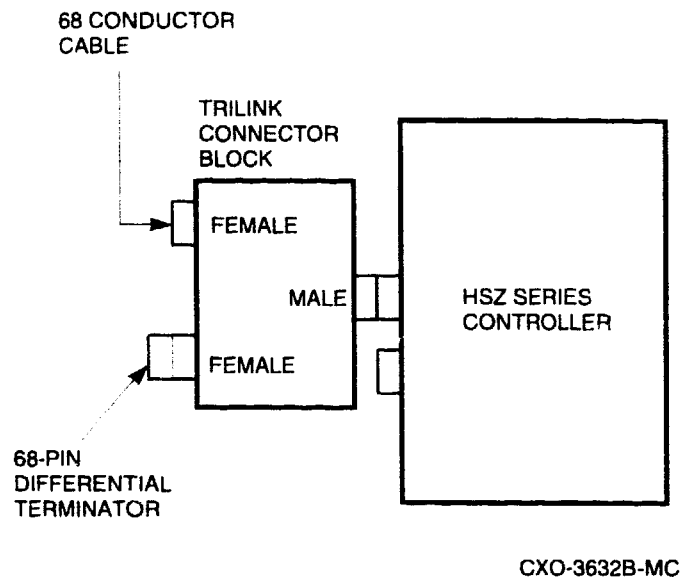
Table 3-2 End-Bus Connections

Step	Procedure
1.	Leave one of the two HSZ10-AA controller connectors open.
2.	Connect the second HSZ10-AA controller connector to the male connector on the trilink connector block.
3.	Connect a BN21L cable to one of the female connectors on the trilink block.
4.	Attach the H879-AA terminator to the other female connector on the trilink block.

Installation

3.1 Basic Configuration Installation

Figure 3–4 Trilink Connector End-Bus Connection



3.1.2 Verifying Cables and Connectors

To verify the physical connections, do the following:

- Ensure that all the cables are securely fastened.
- Ensure that there are no SCSI ID conflicts with the other devices that are connected to the same SCSI bus.
- Ensure that the host SCSI bus is correctly terminated at both ends of the SCSI bus and is not doubly terminated by other SCSI devices. The bus should only be terminated at both ends.
- Ensure that the correct cable and termination scheme, which allows for different SCSI connection schemes, is followed based on the guidelines set forth above.

3.1.3 Powering On the Subsystem

To power on the subsystem do the following:

1. Verify that the ac power switch on the ac distribution unit is in the Off position.
2. Connect the ac input power cable to the ac distribution unit and to the wall receptacle.
3. Verify that the ac input power cables are inserted firmly into the power supply SBBs.
4. Turn on the ac power switch and perform functional verification, as described in Section 3.1.4.

3.1.4 Functional Verification

After powering on the subsystem, verify that all the subsystem components are functioning properly. To do this, perform the following tasks:


1. Verify that both of the LED indicators on all of the power supply SBBs are lit.
2. Verify that all drive SBB activity LEDs initially flash and then go off.

Note

If the drive SBB fault light remains lit, this may not necessarily be a fault condition. It is important to wait until the unit has been given a start command from the host and has gone through its initialization process.

3. Verify that all the HSZ10-AA controller LEDs flash on initial power on and the controller moves to a wait state with binary code 5 and the *Heartbeat* LED is beating. For an illustration of the LED indicators, refer to the *HSZ10-AA Controller Site Preparation Guide* and to Figure 3-5.

Figure 3-5 LED Indicators



SHR-XR3022-GRA

3.2 StorageWorks RAID Array 110 Utility for Novell NetWare Installation

3.2.1 About Installation

The StorageWorks RAID Array 110 Utility for Novell NetWare is installed into the NetWare file server system by the installation program called RMINST.NLM. You do not need to remove an old version of the StorageWorks RAID Array 110 Utility for Novell NetWare since the installation program copies the new version into a different directory than the old version.

RMINST.NLM performs the following tasks when it installs the StorageWorks RAID Array 110 Utility for Novell NetWare:

- Creates the following directories:
 - sys:/system/raidmgr
 - sys:/system/raidmgr/acehelp
- Copies the program and data files from the distribution diskette containing StorageWorks RAID Array 110 Utility for Novell NetWare into the raidmgr subdirectory
- Copies the online help file (raidmgr.nfo) into the sys:/public directory
- Copies the Array Configuration Editor help files into the sys:/system/raidmgr/acehelp subdirectory
- Determines which drivers are in resident memory using the list of supported disk drivers in the raidmgr.cfg file (located in the sys:/system/raidmgr subdirectory)
- Loads the AMD utility if the driver is ASPI (Advanced SCSI Programming Interface) and loaded into memory
- Reports any errors it encounters while copying files, checking drivers, or loading the AMD utilities.

3.2 StorageWorks RAID Array 110 Utility for Novell NetWare Installation

3.2.1.1 StorageWorks RAID Array 110 Utility for Novell NetWare Files

When you install the StorageWorks RAID Array 110 Utility for Novell NetWare, the installation program copies the files listed in Table 3-3 onto the file server.

Table 3-3 Files Copied to the File Server in NetWare

File	Description
raidmgr.nlm	User shell
raidifl.nlm	NetWare file
ace.nlm	Array Configuration Editor
amd.nlm	Array Monitor Daemon
apc.nlm	Automatic Parity Check/Repair
raidmgr.cfg	Configuration file
raidmgr.dat	Menu and form definitions for RAID Manager
raidmgr.msg	Messages displayed by RAID Manager
raidmgr.hlp	Online help for RAID Manager
release.id	Release version
acehelp <dir>	Online help files for the Array Configuration Editor

3.2.2 When to Install the StorageWorks RAID Array 110 Utility for Novell NetWare

Install the StorageWorks RAID Array 110 Utility for Novell NetWare after you have installed the array subsystem and host adapter, if necessary.

3.2.3 Before You Install the StorageWorks RAID Array 110 Utility for Novell NetWare

- If necessary, install the host adapter using the instructions accompanying the host adapter kit.
- Install all the required drivers you need.

3.2.4 Running the Installation Program

There are two ways to install the StorageWorks RAID Array 110 Utility for Novell NetWare:

- If you have an MS-DOS partition on your file server, refer to Section 3.2.4.1 to run the installation program using MS-DOS.
- If you do not have an MS-DOS partition on your file server, refer to Section 3.2.4.2 to run the installation program without using MS-DOS.

Note

The Enter key and the Return key can be used interchangeably. In the procedures that follow, either press the Enter key or press the Return key. The result is the same.

3.2.4.1 Using an MS-DOS Partition

Perform the following steps to run the installation program using an MS-DOS partition:

1. Make sure you have read Section 3.2.3, Before You Install the StorageWorks RAID Array 110 Utility for Novell NetWare.
2. Insert the StorageWorks RAID Array 110 Utility for Novell NetWare distribution diskette in drive A:.
3. Type the following at the NetWare file server console, then press the Enter key:

```
load a:rminst
```

The following screen is displayed:

```
=====
NetWare RAID Manager Installation
=====

Enter DOS drive identifier (default A:):

=====
```

4. Press the Enter key twice. The installation program reads the release.id file from the device, and displays the contents on the screen. This file contains RAID Manager-specific information.

3.2 StorageWorks RAID Array 110 Utility for Novell NetWare Installation

5. Press the Y key and the Enter key to continue with the installation. The installation program does the following:
 - Creates the sys:/system/raidmgr and sys:/system/raidmgr/acehelp directories
 - Copies files into the directories from the distribution diskette
 - Copies the online help files into the sys:/public directory
6. Note when the following message (indicating that installation is complete) is displayed:

```
AUTOEXEC.NCF setup complete
Disk Array Monitor started

NetWare RAID Manager Installation
complete
```

7. If the Disk Array Monitor is not automatically started, unload the StorageWorks RAID Array 110 Utility for Novell NetWare, load AMD, then load the StorageWorks RAID Array 110 Utility for Novell NetWare.

When you have successfully completed the installation, go to Section 5.2 to configure the logical units. If installation is not successful, refer to Section 3.2.5 for installation error messages and their probable causes.

3.2.4.2 Without Using an MS-DOS Partition

Perform the following steps to run the installation program without using an MS-DOS partition.

NOTE

You cannot use MS-DOS if you have made it inaccessible on the file server.

1. Make sure you have read Section 3.2.3, Before You Install the StorageWorks RAID Array 110 Utility for Novell NetWare.
2. Make sure you are using a client workstation that has MS-DOS running and the network software loaded.
3. Log in to the file server from the PC with security privileges equivalent to supervisor.

For example, if the file server name is SERVER1, type the following on the workstation, and then press the Enter key:

```
login SERVER1/supervisor
```

4. Map a virtual volume to the SYS volume.

For example, if the file server name is SERVER1, type the following on the workstation, and then press the Enter key:

```
map s:=SERVER1/SYS:
```

5. Insert the StorageWorks RAID Array 110 Utility for Novell NetWare distribution diskette in drive A:.

6. If necessary, change the default path to the SYS volume path by typing the following, and then press the Enter key:

```
S:
```

7. Change the directory to system by typing the following, and then press the Enter key:

```
cd \system
```

8. Delete the following existing StorageWorks RAID Array 110 Utility for Novell NetWare files by typing the following and pressing the Enter key after each line:

```
del acf.nlm
```

```
del apr.nlm
```

```
del apc.nlm
```

```
del arc.nlm
```

```
del amd.nlm
```

```
del raidmgr.ini
```

Note

If a file is not found, an error message is displayed after you try to delete it. Ignore the error message and continue deleting files.

9. Change the directory to public by typing the following, and then press the Enter key:

```
cd \public
```

10. Delete the following old StorageWorks RAID Array 110 Utility for Novell NetWare file by typing the following, and then press the Enter key:

```
del nwdutil.nfo
```

11. Change the directory to system by typing the following, and then press the Enter key:

```
cd \system
```

12. Create the following directories, if they do not already exist, by typing the following and pressing the Enter key after each line:

```
mkdir raidmgr
```

```
mkdir raidmgr\acehelp
```

3.2 StorageWorks RAID Array 110 Utility for Novell NetWare Installation

13. Copy the following files from the distribution diskette to the directories by typing the following and pressing the Enter key after each line:

```
copy a:\*.* raidmgr
```

```
copy a:\acehelp\*.* raidmgr\acehelp
```

It takes a few moments to copy the files.

14. Delete any occurrences of lines that contain the following commands in the autoexec.ncf file (located in the system directory) using a text editor:

```
load amd
```

15. Add the following commands to the end of the autoexec.ncf file to load AMD:

```
load raidmgr\amd
```

```
search add sys:\system\raidmgr
```

16. Log off the file server.

17. Reboot the server.

When you have successfully installed the StorageWorks RAID Array 110 Utility for Novell NetWare, go to Section 5.2 to configure the LUNs. If installation is not successful, refer to Section 3.2.5 for installation error messages and their probable causes.

3.2.5 Installation Error Messages

Table 3-4 describes installation error messages and probable causes during installation of the StorageWorks RAID Array 110 Utility for Novell NetWare.

Table 3-4 Installation Error Messages

Message	Probable Cause	Action To Take
ERROR Unable to create RAIDMGR and/or RAIDMGR ACEHELP directory.	SYS volume is not mounted, or NetWare is unable to create a directory, or software problem.	First Option: Mount the SYS volume if it is not mounted. Second Option: Try to install again.
ERROR in copying files. Re-run RMINST.	Media error.	First Option: Try to install again. Second Option: Obtain a new copy of the StorageWorks RAID Array 110 Utility for Novell NetWare, and try to install again.

(continued on next page)

Table 3-4 (Cont.) Installation Error Messages

Message	Probable Cause	Action To Take
ERRORS occurred in help copy file. Files not copied: <filenames>	SYS volume is not mounted, or there is a media error.	First Option: Mount the SYS volume if it is not mounted. Second Option: Try to install again. Third Option: Obtain a new copy of StorageWorks RAID Array 110 Utility for Novell NetWare, and try to install again.
ERROR DOS is not loaded.	RMINST procedure has failed.	First Option: Load the utilities manually by copying them from a workstation over the network. Second Option: Reboot the server without performing REMOVE DOS, and reinstall RAID Manager.
ERROR AUTOEXEC.NCF not updated. No space for temporary file.	There is not enough disk space. The temporary file requires 10-12 Kbytes.	First Option: Remove some files to allow for the temporary file to be created.
ERROR AUTOEXEC.NCF not updated. File not found.	There is no autoexec.ncf file available.	First Option: Create an autoexec.ncf and reinstall the StorageWorks RAID Array 110 Utility for Novell NetWare. Second Option: Ignore the error and continue with the installation.
ERROR AUTOEXEC.NCF not updated. File could not be written.	There is a write-protect flag on the autoexec.ncf file or the autoexec.ncf file is damaged.	First Option: Check the permissions on the autoexec.ncf file and remove any write-protects. Second Option: Obtain a new copy of the autoexec.ncf file.
WARNING Array monitor not started. Cause: Disk driver not loaded.	One of the recognized drivers that StorageWorks RAID Array 110 Utility for Novell NetWare supports needs to be loaded.	First Option: Load one of the supported drivers.
WARNING An old version of Array Monitor is running. Complete this installation, then enter "unload AMD" and "load RAIDMGR\AMD."	A file called amd.nlm has already been loaded and you cannot load an nlm file with the same name twice.	First Option: Unload AMD and load RAIDMGR\AMD.
ERROR Array Monitor not started. Cause: AMD.NLM not found.	There is a media error.	First Option: Obtain a new set of diskettes, and repeat the installation procedure.
ERROR RAID Manager Software not found.	The wrong location for the diskette was specified, or there is a media error.	First Option: Specify the diskette location again. Second Option: Obtain a new set of diskettes, and repeat the installation procedure.

3.3 Installation of the StorageWorks RAID Array 110 Utility for MS-DOS

3.3 Installation of the StorageWorks RAID Array 110 Utility for MS-DOS

The StorageWorks RAID Array 110 Utility for MS-DOS is an MS-DOS utility used primarily to configure the RAID array in an MS-DOS and/or SCO UNIX environment. The utility is provided on an MS-DOS Version 5.0 bootable floppy. The kit does not have to be installed.

3.3.1 StorageWorks RAID Array 110 Utility for MS-DOS Kit

The diskette contains the files listed in Table 3-5.

Table 3-5 StorageWorks RAID Array 110 Utility for MS-DOS Kit

Files	Description
COMMAND.COM	MS-DOS Version 5.0
DPTDDL.SYS	KZESA MS-DOS driver
CONFIG.SYS	Configuration file
RAIDMGR.CWA	RAID Manager overlay C-Worthy file
RAIDMGR.EXE	Stand-Alone RAID Manager
RAIDMGR.MSG	RAID Manager message file
RAIDMGR.HLP	RAID Manager help file
RAIDMGR.CFG	RAID Manager configuration file
ACE.EXE	Array Configuration Editor (ACE)
ACF.EXE	Array Configuration Utility (ACF)
FAD.EXE	Array Diagnostic Utility (FAD)
README.PAT	RAID Manager online introduction
READ.ME	RAID Manager online introduction text
PAGE.EXE	Pagination for online introduction
FDISK.EXE	Fdisk utility
ACEHELP <DIR>	Directory of help files
DOS <DIR>	Directory of DOS utilities

3.4 StorageWorks RAID Array 110 Utility for SCO UNIX Installation

The StorageWorks RAID Array 110 Utility for SCO UNIX provides additional functions in the SCO UNIX environment that are not present using the StorageWorks RAID Array 110 Utility for MS-DOS. These include a RAID Monitor for notification of changes in LUN status or drive status, a RAID status utility for displaying current LUN status and drive status, and a Parity Check /Repair utility for checking and repairing parity on the disk array.

This section describes the StorageWorks RAID Array 110 Utility for SCO UNIX, its installation and functionality. Section 12.4 discusses the Parity Check/Repair utility. Section 12.5 describes the array monitor utility, and Section 12.6 the RAID status utility.

3.4.1 About Installation

The StorageWorks RAID Array 110 Utility for SCO UNIX kit is contained on a floppy disk, which contains executables, and shell scripts. Table 3-6 describes the contents of the kit.

Table 3-6 StorageWorks RAID Array 110 Utility for SCO UNIX Kit

File Name	Description
makedpth	Shell script to create device files /dev/dptH0 and /dev/dptH1 used by the utilities
monitor	RAID Monitor executable
monitorsh	Shell script to run the RAID Monitor
parity	RAID Parity Check/Repair executable
paritysh	Shell script to run the Parity Check/Repair utility
raidstat	RAID status utility executable
raidstatsh	Shell script to run the RAID status utility
read.me	Informational file for the user to read describing each of these files
help.txt	Text file for help messages

The next sections describe installation procedures for each of these utilities. Section 12.4, Section 12.5, and Section 12.6 provide further details on each of the utilities, their options, and their error messages.

3.4.1.1 Installing the StorageWorks RAID Array 110 Utility for SCO UNIX Parity Check/Repair Utility

Installation of the StorageWorks RAID Array 110 Utility for SCO UNIX Parity Check/Repair utility involves the following steps:

1. Copy the following files from the floppy disk to any directory of choice:
 - parity
 - paritysh
 - makedpth
 - help.txt
2. Change to superuser mode.
3. Set the PATH variable to point to the directory where the files reside:

```
PATH=$PATH/directory_name
```
4. Type **makedpth** to create the device files in /dev required by the utilities.
5. At this point, the Parity Check/Repair utility can be run. Section 12.4 describes running the RAID Parity Check/Repair utility and its options.

3.4 StorageWorks RAID Array 110 Utility for SCO UNIX Installation

3.4.1.2 Installing the RAID Monitor Utility

Installation of the RAID Monitor utility involves the following steps:

1. Copy the following files from the floppy disk to any directory of choice:
 - monitor
 - monitorsh
 - makedpth
 - help.txt
2. Change to superuser mode.
3. Set the PATH variable to point to the directory where the files reside:
`PATH=$PATH/directory_name`
4. Type **makedpth** to create the device files in /dev required by the utilities.
5. At this point, the RAID Monitor can be started. Section 12.5 describes running the RAID Monitor and its options.

3.4.1.3 RAID Status Utility

Installation of the RAID status utility involves the following steps:

1. Copy the following files from the floppy disk to any directory of choice:
 - raidstat
 - raidstatsh
 - makedpth
 - help.txt
2. Change to superuser mode.
3. Set the PATH variable to point to the directory where the files reside:
`PATH=$PATH/directory_name`
4. Type **makedpth** to create the device files in /dev required by the utilities.
5. At this point, the RAID status utility can be run. Section 12.6 describes running the RAID status utility and its options.

For additional information regarding the StorageWorks RAID Array 110 Utility for OpenVMS VAX installation, refer to the *DEC RAID OpenVMS VAX Utility Release Notes and Installation Guide*.

3.5 StorageWorks RAID Array 110 Utility for OpenVMS VAX Installation Procedure

The software kit (HSZ10010) is distributed on a TK50 cartridge and CDROM and contains utilities to support the HSZ10-AA controller under OpenVMS VAX Version V5.5-2.

3.5.1 Installation Requirements

The HSZ10-AA controller is a SCSI-based array controller supporting RAID 0, 1 and 5 disk arrays. The requirements for running the HSZ10-AA controller on your OpenVMS VAX system are:

1. OpenVMS VAX Version V5.5-2. The HSZ10-AA controller is not supported with versions prior to V5.5-2.
2. A supported StorageWorks RAID Array 110 Subsystem, which uses the HSZ10-AA controller. The subsystem can be connected to the host prior to installation of the StorageWorks RAID Array 110 Utility for OpenVMS VAX, thus avoiding an additional shutdown of the system to install the subsystem after the utility installation is complete.
3. The DEC SCSI Tagged Command Queuing (TCQ) Drivers software kit. This software kit includes the tagged command drivers and support for the HSZ10-AA controller.
4. The StorageWorks RAID Array 110 Utility for OpenVMS VAX software kit. This kit includes the utilities needed to configure, run array Parity Check /Repair, and monitor the array.
5. A supported OpenVMS VAX configuration. Supported systems and hardware are described in Table 3-7.

3.5 StorageWorks RAID Array 110 Utility for OpenVMS VAX Installation Procedure

Table 3-7 TCQ and SCSI Port Driver Support for VAX Systems

Systems	Port Drivers						Adapter
	Queuing Available			Queuing Not Available			
	PKB	PKC	PKR	PKI	PKS	PKN	
MicroVAX 3100							
Model 30	-	Y	-	-	-	-	Native
Model 40	-	Y	-	-	-	-	Native
Model 80	-	Y	-	-	-	-	Native
Model 90	Y	-	-	-	-	-	Native
VAXstation 4000							
Model VLC	-	Y	-	-	-	-	Native
Model 60	-	Y	Y	-	-	N	Native, PMAZ
Model 90	Y	-	Y	-	-	N	Native, PMAZ
VAX System 4000							
Model 100	Y	-	-	N	-	-	Native, KZQSA

Key to Port Drivers

Y—The StorageWorks RAID Array 110 Subsystem can be configured on this port and system.
 N—The StorageWorks RAID Array 110 Subsystem is *not* functional on these ports and systems, although the TCQ kit can be installed and existing SCSI devices will work.

Note the following as well:

- VAXstation 3100-series computers (Models 10, 20, 30, 40, 38, 48, and 76) do not support TCQ. However, if the TCQ kit is installed on one of these systems, existing SCSI devices will still work on the PKN port driver (with a native adapter).
- VAXstation 3520 and VAXstation 3540 computers do not support TCQ. However, if the TCQ kit is installed on one of these systems, existing SCSI devices will still work on the PKS port driver (with a native adapter).
- VAX 4000-series computers other than the VAX 4000 Model 100, do not support TCQ. However, if the TCQ kit is installed on one of these systems, existing SCSI devices will still work on the PKI port driver (with a KZQSA adapter).

3.5.2 Preparing Your System for the Installation

Before performing the installation, follow the steps in this section.

1. Make an image backup copy of the system disk. Backing up the system disk preserves the original system disk in case the install fails. An image backup also stores files contiguously, creating more free blocks on the system disk.

CAUTION

A system failure at a critical point in the installation procedure might render the system unbootable.

2. If you are installing this kit in a VAXcluster environment, this installation procedure will require a reboot of the entire VAXcluster. Digital recommends that you shut down all of the satellite nodes in the cluster at this time.
3. Log in to the SYSTEM account.
4. Make sure the system disk has enough free blocks. The HSZ10010 installation requires 15,000 free blocks. To determine how many free blocks are available, enter the following command:

```
$ SHOW DEVICE/FULL SYS$SYSDEVICE:
```

If your system disk has a sufficient number of free blocks, proceed to the next section.

Note

Digital recommends that no more than 70% of your disk be utilized at any given time. While this installation procedure requires only 15,000 free blocks, your system disk should have a significant amount of free space available to ensure optimal utilization. For more information on disk management, refer to the *VMS Systems Managers Manual* and the *Guide to VMS Performance Management*.

If your system disk does not have sufficient free blocks, create free space by doing the following:

- a. Delete or purge all unwanted files from the system disk.
- b. If there still are not enough free blocks, copy the following files to another disk or tape and delete them from the system disk:
 - All files with file types of JNL, MAP, and LOG
 - All files in SYS\$EXAMPLES and SYS\$TEST
- c. Once you are logged in to the SYSTEM account, make sure that you are the only user:
 - Notify current users that they must log out by entering the following command:

```
$ REPLY/ALL/BELL/SHUTDOWN "Log out for HSZ10 installation please..."
```
 - Prevent nonprivileged users from logging in by entering the following command:

```
$ SET LOGINS/INTERACTIVE=0
```

If you are using the InfoServer for your installation, proceed to Section 3.5.3.

If your system includes DECnet–OpenVMS VAX software and you want to shut it down, enter the following commands:

```
$ RUN SYS$SYSTEM:NCP
NCP> SET EXECUTOR STATE OFF
NCP> EXIT
```

- d. If you are not using the InfoServer, proceed to Section 3.5.5. If you are using the InfoServer, proceed to the next section.

3.5.3 Using the InfoServer During the Installation

If you are using the InfoServer for your installation, you must start the InfoServer Client software and make the CD drive accessible to your system. This section describes these tasks.

To start the InfoServer Client software, complete the following steps:

1. If you use DECnet software, make sure it is already started. To check the status of your DECnet software, enter the DCL command **SHOW NETWORK**. If the system displays the following message, DECnet has not been started:

```
%SHOW-I-NONET, network unavailable
```

Start up the network with the following commands and go to step 3:

```
$ RUN SYS$SYSTEM:NCP
NCP> SET EXECUTOR STATE ON
NCP> EXIT
```

For more information on networking see the *VMS Networking Manual*.

2. If you are not using DECnet software on your system, you must define the System Generation Utility (SYSGEN) parameter **SCSNODE**. To see whether **SCSNODE** is defined on your system, enter the following commands:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> SHOW SCSNODE
```

If the SYSGEN displays a value for **SCSNODE**, proceed to the next step. If SYSGEN does not display a value for **SCSNODE**, enter the following commands (where *node* is a 1- to 6-character unique node name):

```
SYSGEN> SET SCSNODE node
SYSGEN> WRITE CURRENT
SYSGEN> EXIT
```

After you set **SCSNODE**, you must reboot the system. For more information about booting the system, see the OpenVMS VAX upgrade and installation supplement for your VAX computer.

3. Start the InfoServer Client software by entering the following command:

```
$ @SYS$STARTUP:ESS$STARTUP CLIENT
```

Note

Ensure your system node name is correct. If the InfoServer Client software does not find the node name of your system, it will not start.

The **CLIENT** parameter loads the client driver, **ESS\$DADDRIVER.EXE**, and the InfoServer transport driver, **ESS\$LASTDRIVER.EXE**.

4. To start the InfoServer Client software each time the system boots, add the following line to **SYS\$MANAGER:SYSTARTUP_V5.COM**:

```
$ @SYS$STARTUP:ESS$STARTUP CLIENT
```

5. As the startup procedure executes, it displays the following messages:

```
%LASTCP-I-VERSION, LASTDRIVER X1.5 is stopped
%LASTCP-I-ADAINIT, Initializing adapter xxx for LASTDRIVER
%LASTCP-I-STARTED, LASTDRIVER X1.5 started on node yyy
```

6. After you start the InfoServer Client software, you must make the CD drive accessible to your system as follows:

- a. Insert the distribution CD in the CD drive connected to the InfoServer 100.
- b. Enter the following commands:

```
$ RUN SYS$SYSTEM:ESS$LADCP
LADCP> BIND HSZ10010
%LADCP-I-BIND, service bound to logical unit DADSHSZ10010 (_DADn:)
LADCP> EXIT
```

Make note of the device name `_DADn:` in the previous display. You will have to specify this device name during the update.

7. To make sure that the `SYSTEM` account has sufficient quotas and limits, use the OpenVMS Authorize Utility as follows:

- a. Enter the following commands:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN AUTHORIZE
UAF> SHOW SYSTEM
```

- b. Compare the `SYSTEM` account's limits and quotas to the following minimum required values:

Quota Name	Minimum Values
Open file quota (FILLM)	100
Buffered I/O limit (BIOLM)	18
Direct I/O limit (DIOLM)	18
AST limit (ASTLM)	24
Enqueue quota (ENQLM)	200
Buffered byte quota count (BYTLM)	32768

- c. If necessary, adjust the limits and quotas until they are equal to or greater than the required values. You can change each value by entering a command in the following format:

`UAF> MODIFY SYSTEM/limit=new_value`

 For example:

`UAF> MODIFY SYSTEM/DIOLM=18`
 - d. Exit the Authorize Utility by entering the following command:

`UAF> EXIT`
 - e. If you adjust any of the SYSTEM account's parameter or quota values, log out and log in again so that the new values take effect.
8. Once you are logged in to the SYSTEM account, make sure that you are the only user by completing the following steps:
- a. Enter the following command to notify current users that they must log out:

`$ REPLY/ALL/BELL/SHUTDOWN "Log out for DEC SCSI TCQ Driver update."`
 - b. Enter the following command to prevent nonprivileged users from logging in:

`$ SET LOGINS/INTERACTIVE=0`
9. Next, continue as follows:
- If you are updating using the InfoServer system, proceed to Section 3.5.4.
 - If you are not using the InfoServer, proceed to Section 3.5.5.
 - If you are not using the InfoServer system *and* you want to shut down the DECnet software on your system, enter the following commands and then go to Section 3.5.5:

`$ RUN SYS$SYSTEM:NCP`
`NCP> SET EXECUTOR STATE OFF`
`NCP> EXIT`

3.5.4 Using the InfoServer System to Update Your System

If you are using the InfoServer system to update your operating system, you must start the InfoServer Client software and make the CD drive accessible to your system. This section describes these tasks.

To start the InfoServer Client software, perform the following steps:

1. Make sure the system parameter SCSNODE is defined on your system by entering the following commands:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE CURRENT
SYSMAN> PARAMETERS SHOW SCSNODE
```

If SYSMAN displays a value for SCSNODE, proceed to the next step.

If SYSMAN does *not* display a value for SCSNODE, enter the following commands, where *node* is a unique node name up to six characters in length:

```
SYSMAN> PARAMETERS SET SCSNODE node
SYSMAN> PARAMETERS WRITE CURRENT
SYSMAN> EXIT
```

After you set SCSNODE, you must reboot the system. For more information about booting the system, see the upgrade and installation supplement for your VAX computer. For more information about SCSNODE as it pertains to VAXcluster batch and print queues, see the *VMS DCL Dictionary* or the *VMS VAXcluster Manual*.

2. If you use DECnet software, make sure that it is already started. To check the status of your DECnet software, enter the following command:

```
$ SHOW NETWORK
```

If the system displays the following message, DECnet is not running:

```
%SHOW-I-NONET, network unavailable
```

Start the network by entering the following commands:

```
$ @SYS$MANAGER:STARTNET
```

For more information about DECnet network software, see the *Guide to DECnet-VAX Networking*.

3. Start the InfoServer Client software by entering the following command:

```
$ @SYS$STARTUP:ESS$STARTUP CLIENT
```

The CLIENT parameter loads the client driver, ESS\$DADDRIVER.EXE, and the InfoServer transport driver, ESS\$LASTDRIVER.EXE.

As the startup procedure executes, it displays the following messages:

```
%LASTCP-I-VERSION, LASTDRIVER X1.5 is stopped
%LASTCP-I-ADAINIT, Initializing adapter xxx for LASTDRIVER
%LASTCP-I-STARTED, LASTDRIVER X1.5 started on node node
```

To start the InfoServer Client software each time the system boots, add the following line to SYS\$MANAGER:SYSTARTUP_V5.COM:

```
@SYS$STARTUP:ESS$STARTUP CLIENT
```

4. After you start the InfoServer Client software, you must make the CD drive accessible to your system by completing the following steps:
 - a. Insert the OpenVMS distribution CD in the drive connected to the InfoServer system.

3.5 StorageWorks RAID Array 110 Utility for OpenVMS VAX Installation Procedure

- b. Enter the following commands:

```
$ RUN SYS$SYSTEM:ESS$LADCP
LADCP> BIND TCQU1010
%LADCP-I-BIND, service bound to logical unit DADSTCQU1010
(_DADn:)
LADCP> EXIT
```

Make note of the device name `_DADn`: because you must specify this device name during the update procedure.

3.5.5 Matching Update Procedures to System Configurations

Different system configurations require slightly different update procedures. The following list indicates the possible system configurations and the section to which you should refer:

- VAXcluster environment (not including mixed-version VAXclusters):
Section 3.5.5.1
- Mixed-version VAXclusters: Section 3.5.5.1. In mixed-version VAXclusters, only reapply the update to other system disks that are running OpenVMS Version V5.5-2.
- Mixed-architecture VAXclusters: Section 3.5.5.1. In mixed-architecture VAXclusters, only reapply the update to other system disks for OpenVMS VAX systems running Version V5.5-2.
- Local area VAXclusters with one boot server and two system disks:
Section 3.5.5.2
- Standalone system: Section 3.5.6

3.5.5.1 Updating VAXcluster Environments

Use the following procedure to update all VAXcluster environments and local area VAXclusters with one boot server and two system disks (described in Section 3.5.5.2).

Note

After completing this procedure, all the systems in your VAXcluster environment will be running the same version of the OpenVMS VAX operating system.

1. Make sure that you have prepared your system for the update as described in Section 3.5.2.
2. Log in to the SYSTEM account on a node that uses the system disk you are updating.
3. Shut down all other nodes in the cluster that boot from the system disk.
4. Apply the update according to the instructions in Section 3.5.6.
5. If your VAXcluster environment uses several system disks, repeat Steps 1 through 4 in this section for each system disk in the VAXcluster environment.

When the update is complete, perform the post-installation instructions in Section 3.5.7.

3.5.5.2 Updating a Local Area VAXcluster System with One Boot Server and Two System Disks

To update a local area VAXcluster system with one boot server and two system disks, perform the following steps:

1. Make sure that you have prepared your system for the update as described in Section 3.5.2.
2. Log in to the **SYSTEM** account on the boot server.
3. Shut down all other nodes in the cluster that boot from the first system disk.
4. Apply the update to the first system disk, according to the instructions in Section 3.5.6.
5. To update the second disk, perform the following steps:
 - a. Log in to the **SYSTEM** account on a satellite node that boots from the second system disk.
 - b. Shut down all other nodes in the cluster that boot from the second system disk.
 - c. Apply the update to the second system disk, according to the instructions in Section 3.5.6.

When the update is complete, perform the post-installation instructions in Section 3.5.7.

3.5.6 Installing the StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit

To install the StorageWorks RAID Array 110 Utility for OpenVMS VAX kit, complete the following steps:

1. Place the distribution media in the drive. This distribution media contains two savesets named **HSZ10010.A** and **HSZ10010.B** which are the kits for the StorageWorks RAID Array 110 Utility for OpenVMS VAX. (If you are using the InfoServer, you should have already inserted the CD and completed the steps in Section 3.5.3).
2. To start the normal **VMSINSTAL** command procedure, enter the following command:

```
$ @SYS$UPDATE:VMSINSTAL HSZ10010 ddcu:
```

where *ddcu* is the device name of the source drive that holds the installation distribution media.

The following example shows the command line you enter if your distribution media is loaded on a TK50 tape cartridge drive on controller B, with unit number 3, and you are installing **HSZ10010**:

```
$ @SYS$UPDATE:VMSINSTAL HSZ10010 MKA500:
```

If you are updating from a CD drive connected to an InfoServer, the device name is **DADn**. The *n* part of the device name was displayed when you bound to the distribution CD as described in Section 3.5.3.

3. As the installation procedure begins, **VMSINSTAL** displays the following messages:

3.5 StorageWorks RAID Array 110 Utility for OpenVMS VAX Installation Procedure

Installation

VAX/VMS Software Product Installation Procedure V5.5-2

It is 21-APR-1993 at 23:50.

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

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

If you have backed up the system disk, press the Return key and continue.

If you have NOT yet backed up your system disk, do the following:

- a. **Type NO and press the Return key. VMSINSTAL returns to DCL level to permit you to perform the backup.**
- b. **Back up your system disk (see Step 1 in Section 3.5.2).**
- c. **When the backup is complete, restart the installation procedure at Step 1 of this section.**

4. The procedure displays the following message:

Please mount the first volume of the set on ddcu:

* Are you ready?

Make sure the distribution media is in the correct drive, and enter YES.

5. VMSINSTAL continues with the following display:

%MOUNT-I-MOUNTED. HSZ10010 mounted on ddcu:
The following products will be processed:

HSZ10 V1.0

Beginning installation of HSZ10 V1.0 at 23:51.

%VMSINSTAL-I-RESTORE, Restoring product save set A...

The system now displays the following message:

This is the DEC RAID OpenVMS VAX Utility Installation Procedure.

* Do you wish to install the DEC RAID OpenVMS VAX Utility on your system [YES]?

If you are trying to deinstall the HSZ10010 kit, answer NO to this question and go to Section 3.6 now. If you answer YES the system displays the following message:

This Installation procedure will:

* Install the DEC RAID OpenVMS VAX Utility.

In order to install the DEC RAID OpenVMS VAX Utility, a directory structure called HSZ10\$UTIL is created in VMS\$COMMON (i.e.: SYSS\$COMMON:[HSZ10\$UTIL] and the following logicals are defined:

HSZ10\$DIR	- Default location of Array Parity Check and Array Parity Repair log files.
HSZ10\$LOGS	- Default location of Array Status Monitor log files.
HSZ10\$CONFIG	- Location of Predefined Configuration files.
HSZ10\$HELP	- Location of Array Configuration Editor Help files.
HSZ10\$EEPROM	- Location of EEPROM files

* Do you wish to continue the DEC RAID OpenVMS VAX Utility Installation? [YES]

If you do not want to continue with the installation, or if you cannot install the StorageWorks RAID Array 110 Utility for OpenVMS VAX on your system, type NO, press the Return key and the installation procedure will abort. If you are ready to proceed, type YES and press the Return key.

This installation will take approximately 5 minutes to complete.

No further questions will be asked.

Now Installing Utility Images...

%VMSINSTAL-I-RESTORE, Restoring product save set B ...

Now Creating Utility Directories...

%VMSINSTAL-I-SYSDIR, This product creates system directory [HSZ10\$UTIL].

%VMSINSTAL-I-SYSDIR, This product creates system directory [HSZ10\$UTIL.LOGS].

%VMSINSTAL-I-SYSDIR, This product creates system directory [HSZ10\$UTIL.ACEHELP].

%VMSINSTAL-I-SYSDIR, This product creates system directory [HSZ10\$UTIL.CCFILES].

%VMSINSTAL-I-SYSDIR, This product creates system directory [HSZ10\$UTIL.EEPROM].

Now Installing Utility Files...

Now Adding Help Files to System Help Library...

Now Adding Commands to System DCLTABLES...

Now Invoking HSZ10\$STARTUP.COM to define logicals...

-----REMINDER-----

Don't forget to add a call to HSZ10\$STARTUP.COM in your SYSTARTUP_V5.COM to connect any Logical Units and start the Array Status Monitor processes.

Please consult the DEC RAID OpenVMS VAX Utility Installation document for further information about starting the DEC RAID OpenVMS VAX Utility.

HELP is available online by typing

\$HELP RAID_HSZ10

%VMSINSTAL-I-MOVEFILES, Files will now be moved to their target directories...

Installation of HSZ10 V1.0 completed at 23:54

6. Proceed to Section 3.5.7.

3.5.7 Tasks to Perform After the Installation

After you have successfully installed the StorageWorks RAID Array 110 Utility for OpenVMS VAX kit on your system, you must perform the following steps:

1. If needed, configure the HSZ10-AA controller using the Array Configuration Editor (ACE). Refer to the *DEC RAID OpenVMS VAX Utility User's Guide* for further information about configuring, using, and maintaining HSZ10-AA disk arrays.
2. Edit the SYSTARTUP_V5.COM file and add a call to HSZ10\$STARTUP.COM for every logical unit configured in the HSZ10-AA controller. Each call to HSZ10\$STARTUP.COM should use the P1 Keyword START and include a P2 and P3 parameter describing the device name and the polling interval in minutes for the Array Status Monitor.

Note

All calls to HSZ10\$STARTUP.COM should come before mounting any disks.

3.5 StorageWorks RAID Array 110 Utility for OpenVMS VAX Installation Procedure

```

SYSTARTUP_V5.COM
.
.
Purge Operator Logs
.
@SYS$STARTUP:HSZ10$STARTUP START 'ddcu' 5
@SYS$STARTUP:HSZ10$STARTUP START 'ddcu' 5
.
.
Mount Disks

```

3. After this, reboot the system to connect the defined logical units. Upon rebooting, ensure that the Array Status Monitor is started for each HSZ10-AA controller that was defined in the SYSTARTUP_V5.COM file. This can be determined by watching the OPCOM messages displayed as the system boots up.

3.6 Deinstalling the StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit

To deinstall the StorageWorks RAID Array 110 Utility for OpenVMS VAX kit you must follow the procedure in Section 3.5.2 and Section 3.5.5.

Note

Before deinstalling the StorageWorks RAID Array 110 Utility for OpenVMS VAX software, the HSZ10-AA controller must be deinstalled.

The procedure to deinstall the HSZ10-AA controller is as follows:

1. Edit the SYSTARTUP_V5.COM file and remove all calls to the HSZ10\$STARTUP.COM file for the logical units on the HSZ10-AA controller you want to remove from the system.
2. Shut down the system and physically remove the subsystem from the SCSI bus.

To deinstall the StorageWorks RAID Array 110 Utility for OpenVMS VAX kit, complete the following steps:

1. Place the distribution media in the drive. (If you are using the InfoServer, you should have already inserted the CD and completed the steps in Section 3.5.3).
2. Start the normal VMSINSTAL command procedure by entering the following command:

```
$ @SYS$UPDATE:VMSINSTAL HSZ10010 ddcu
```

where *ddcu* is the device name of the source drive that holds the installation distribution media.

The following example shows the command line you enter if your distribution media is loaded on a TK50 tape cartridge drive on controller B, with unit number 3, and you are installing HSZ10010:

```
$ @SYS$UPDATE:VMSINSTAL HSZ10010 MUB3
```

If you are updating from a CD drive connected to an InfoServer, the device name is DAD n . The n part of the device name was displayed when you bound to the distribution CD as described in Section 3.5.3.

3. As the installation procedure begins, VMSINSTAL displays the following messages:

VAX/VMS Software Product Installation Procedure V5.5-2

It is 21-APR-1993 at 23:50.

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

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

If you have backed up the system disk, press the Return key and continue.

If you have NOT yet backed up your system disk, do the following:

- a. Type NO and press the Return key. VMSINSTAL returns to DCL level to permit you to perform the backup.
- b. Back up your system disk (see Step 1 in Section 3.5.2).
- c. When the backup is complete, restart the installation/deinstallation procedure at Step 1 of this section.

4. The procedure displays the following message:

Please mount the first volume of the set on ddcu:

* Are you ready?

Make sure the distribution media is in the correct drive, and enter YES.

5. VMSINSTAL continues with the following display:

%MOUNT-I-MOUNTED. HSZ10010 mounted on ddcu:

The following products will be processed:

HSZ10 V1.0

Beginning installation of HSZ10 V1.0 at 23:51.

%VMSINSTAL-I-RESTORE, Restoring product save set A...

The system now displays the following message:

This is the DEC RAID OpenVMS VAX Utility Installation Procedure.

* Do you wish to install the DEC RAID OpenVMS VAX Utility on your system [YES]?

Answer NO to this question.

The system now asks the following question:

* Do you wish to remove the DEC RAID OpenVMS VAX Utility from your system [NO]?

Answer YES to this question and the system proceeds with the following:

This Deinstallation procedure will:

- * Deinstall the DEC RAID OpenVMS VAX Utility.

The deinstallation procedure will remove all DCL command verbs, help libraries, and command files necessary to run the DEC RAID OpenVMS VAX Utility.

In order to complete the deinstallation of this utility you will have to manually:

3.6 Deinstalling the StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit

- 1) Delete the directory structure and files located in `SYSSCOMMON:[HSZ10$UTIL]`.
- 2) Remove all references to `HSZ10$STARTUP.COM` from `SYSTARTUP_V5.COM`.

after this procedure is done.

* Do you wish to continue DEC RAID OpenVMS VAX Utility Deinstallation [YES]?

If you do not want to continue with the deinstallation, type NO, press the Return key and the installation procedure will abort. If you are ready to proceed, type YES and press the Return key.

No further questions will be asked.

Now Deinstalling Utility Images...

Now Removing Help Modules From System Help Library...

Now Removing Command Verbs From System DCLTABLES...

Installation of HSZ10 V1.0 completed at 23:54

3.6.1 Post Deinstallation Procedure

To complete the deinstallation perform the following steps:

1. Remove all references to `HSZ10$STARTUP.COM` from the system startup files.
2. Delete the directory structure `SYSSCOMMON:[HSZ10$UTIL]`.

Diagnostics and Software

This chapter contains information on the following:

- StorageWorks RAID Array 110 utilities
- Array Configuration Editor (ACE) Utility
- Array Parity Check (APC) and Array Parity Repair (APR) Utilities for OpenVMS VAX
- Array Status Monitor (ASM) Utility for OpenVMS VAX

Note

The Array Parity Check (APC), Array Parity Repair (APR), and the Array Status Monitor (ASM) utilities for all other operating systems are described in Chapter 12.

4.1 StorageWorks RAID Array 110 Utilities

This section contains information on the following:

- What are the StorageWorks RAID Array 110 utilities?
- Using the StorageWorks RAID Array 110 Utility for OpenVMS VAX
- Using the StorageWorks RAID Array 110 utilities under MS-DOS, SCO UNIX, and NetWare
- Moving around in the StorageWorks RAID Array 110 utilities under MS-DOS, SCO UNIX, and NetWare
- StorageWorks RAID Array 110 utilities menus under MS-DOS, SCO UNIX, and NetWare
- RAID Manager LUN and drive status
- RAID Manager device names

4.1.1 What are the StorageWorks RAID Array 110 Utilities?

The StorageWorks RAID Array 110 utilities allow you to configure, monitor, and repair disk arrays. There are three versions of the utilities:

- StorageWorks RAID Array 110 Utility for OpenVMS VAX—this is an OpenVMS VAX program provided on a TK50 cartridge or a CDROM and contains utilities under OpenVMS VAX Version 5.5-2.
- StorageWorks RAID Array 110 Utility for MS-DOS—this is an MS-DOS program provided on a bootable MS-DOS Version 5.0 floppy for use under the MS-DOS and SCO UNIX operating systems.

- StorageWorks RAID Array 110 Utility for Novell NetWare—this is an NLM used for most array purposes in NetWare. It is loaded to your NetWare system during the StorageWorks RAID Array 110 utilities installation.

4.1.2 Using the StorageWorks RAID Array 110 Utility for OpenVMS VAX

There are two ways to use the StorageWorks RAID Array 110 Utility for OpenVMS VAX as follows:

- You can use the ACE Window to configure, monitor, and repair disk arrays.
- You can use the command-line interface, which includes the Array Parity Check (APC), Array Parity Repair (APR), and Array Status Monitor (ASM) utility commands.

4.1.3 Using the StorageWorks RAID Array 110 Utility for MS-DOS

Use the StorageWorks RAID Array 110 Utility for MS-DOS to do the following:

- Configure or reconfigure a LUN in the MS-DOS or SCO UNIX environment using ACE.
- Configure or reconfigure a LUN in the NetWare environment if that unit contains (or will contain) the operating system.
- Restore the LUN containing your operating system if drives fail (one or more drives on a RAID 0 LUN, 2 or more drives with other RAID levels). Note that this involves reformatting the LUN instead of reconstructing the data on it.
- Download controller firmware to a controller attached to the LUN containing your operating system.
- Access LUNs after system boot, if for some reason you cannot access any of your LUNs any other way.

4.1.4 Using the StorageWorks RAID Array 110 Utility for Novell NetWare

Use the StorageWorks RAID Array 110 Utility for Novell NetWare to perform all array tasks that do not involve a LUN containing the operating system. Specifically, use the StorageWorks RAID Array 110 Utility for Novell NetWare to do the following:

- Configure, reconfigure, and modify LUNs that do not contain the operating system using ACE.
- Check and repair array parity on LUNs.
- Check device status for the array devices on your system.
- Change RAID Manager configuration parameters (scheduled parity time, parity file name, and so on).
- Restore a RAID 1 or RAID 5 LUN after a single disk failure (on any LUN, even one containing the operating system).

4.2 Getting Started with the StorageWorks RAID Array 110 Utility for OpenVMS VAX

The StorageWorks RAID Array 110 Subsystem is shipped to you preconfigured as a single RAID 5 LUN. For more discussion of RAID and LUNs, refer to the *StorageWorks RAID Array 110 Subsystem User's Guide*. The StorageWorks RAID Array 110 Utility for OpenVMS VAX software allows you to change this configuration and reallocate the disk storage into multiple LUNs with various RAID levels if desired. For example, you can change the factory configuration of a single large RAID 5 LUN (see Figure 4-5) to two smaller LUNs, a two-drive mirrored RAID 1 configuration, and a three-drive RAID 5 configuration.

There are two ways to use the StorageWorks RAID Array 110 Utility for OpenVMS VAX as follows:

- You can use the ACE Window to configure, monitor, and repair disk arrays.
- You can use the command-line interface, which includes the APC, APR, and ASM utility commands.

This chapter explains first how to use the ACE Window.

This chapter contains the following information:

- Before you begin
- ACE Window features
- Using ACE
- Other getting started tasks

4.2.1 Before You Begin

If you intend to use the preconfigured system as it is, be certain that when the unit was installed, the tasks described as, "Tasks to Perform after the Installation," in the *DEC RAID OpenVMS VAX Utility Release Notes and Installation Guide* were completed during installation. Remember to reboot the system to connect the defined LUNs.

For LUNs to be recognized by the operating system, they must be connected using SYSGEN. Add a line in the system startup file (SYS\$STARTUP:SYSTARTUP_V5.COM) for each LUN that you plan to use.

Naming conventions for devices are described in Section 4.2.2.

For example, you can add the following lines to the system startup file:

```

$ @SYS$STARTUP:HSZ10$STARTUP.COM START DKA500
$ @SYS$STARTUP:HSZ10$STARTUP.COM START DKA501
$ @SYS$STARTUP:HSZ10$STARTUP.COM START DKA502

```

In this case, you can configure up to three LUNs, numbers 0, 1, and 2, on SCSI ID 5.

Note

It is important to have a LUN 0. The system does not recognize a device at that SCSI ID without a LUN 0.

4.2.2 Device Names

The device name used by the StorageWorks RAID Array 110 Utility for OpenVMS VAX is the standard OpenVMS VAX device name:

```

-----OpenVMS VAX Device Prefix ("DK")
-----OpenVMS VAX Device Bus Adapter
DKA400
-----OpenVMS VAX Logical Unit Number (0-7)
-----OpenVMS VAX Device SCSI ID (0-7)

```

4.3 ACE Window Features

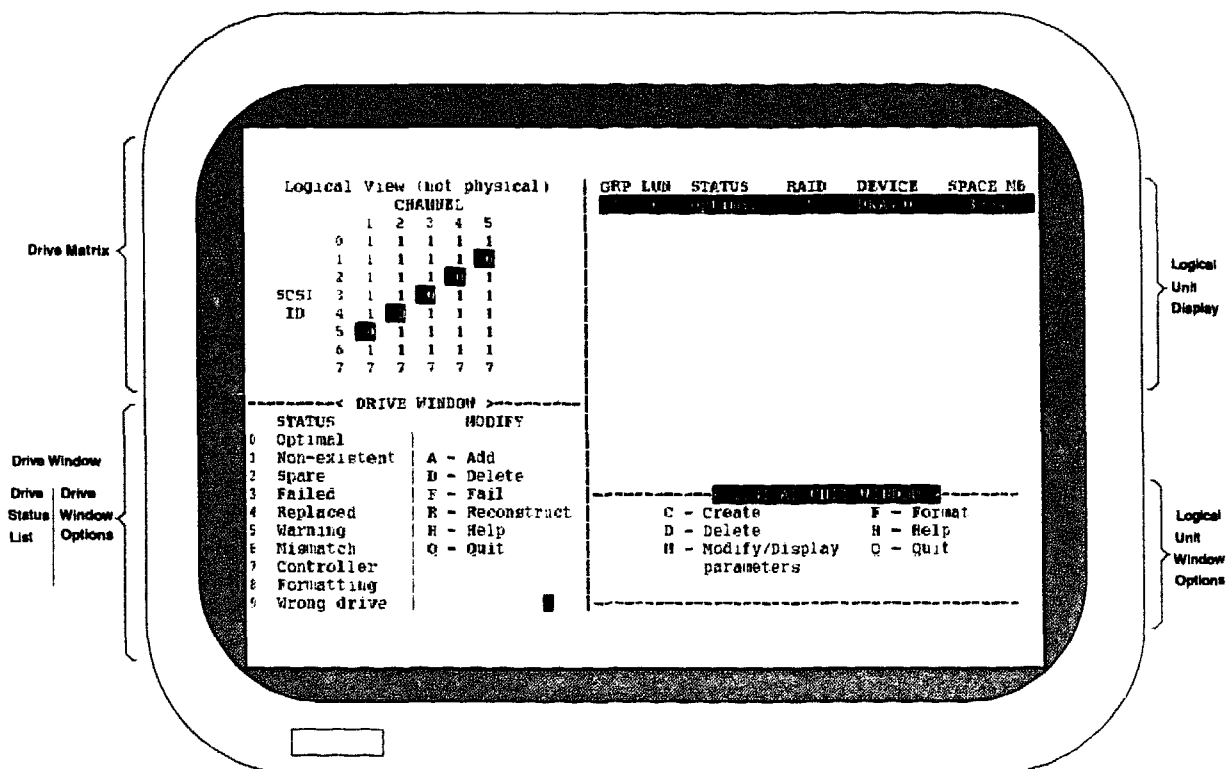
ACE is used for configuring LUNs on your array subsystem. The ACE utility also provides a snapshot view of the current LUN and drive configuration.

When you invoke the ACE utility, ACE displays a window made up of two parts:

- The Drive window (the left-half of the screen)
- The Logical Unit window (the right-half of the screen)

See Figure 4-1.

Figure 4-1 Annotated ACE Window



SHR-XR3031-QRA

4.3.1 ACE Drive Window

The Drive window is divided as follows:

- The upper-half of the window displays a matrix of channels and SCSI IDs showing the current drive status.
- The lower-left side of the window contains a key for the status definitions. For a detailed description of the status codes, refer to Section 4.4.3.2.
- The lower-right side of the window lists the functions you can perform when the Drive window is active. For a detailed description of the functions you can perform from the active Drive window, refer to Section 4.4.3.3.

You can use the Drive window to do the following:

- View the status of all the drives connected to a particular array controller
- View the location of the array controller itself
- Add, delete, fail, or reconstruct drives

The functions of the Drive window are advanced features (for example, you can use the Drive window to add and delete drives and to reconstruct drives). These features are not discussed in detail in this chapter. For more information regarding the Drive window functions, refer to Section 4.4.3.

From the Drive window, use the Tab key to switch to the ACE Logical Unit window described in Section 4.3.2.

4.3.2 ACE Logical Unit Window

The Logical Unit window displays the following information:

- The drive groups of the LUNs (there may be more than one LUN per drive group)
- The LUNs (and drive groups) configured on the array controller you are accessing
- LUN status (see Section 4.4.4.2 for the possible statuses)
- RAID levels of LUNs
- LUN device names
- LUN size (in Mbytes)
- Disk space remaining in spare drives (this is drive group 0)
- Disk space remaining in drive groups with configured LUNs

For more information regarding ACE, refer to Section 4.4.

4.3.3 Using ACE

To start ACE in the OpenVMS VAX environment, type the following command:

```
S ACE device-name
```

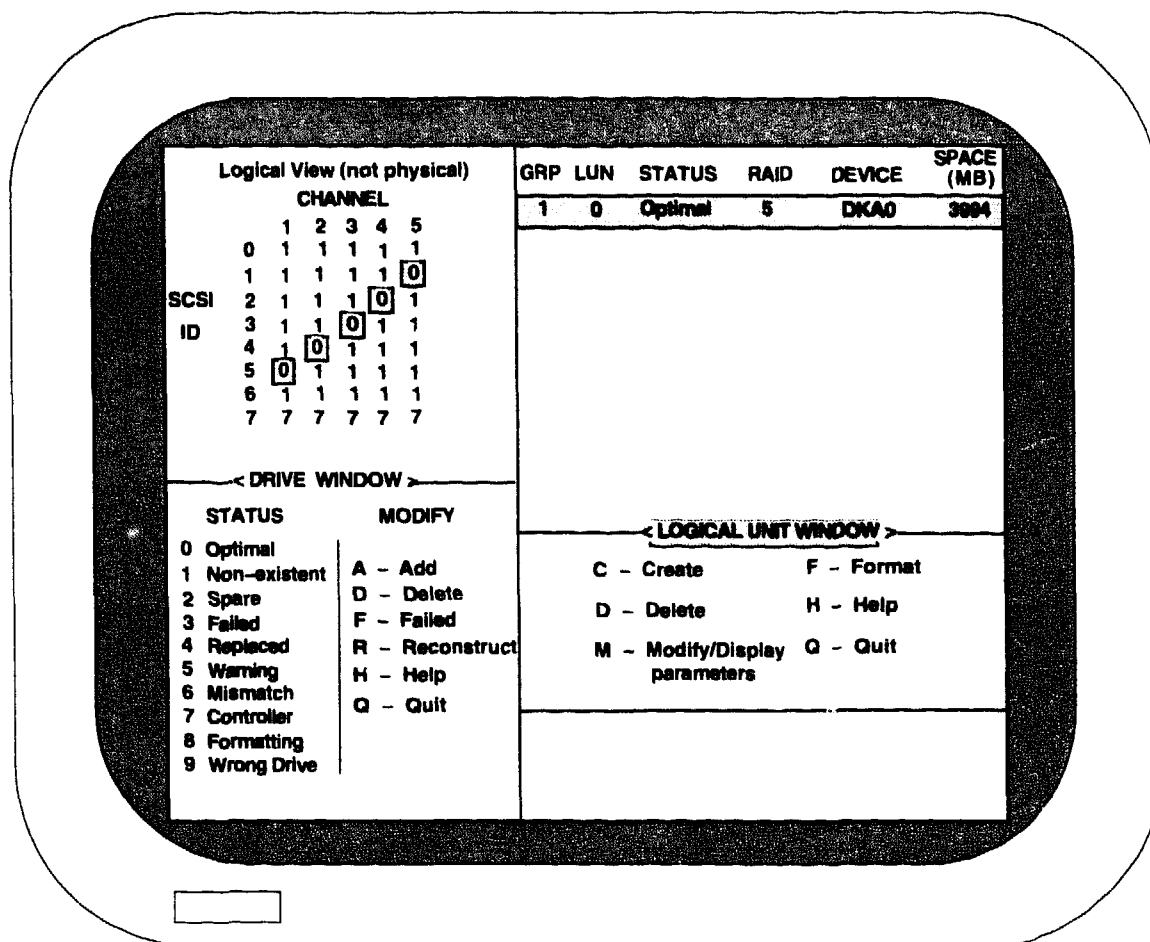
Where device-name is the OpenVAX VMS device name of the array controller, for example, DKA500.

The sys' em displays the ACE window as shown in Figure 4-2. Refer to Section 4.4.2.1 for a description of that display. See Section 4.4.2.2 for a description of how to move around in the screen display.

Diagnostics and Software

4.3 ACE Window Features

Figure 4-2 ACE Window



SHR-XR3027-GRA

The following sections discuss how to use ACE and include the following information:

- Getting Help with ACE
- Deleting a LUN using ACE
- Creating a LUN using ACE
- Exiting the ACE Window

4.3.3.1 Getting Help with ACE

To get help on ACE at the DCL (\$) prompt, type the following command:

```
$ HELP RAID_HSZ10_UTIL
```

The system displays the help options for ACE.

To get help on ACE from the ACE window, invoke ACE and then press the H key. ACE displays the Help options for the Logical Unit window shown in Figure 4-3.

Figure 4-3 Example of a HELP Screen

```
-----< HELP FOR (CONFIGURE LOGICAL) UNIT >-----  
  
N - Navigating In Ace  
O - Overview - Configuring a Logical Unit  
C - Creating A Logical Unit  
D - Deleting A Logical Unit  
M - Modifying/Displaying Parameters  
F - Formatting A Logical Unit  
H - Help Using The Help  
Q - Quit  
  
Enter a Help option, or "q" to quit: █
```

If you enter M(odify), for example, ACE displays help for that option as shown in Figure 4-4.

Figure 4-4 HELP Screen for Modifying and Displaying Parameters

```
-----Modifying/Displaying Parameters-----  
EXPLANATION OF OPTION  
  
Use this option to modify or display the values for the logical  
unit parameters. Note that if you modify the logical unit size,  
segment size, or the segment zero size, or are creating a logical  
unit, the array controller automatically formats the logical  
unit. This means you must back up your data file by file before  
modifying a logical unit parameter.  
  
NOTE: Procedures to modify and display the logical unit  
parameters are included in the help screens following  
the explanation of this option.  
  
- <n> to move down one screen <p> to move up one screen "q" to QUIT -----
```

To leave Help and return to the ACE Window, type Q.

4.3.3.2 Deleting a LUN Using ACE

Because of the factory preconfiguration of the StorageWorks RAID Array 110 Subsystem, you may need to delete a LUN before you can change the configuration. You cannot the change size or RAID level without deleting the existing LUN and then creating the new LUN.

Important

Note that deletion of a LUN results in loss of data on that LUN. See Figure 4-6.

If the device is mounted, dismount the device before you attempt to delete the LUN.

Diagnostics and Software

4.3 ACE Window Features

To delete a LUN, invoke the ACE utility. From the Logical Unit window, use the Spacebar to move the cursor to the LUN you want to delete as shown in Figure 4-5 and type D.

Figure 4-5 Select LUN to be Deleted

The screenshot shows a terminal window titled "Logical Unit". At the top, it says "Logical View (not physical)". Below this is a table with columns: CHANNEL, 1, 2, 3, 4, 5. The rows are labeled "STATUS" and "LUN". The "STATUS" row shows "Optimal" for all LUNs. The "LUN" row shows "0" for all LUNs. Below the table is a "STATUS" section with a list of options: 0 Optimal, 1 Non-existent, 2 Spare, 3 Failed, 4 Replaced, 5 Warning, 6 Mismatch, 7 Control test, 8 Format and, 9 Wrong drive. To the right of the table is a "PROPERTY" section with a list of options: A - Add, D - Delete, F - Fail, R - Reconstruct, H - Help, Q - Quit. At the bottom right is a "LOGICAL UNIT WINDOW" section with a list of options: C - Create, D - Delete, M - Modify/Display Parameters. The status bar at the bottom shows "GRP LUN STATUS RAID DEVICE SPACE RM" with values "1 0 Optimal 5 RAID0 3994".

CHANNEL	1	2	3	4	5
STATUS	0	1	1	1	1
LUN	0	1	1	1	1

STATUS

- 0 Optimal
- 1 Non-existent
- 2 Spare
- 3 Failed
- 4 Replaced
- 5 Warning
- 6 Mismatch
- 7 Control test
- 8 Format and
- 9 Wrong drive

PROPERTY

- A - Add
- D - Delete
- F - Fail
- R - Reconstruct
- H - Help
- Q - Quit

LOGICAL UNIT WINDOW

- C - Create
- D - Delete
- M - Modify/Display Parameters

GRP LUN STATUS RAID DEVICE SPACE RM

1 0 Optimal 5 RAID0 3994

When you type D, the system displays a warning message as shown in Figure 4-6.

Figure 4-6 Deletion Warning Message

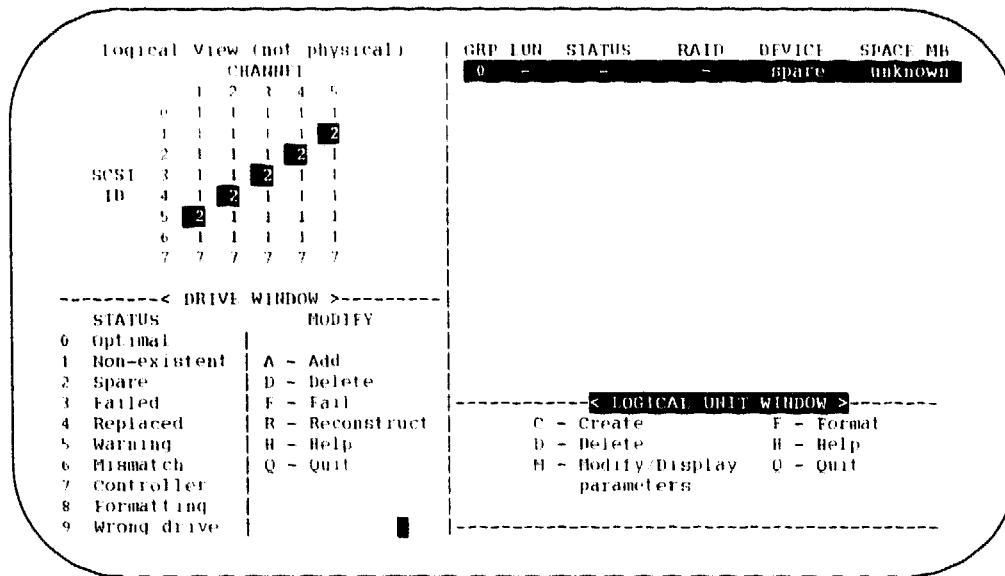
The screenshot shows a warning message box with a dashed border. The text inside reads: "Deletion of the selected logical unit will result in loss of data. Press 'c' to continue. Press 'q' to quit." Below the text is a cursor pointing to the 'c' option.

Deletion of the selected logical unit will result in loss of data.
Press 'c' to continue.
Press 'q' to quit.

Type C to continue to delete the selected LUN.

When you confirm the deletion by typing C, the system displays a window as shown in Figure 4-7.

Figure 4-7 After the Deletion of a LUN Screen



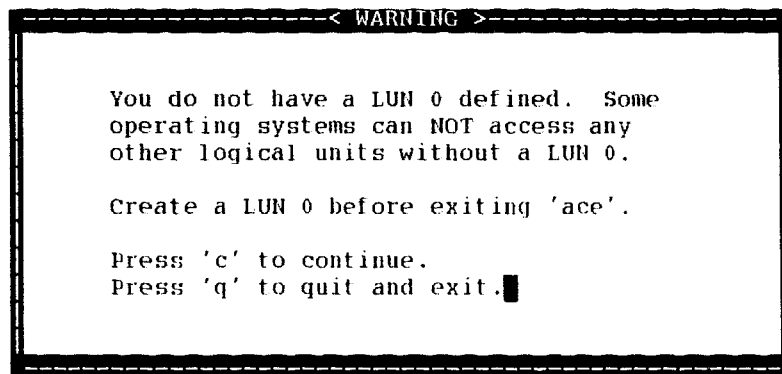
Note that the drive status indicated in the upper-half of the Drive window is now 2 and if you refer to the lower-left side of the Drive window, you see that the status indicates that the drive status is now *Spare*. You can use the space on the drive that you have freed up to create one or more LUNs.

Note

It is recommended that you recreate a LUN 0 before you exit the ACE utility.

If you attempt to exit ACE before you create a new LUN 0, the system displays the Warning message shown in Figure 4-8.

Figure 4-8 Exiting Without LUN 0 Warning



4.3.3.3 Creating a LUN Using ACE

To create a LUN, invoke the ACE utility. From the Logical Unit window, type C. The system displays a window as shown in Figure 4-9.

Figure 4-9 Predefined Configurations

```
-----< PRE-DEFINED CONFIGURATIONS >-----  
  
Enter RAID level choice, or enter  
"q" to quit. █  
  
0 - RAID Level 0  
1 - RAID Level 1  
3 - RAID Level 3  
5 - RAID Level 5
```

Enter the RAID level for the LUN (for example, 1). The system displays a window as shown in Figure 4-10.

Figure 4-10 Create LUN Screen

```
-----< CREATE LUN >-----  
  
RAID Level [ 1 ]  
  
Drive  
  
1      Channel = 5      ID = 1  
2      Channel = 4      ID = 2  
3      Channel = █      ID =  
4      Channel =         ID =  
5      Channel =         ID =  
6      Channel =         ID =  
7      Channel =         ID =  
8      Channel =         ID =  
9      Channel =         ID =  
10     Channel =         ID =  
  
Enter the channel and ID number.  
  
Press <Return> at "Channel" to end  
drive selection, or press "q" to  
quit.
```

Enter the channel and SCSI ID for each drive in the new LUN. Press the Return key. The system prompts you to confirm the channel and SCSI ID. At the *Is the data OK? (y/n/q)* prompt, type Y. The system displays a window as shown in Figure 4-11.

Figure 4-11 Modify/Display Parameters Screen

```

-----< MODIFY/DISPLAY PARAMETERS >-----

LUN 0 Status = creating
LUN 0 Drives = (5,1) (4,2)

PARAMETER                                VALUE

RAID Level                               1
Logical Block Size (bytes)               512
Logical Unit Size (mb)                   998
Segment Size (blocks)                   512
Segment Zero Size (blocks)               0
Delay Interval (tenth's of a second)    1
Blocks Per Delay Interval               256

Are the values shown OK? (y/n/q) █

```

At the *Are the values shown OK? (y/n/q)* prompt, type Y. The system displays an informational message as follows:

Formatting Logical Unit 0

The system automatically formats the LUN. The time it takes to format a LUN depends upon the RAID level selected and the size of the LUN. The informational message remains on the screen until the LUN formatting is complete. The system then displays a window as shown in Figure 4-12.

Figure 4-12 LUN Creation Completed

Logical View (not physical)						GRP	LUN	STATUS	RAID	DEVICE	SPACE MB
CHANNEL1						0	-	-	-	SPARE	unknown
	1	2	3	4	5	1	0	optimal	1	DKA400	998
	0	1	1	1	1						
	1	1	1	1	1						
	2	1	1	1	0						
SCSI	3	1	1	2	1						
ID	4	1	2	1	1						
	5	2	1	1	1						
	6	1	1	1	1						
	7	7	7	7	7						

< DRIVE WINDOW >	
STATUS	MODIFY
0 optimal	
1 Non-existent	A - Add
2 Spare	D - Delete
3 Failed	F - Fail
4 Replaced	R - Reconstruct
5 Warning	H - Help
6 Mismatch	Q - Quit
7 Controller	
8 Formatting	
9 Wrong drive	

< LOGICAL UNIT WINDOW >	
C - Create	F - Format
D - Delete	H - Help
H - Modify/Display parameters	Q - Quit

Note that the LUN is now in an optimal state. The LUN was successfully created.

4.3.3.4 Exiting the ACE Window

To exit the ACE menu, type Q. The system returns you to the DCL (\$) prompt.

4.3.4 Other "Getting Started" Tasks

Once you have set up your LUNs, use the Array Parity Check (APC) and the Array Parity Repair (APR) Utilities. These utilities have a command-line interface which is described in Section 4.5.

4.3.4.1 Description of APC and APR

The Parity Check/Repair utilities perform the following functions:

- Scans the LUN and checks the array parity for each block in the LUN. On a RAID 1 LUN, parity check compares the data on each mirrored pair, block by block.
- Repairs any array parity errors found during the parity check. On a RAID 1 LUN, the array controller changes the data on the mirror disk to make it match the data on the data disk. On a RAID 5 LUN, the controller changes the parity segment so that it is consistent with the data segments.

Note that if the array parity errors result from corrupted data, the data is not repaired, only the array parity. Also, note that you may still lose some data as a result of the power failure or abnormal shutdown, especially if you do not have an uninterruptible power supply (UPS). Data cached in buffers are lost and cannot be reconstructed if you do not have a UPS. This is one of the reasons you should always maintain back-up files, even with a redundant array.

4.3.4.2 When You Should Use APC and APR

You should plan to use the APC and APR utilities as follows:

- After you reconfigure a LUN or add/delete a drive
- To check original parity on a new drive after you remove it from its shipping case
- As part of the regularly scheduled maintenance for your system

The frequency with which you use the APC and APR utilities is at your discretion. You should use the APC and APR utilities whenever a read/write operation is not completed normally, such as when a power failure occurs during an operation.

4.3.4.3 Using APC and APR

To run a parity check on LUN DKA203 and generate a file on bad parity blocks called PARITY_CHECK.LOG, type the following command:

```
S APC DKA203 /OUTPUT=PARITY_CHECK.LOG
```

To run a parity repair on LUN DKA203 using the bad parity block file PARITY_CHECK.LOG, type the following command:

```
S APR DKA203 /INPUT=PARITY_CHECK.LOG
```

4.3.5 Starting the StorageWorks RAID Array 110 Utilities Under MS-DOS, SCO UNIX, and NetWare

There are two ways to start StorageWorks RAID Array 110 utilities, depending on whether you are running the StorageWorks RAID Array 110 Utility for MS-DOS from diskette or running the StorageWorks RAID Array 110 Utility for Novell NetWare.

4.3.5.1 Starting the StorageWorks RAID Array 110 Utility for MS-DOS

To run StorageWorks RAID Array 110 Utility for MS-DOS, perform the following steps:

1. Bring down your operating system in an orderly fashion, if necessary.
2. Insert the diskette with the StorageWorks RAID Array 110 Utility for MS-DOS in Drive A:
3. Reboot your system.
4. Type **raidmgr** and press the Return key.
5. The system starts to load StorageWorks RAID Array 110 Utility for Novell NetWare from the diskette. There is a delay while RAID Manager scans all the SCSI buses on your system, looking for arrays. This delay can take 5-10 seconds per adapter. During this scan, RAID Manager may prompt you if it does not detect any arrays on your system (for example, arrays may not be turned on or are not yet ready).
6. After StorageWorks RAID Array 110 Utility for MS-DOS has located all the arrays on your system, the opening menu is displayed. See Section 4.3.7 for a description of the RAID Manager™ menus.

4.3.5.2 Starting the StorageWorks RAID Array 110 Utility for Novell NetWare

To run the StorageWorks RAID Array 110 Utility for Novell NetWare, perform the following steps:

1. From the system console, type **load raidmgr** and press the Enter key.
2. There is a delay while the RAID Manager scans all the SCSI buses on your system, looking for arrays. This delay can take 5-10 seconds per adapter.
3. After StorageWorks RAID Array 110 Utility for Novell NetWare has located all the arrays on your system, the opening menu is displayed. See Section 4.3.7 for a description of the RAID Manager menus.

Note

Do not leave RAID Manager or the NetWare Install utility running because the Array Monitor Daemon (AMD) will not work.

4.3.6 Moving Around in the StorageWorks RAID Array 110 Utilities Under MS-DOS, SCO UNIX, and NetWare

Table 4-1 lists how to move around in the RAID Manager menus under the MS-DOS, SCO UNIX, and NetWare operating systems.

Table 4-1 Moving within the RAID Manager Menus

To:	Press:
Select an option	The up and down arrow keys to highlight the option you want, then press the Enter key.
Return to the previous menu	The Escape key.
Exit the StorageWorks RAID Array 110 Utility for MS-DOS	The Esc key until a popbox appears asking whether you want to exit StorageWorks RAID Array 110 Utility for MS-DOS. To exit RAID Manager, select Y. To remain in RAID Manager, select N. To return to the Disk Array Services menu, press the Esc key.

4.3.7 Utility Menus Under MS-DOS, SCO UNIX, and NetWare

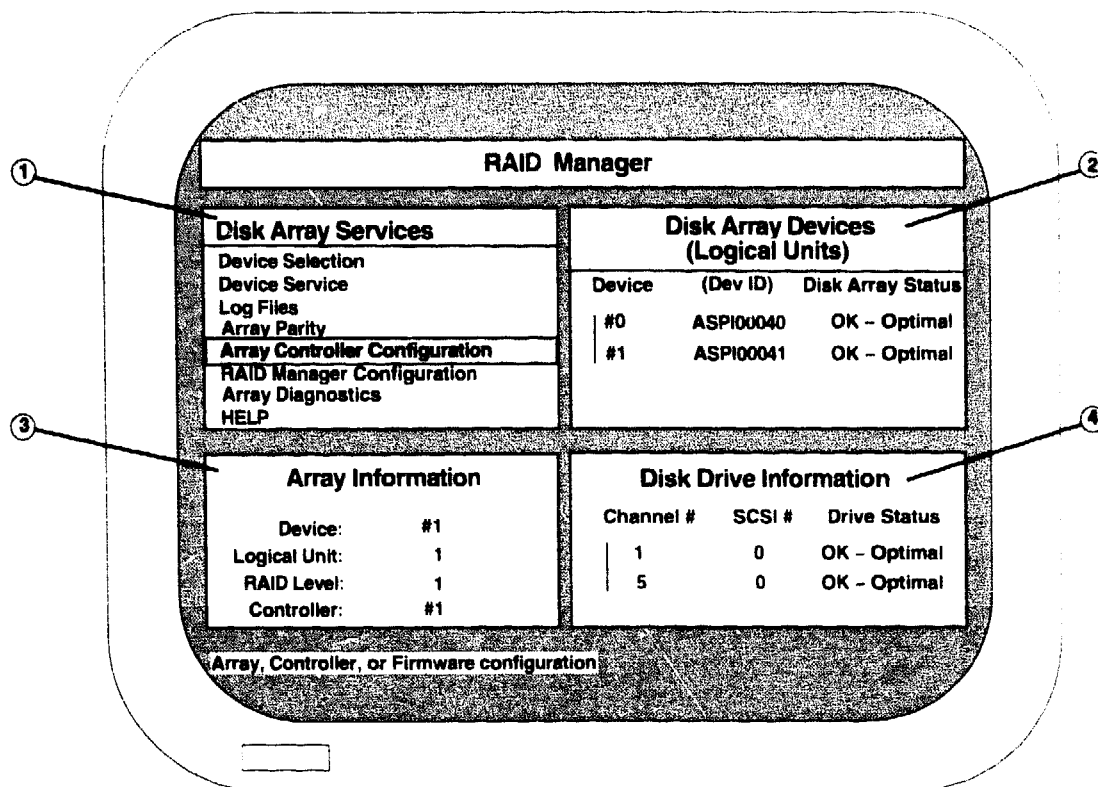
StorageWorks RAID Array 110 utilities operations are performed by selecting functions from the various RAID Manager menus. Most of the menus are self-explanatory and use the same basic type of display. The following sections detail the displays used by the RAID Manager and give a map showing what selections you must make from the menus to perform a given function.

Two menus are given, the initial RAID Manager window (Disk Array Services, Figure 4-13 and the Array Controller Configuration, Figure 4-14). The major difference between the two windows is that the Disk Array Services Window shows all the LUNs on your system, whereas the Array Controller Configuration Window shows only the LUNs attached to a given array controller.

4.3.7.1 Disk Array Services Window

Figure 4-13 shows the Disk Array Services menu. The following is a brief description of each menu item. Note that some options are available only on the StorageWorks RAID Array 110 Utility for Novell NetWare, while others are accessible by both the NetWare and the StorageWorks RAID Array 110 Utility for MS-DOS.

Figure 4–13 Disk Array Services Window



SHR-XR3024-GRA

① Disk Array Services Window

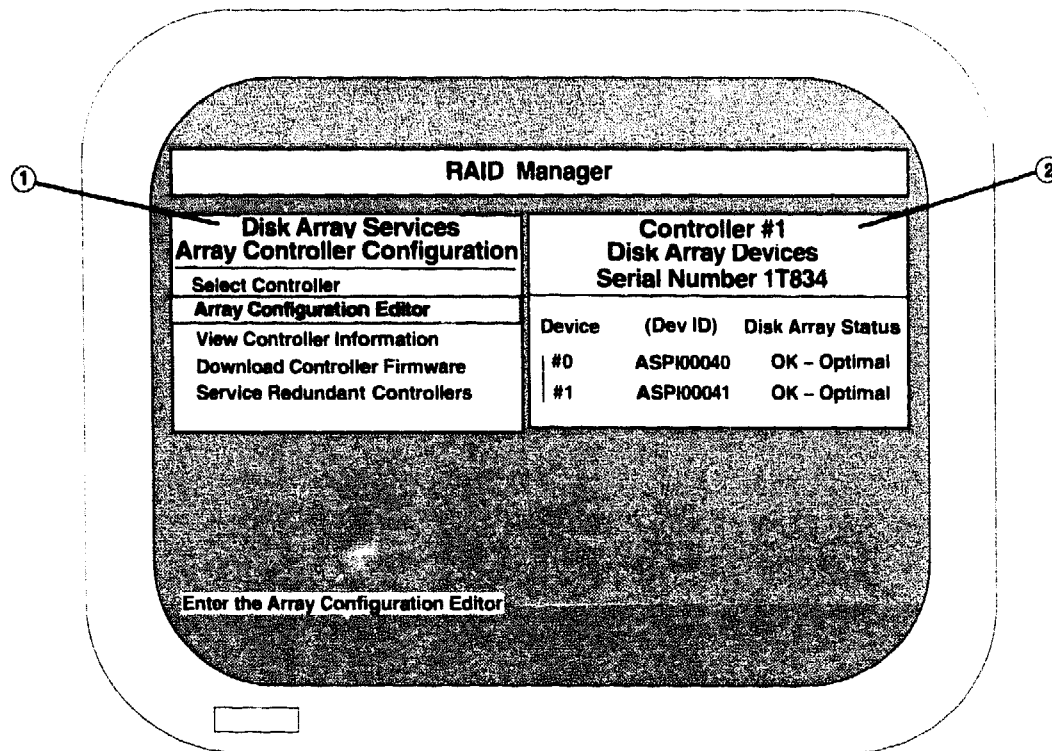
- **Disk Array Services Menu**—Lists the functions you can select. Only the highlighted items can be selected on your system.
- **Device Selection**—Selects the LUN you want to access or for which you want to display information. These options are not be highlighted if only a single LUN exists.
- **Device Service (NetWare only)**—Restores a LUN after one or more disk failures (see Section 8.7).
- **Log Files (NetWare only)**—Displays the RAID Manager Activity and Error logs (see Section 12.3).
- **Array Parity (NetWare only)**—Checks and repairs array parity (NetWare only, see Section 12.2).
- **Array Controller Configuration**—Configures LUNs (see Section 5.3).
- **RAID Manager Configuration (NetWare only)**—Sets the RAID Manager configuration.
- **Array Diagnostics**—Runs diagnostics on the disk array.
- **Help**—Displays help. Note that you can get online help at any point by pressing the F1 key. Press the Esc key to exit help.

- ② **Disk Array Devices Display**—Shows all the disk array devices LUNs attached to your system.
 - **Device**—The device number of the LUN. For the StorageWorks RAID Array 110 Utility for MS-DOS, this number is assigned by the RAID Manager (see Section 4.3.9.1). In the StorageWorks RAID Array 110 Utility for Novell NetWare, this number is assigned by NetWare and is the same device number used for all NetWare operations (see Section 4.3.9.2).
 - **Dev ID**—The device ID of the LUN. For the StorageWorks RAID Array 110 Utility for MS-DOS, this device ID is assigned by the RAID Manager (see Section 4.3.9.1). In the StorageWorks RAID Array 110 Utility for Novell NetWare, this is the same device ID used for all NetWare operations (see Section 4.3.9.2).
 - **Disk Array Status**—The current status of the LUN. LUNs marked as anything other than *OK - Optimal* require immediate attention (see Section 4.3.8 for a description of the statuses displayed).
- ③ **Array Information Display**—Gives further information on the LUN currently highlighted in the Disk Array Devices display.
 - **Device**—The device number of the LUN (the same device number displayed in the Disk Array Devices display).
 - **Logical Unit**—The logical unit number of the logical unit.
 - **RAID Level**—The RAID level of the LUN.
 - **Controller**—The RAID Manager-assigned number of the array controller attached to the LUN. RAID Manager numbers array controllers sequentially, in the order in which they are found.
- ④ **Disk Drive Information Display**—Gives information on the drives contained in the LUN currently highlighted in the Disk Array Devices display.
 - **Channel #**—The channel number of the drive.
 - **SCSI #**—The SCSI ID of the drive.
 - **Drive Status**—The current status of the drive. Drives marked as anything other than *OK - Optimal* require immediate attention (see Section 4.3.8 for a description of the statuses displayed).

4.3.7.2 Array Controller Configuration Window

Figure 4-14 shows the Array Controller Configuration Window. The following is a brief description of each menu item.

Figure 4–14 Array Controller Configuration Window



SHR-XR3025-GRA

- ① **Array Controller Configuration Menu**—Lists the functions you can select. Only the highlighted items can be selected.
 - **Select Controller**—Selects the array controller you want to access. In a single controller configuration, you cannot select this option.
 - **Array Configuration Editor**— Configures LUNs attached to the highlighted controller (see Section 4.4 and Section 5.3).
 - **View Controller Information**—Views information on the highlighted controller.
 - **Download Controller Firmware**—Downloads new array controller firmware to the highlighted array controller.
 - **Service Redundant Controllers**—Performs redundant path management. This feature is currently not available.
- ② **Disk Array Devices Display**—Shows the disk array devices (LUNs) attached to the currently selected controller. (The controller is identified by controller number and serial number.) To change the currently selected controller, use the *Select Controller* function on the Disk Array Services menu.
 - **Controller #**—The RAID Manager-assigned number of the array controller attached to the LUN. RAID Manager numbers array controllers sequentially, in the order in which they are found.
 - **Controller serial number**—The serial number of the array controller attached to the LUN displayed.

- **Device #**—The device number of the LUN. For the StorageWorks RAID Array 110 Utility for MS-DOS, this number is assigned by RAID Manager (see Section 4.3.9.1). In the StorageWorks RAID Array 110 Utility for Novell NetWare, this number is assigned by NetWare and is the same device number used for all NetWare operations (see Section 4.3.9.2).
- **Dev ID**—The device ID of the LUN. For the StorageWorks RAID Array 110 Utility for MS-DOS, this device ID is assigned by the RAID Manager (Section 4.3.9.1). In the StorageWorks RAID Array 110 Utility for Novell NetWare, this is the same device ID used for all NetWare operations (see Section 4.3.9.2).
- **Disk Array Status**—The current status of the LUN. LUNs marked as anything other than *OK - Optimal* require immediate attention (see Section 4.3.8 for a description of the status displayed).

4.3.7.3 StorageWorks RAID Array 110 Utilities Road Map

The following table lists the StorageWorks RAID Array 110 utilities tasks, the RAID Manager menu selections required from the opening menu to perform the task, and where to look in this guide for more information.

Task	Menu Selections	Section(s)
Change Array Monitor Daemon (AMD) parameters (NetWare)	RAID Manager Configuration General Configuration	Section 12.3
Change scheduled parity check time (NetWare)	RAID Manager Configuration General Configuration or Array Parity Change Automatic Parity Time	Section 12.2
Check array parity (NetWare)	Array Parity Check and Repair Parity or Check Parity (no repair)	Section 12.2
Configure a LUN	Array Controller Configuration Array Configuration Editor	Section 4.4 and 5.3
Delete a LUN	Array Controller Configuration Array Configuration Editor	Section 4.4 and 5.3
Format a LUN after drive failure (NetWare)	Device Service Format disk array device	Section 8.7
Get help for selected function	Select function, then press the F1 key	
Get help with procedures	Help	
Modify a LUN	Array Controller Configuration Array Configuration Editor	Section 4.4 and 5.3
Read activity log (NetWare)	Status Messages View RAID Manager Activity Log	Section 12.3
Read error messages (NetWare)	Status Messages View RAID Manager Error Log	Section 12.3
Read parity report (NetWare)	Array Parity View Parity Report	Section 12.2
Replace a failed drive (NetWare)	Device Service Replace the disk drive	Section 8.7

4.3.8 RAID Manager LUN and Drive Status

Table 4–2 explains the LUN status displayed in the Disk Array Devices window.

Table 4–2 LUN Status from the Disk Array Devices Window

Status	Meaning
Degraded	A drive in the LUN has failed and the LUN is now in degraded mode (RAID 1 and 5 only). The LUN is operational in the degraded mode, but the failed drive should be replaced as soon as possible. Select the LUN and check the Disk Drive Information window to determine the status of the drives.
Degraded—Warning	A drive in the degraded LUN has been put in warning. Note that this means one drive has failed and a second drive is about to fail. In a RAID 1 LUN, this status indicates the most serious condition possible. Steps must be taken immediately to replace the bad drives to prevent the permanent loss of data.
Drive Failures	More than one drive in a RAID 5 LUN has failed, and the LUN is no longer operating. All data on the LUN has been lost.
Formatting	The LUN is not available because it is being formatted.
Ok - Optimal	The LUN is fully operational.
Ok - Reconstructing	The LUN is functioning in degraded mode and a drive in the unit is being reconstructed. Select the LUN and check the Disk Drive Information window to determine the status of the drives.
Ok - Warning	The LUN is still optimal, but one or more drives are in a warning state because of a read or write error. Note that in a RAID 0 LUN, the LUN may not be usable. Select the LUN and check the Disk Drive Information window to determine the status of the drives.
Wait For Format	The LUN is not accessible because it needs to be formatted. In NetWare, use the <i>Format disk array device</i> selection in the <i>Device Service</i> menu to format the LUN. In the StorageWorks RAID Array 110 Utility for MS-DOS, use the format option in the Array Configuration Editor.
Wrong Drive	The wrong drive was replaced (that is, the drive that was replaced was not a failed drive but was an optimal drive).

Table 4–3 explains the status of drives as displayed in the Disk Drive Information window.

Table 4–3 Drive Status

Status	Meaning
Ok - Optimal	The drive is functioning correctly.
Failed	The drive has failed and is no longer functioning. The drive should be replaced as soon as possible.
Formatting	The drive is currently being formatted.
Newly Replaced	The drive was just replaced.
Reconstructing	The data on the drive is currently being reconstructed.
Warning	<p>The drive has been put into a warning state as the result of a read or write error. The severity of this status depends to some extent on the RAID level of the LUN:</p> <ul style="list-style-type: none"> • On all RAID levels, this status could indicate a minor read or write error, in which case the LUN is still usable. • In RAID level 0, this is the most serious status the array controller assigns to a drive. The LUN may no longer be accessible. • In RAID 1, this status is assigned if the second drive in a mirrored pair fails after the first drive has failed. The LUN may no longer be accessible. <p>In all cases, the <i>Warning</i> drive should be replaced as soon as possible.</p>

4.3.9 RAID Manager Device Names

The RAID Manager displays a device number and a device ID for each array LUN on your system. These names are for reference purposes only, as you do not have to enter them. Both the device numbers and the device IDs are different depending on whether you are running the StorageWorks RAID Array 110 Utility for MS-DOS or the StorageWorks RAID Array 110 Utility for Novell NetWare.

4.3.9.1 StorageWorks RAID Array 110 Utility for MS-DOS Device Names

During the initial SCSI bus scan, StorageWorks RAID Array 110 Utility for MS-DOS assigns a device number and a device name (Dev ID) to each array LUN it finds on your system.

- The device number is assigned for reference purposes. You have as many device numbers as you have array LUNs attached to your system. Other non-array storage devices are not included in this numbering. The StorageWorks RAID Array 110 Utility for MS-DOS assigns these numbers in the order it detects the LUN (which depends on the type of host adapter and where it is installed).
- The device name (Dev ID) identifies the hardware location of the LUN. The device name used by the StorageWorks RAID Array 110 Utility for MS-DOS has the following format:

`adapter_type device_number`

For example:

`ASPI00040`

- **adapter_type**—This identifies the interface type of the host adapter connected to the controller. The currently supported interface is ASPI (Advanced SCSI Programming Interface).
- **device_number**—This 5-digit number provides the hardware address location of the LUN. The definition of these digits is shown in the following table.

Digit	Meaning
1	I/O bus number
2	Controller number
3	SCSI bus number
4	Controller SCSI ID
5	Logical unit number

In this example, the I/O bus number is 0, the controller number is 0, the SCSI bus number is 0, the controller SCSI ID is 4, and the LUN is 0.

4.3.9.2 NetWare Device Name

NetWare assigns each LUN a device number and a device ID.

- The device number is assigned for reference purposes. This number is the same number used by NetWare to refer to the LUN in other NetWare utilities (such as Install). LUNs are numbered along with the other storage devices on your system.

- The device ID is the NetWare-assigned device ID of the LUN. This device ID is the same ID used by NetWare to refer to the LUN in other NetWare utilities (such as Install). The digits in the device ID have the following meanings:
 - The first and second digits indicate the type of host adapter. These digits are assigned by NetWare. For example, 8B is the KZESA host adapter.
 - The third digit is the board number assigned to the host adapter by the NetWare system.
 - The fourth digit is the SCSI ID of the array controller.
 - The fifth, or last, digit is the LUN.

For example, a device ID of 8B061 identifies the host adapter as a KZESA, board 0, array controller SCSI ID 6, LUN 1.

4.4 Array Configuration Editor (ACE)

This section contains information on the following:

- What is ACE?
- Starting ACE
- Drive window
- Logical Unit window

4.4.1 What is ACE?

The Array Configuration Editor (ACE) is used for configuring LUNs on your array subsystem. Although your array is shipped preconfigured from the factory, you can change that configuration. For example, if your subsystem came with three preconfigured, 5-drive, RAID 5 LUNs, you can change one of those units to a RAID 0 LUN. Use ACE to change the configuration.

Although ACE allows you to do other functions (such as reconstruct data on a drive), some of these functions can more easily be done using other functions in the RAID Manager. The only operation you must use ACE for is configuring or reconfiguring LUNs.

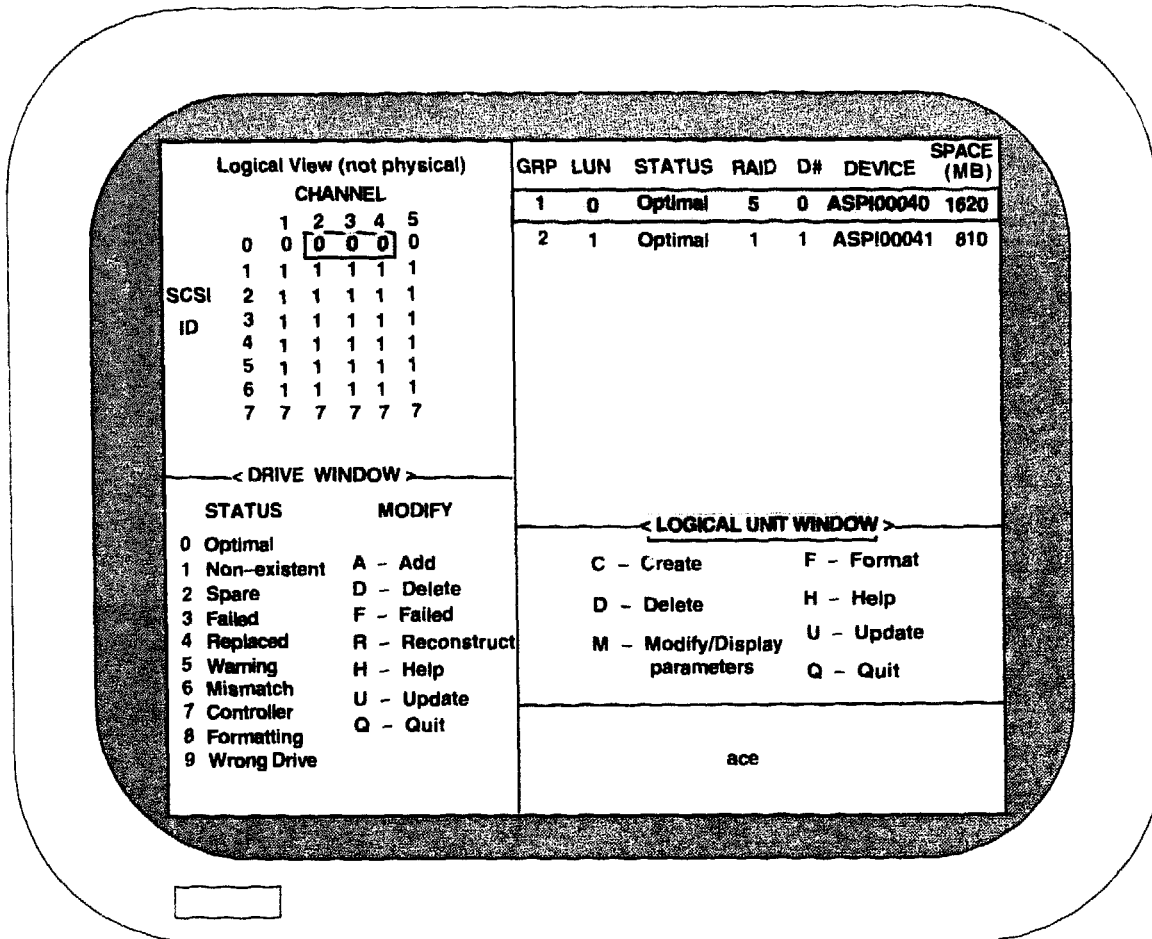
For more information on how to configure or reconfigure LUNs, see Section 5.3.

4.4.2 Starting ACE

ACE is an option on the RAID Manager menus. To start ACE:

1. Start the RAID Manager. See Section 4.1 for information on starting the RAID Manager. Note that you need to use the StorageWorks RAID Array 110 Utility for MS-DOS in the MS-DOS or SCO UNIX environments, or in NetWare if the LUN you want to configure contains (or will contain) your operating system.
2. From the *Disk Array Services* menu select the *Array Controller Configuration* option.
3. From the *Array Controller Configuration* menu select the *Array Configuration Editor* option.
4. The ACE screen displays. Refer to Section 4.4.2.1 for a description of that display. See Section 4.4.2.2 for a description of how to move around in the screen display.

Figure 4-15 ACE Window



SHR-XR3026-GRA

4.4.2.1 ACE Screen Window

Figure 4-1 shows the ACE display. The display is made up of two parts as follows:

- The Drive window (the left-half of the screen)
- The Logical Unit window (the right-half of the screen)

You can tell which window is active by the label (for example, *Drive Window* or *Logical Unit Window*) that is currently highlighted. When you start ACE, you are in the Logical Unit window.

The active window determines which array devices you can select and what parts of the other window are highlighted.

Diagnostics and Software

4.4 Array Configuration Editor (ACE)

Table 4–4 Understanding the ACE Window

When the Logical Unit Window is Active...	
You Select:	Other Window Highlights:
A LUN, Drive Group 0, or a row with space remaining to be configured by pressing the arrow keys or the space bar.	The drives contained in the selected LUN. When you select Drive Group 0, all drives that have a <i>Spare</i> status (2) highlight.
When the Drive Window is Active...	
You Select:	Other Window Highlights:
An individual drive by pressing the arrow keys, the space bar, or the Enter key.	The drive group that contains the selected drive.

4.4.2.2 Navigating in ACE

To navigate or select elements in the ACE, use the keys shown in Table 4–5.

Table 4–5 ACE Navigation Keys

Key	Function in the Drive Window	Function in the Logical Unit Window
Tab	Moves the cursor to the Logical Unit window.	Moves the cursor to the Drive window.
Spacebar	Moves the cursor to all available drives in the drive matrix.	Moves the cursor down in the LUN display.
Down Arrow	Moves the cursor down in the drive matrix.	Moves the cursor down in the LUN display. Moves the cursor to each LUN that can be changed.
Up Arrow	Moves the cursor up in the drive matrix.	Moves the cursor up in the LUN display.
Right Arrow	Moves the cursor right in the drive matrix.	None
Left Arrow	Moves the cursor left in the drive matrix.	None
Enter	Moves the cursor to each drive that is configured into a drive group.	Moves the cursor to each LUN parameter. Executes commands. For example, when you change the LUN parameters, you type Y and press the Enter key.

(continued on next page)

Table 4–5 (Cont.) ACE Navigation Keys

Key	Function in the Drive Window	Function in the Logical Unit Window
Q	Quits the following: ACE The window for selecting a RAID level The window for selecting the Channel and ID of drives <i>Modify/Display Parameters</i> option <i>Help</i> option	ACE The window for selecting a RAID level. The window for selecting the Channel and ID of drives. <i>Modify/Display Parameters</i> option. <i>Help</i> option.

4.4.2.3 Quitting ACE

Press the Q key to quit ACE when the cursor is in the drive window or the Logical Unit window, as indicated in Table 4–5.

4.4.3 Drive Window

The Drive window is the left-half of the ACE utility display (Figure 4–15). You are in the Drive window when the *Drive Window* label in the middle of the window is highlighted. You switch between this window and the Logical Unit window using the Tab key.

Use the Drive window to:

- View the status of all the drives connected to a particular array controller
- View the location of the array controller
- Modify drive status

The upper-half of the Drive window, the Drive Matrix, displays the current drive status. See Section 4.4.3.1 for more information about the Drive Matrix. Notice that an abbreviated definition of each drive status is always displayed in the lower left-hand corner of the Drive window, labeled *Status*. For example, the drive with SCSI Channel 1, SCSI ID 0 is an optimal drive—its status is 0, or *optimal*. For descriptions of each drive status, see Section 4.4.3.2.

The Drive window also shows you the operations you can perform while the Drive Window is active. These operations, or options, are displayed under the *Modify* label.

4.4.3.1 Drive Matrix

The Drive Matrix is located in the upper-half of the Drive window. The Matrix shows the status of each potential drive position on the array controller you are accessing. See Section 4.4.3.2 for a description of the drive status displayed in the Drive Matrix.

The Drive Matrix displays the drives according to a logical order, by SCSI ID, not according to the physical location of the ranks.

The Drive Matrix also indicates the location of the array controller, which has the status *controller* (status of 7), as shown in Figure 4–15.

Diagnostics and Software

4.4 Array Configuration Editor (ACE)

4.4.3.2 ACE Drive Status

ACE displays the following possible status for drives in the Drive Matrix.

Drive Status	Description
0 - Optimal	The drive is operating at an optimal level.
1 - Non-existent	No drive is physically connected to the array at this position.
2 - Spare	The drive is connected to the array, but not to a LUN.
3 - Failed	The drive was failed by the array controller or by the user and must be replaced.
4 - Replaced	The drive has just been replaced, is formatting, or is being reconstructed.
5 - Warning	<p>The drive has been put into a warning state as the result of a read or write error. The severity of this status depends to some extent on the RAID level of the LUN:</p> <ul style="list-style-type: none">• On all RAID levels, this status could indicate a minor read or write error, in which case the LUN is still usable.• In RAID level 0, this is the most serious status the array controller will assign to a drive. The LUN may no longer be accessible.• In RAID 1, this status is assigned if the second drive in a mirrored pair fails after the first drive has failed. The LUN may no longer be accessible. <p>In all cases, the <i>Warning</i> drive should be replaced as soon as possible.</p>
6 - Mismatch	The array controller sensed that the drive has an unexpected sector size, capacity, serial number, SCSI Channel, or ID. That is to say that these are different from what the array controller expected.
7 - Controller	This status is not a drive status—it indicates the location of the array controller in the array.
8 - Formatting	The drive is currently being formatted.
9 - Wrong drive	The wrong drive was replaced.

4.4.3.3 Drive Window Options

The following options are given in the Drive window. For more information on how to perform drive options, see Section 8.6.

- **Add a Drive**—Use the *Add* drive option to add a drive that has a status of *Non-existent* (1). After adding the drive, the drive status changes to *Spare* (2).
- **Delete a Drive**—Use the *Delete* drive option to delete a drive that has a status of *Spare* (2). After physically removing the drive, use this option to change the drive status of the removed drive to *Non-existent* (1).

This option can be bypassed by rebooting the subsystem or server. The status on the removed drive changes to *Non-existent* (1).

- **Fail a Drive**—Use the *Fail* drive option to fail a drive with a drive status of *Warning* (5) or *Optimal* (0). After failing the drive, the status changes to *Failed* (3).

Important Note:

Do not fail a drive in a RAID 0 LUN because there is no parity. Do not fail a drive in a RAID 1, 3, or 5 LUN if the LUN is already degraded.

- **Reconstruct Drive Data**—Normally, you do not need to initiate drive data reconstruction because the array controller automatically reconstructs a replaced drive. However, if your array does not automatically reconstruct replaced drives, you must use this selection to start reconstruction. You can only reconstruct drive data on a RAID 1 or 5 LUN with a single-drive failure.
- **Update Drive Status (NetWare only)**—The ACE display is updated in NetWare every 10 seconds. This option allows you to rescan the devices and update the drive status.
- **Help**—Help menus are available for drive options.
- **Quit (NetWare only)**—Exit the ACE utility when the Drive window is selected.

4.4.4 Logical Unit Window

The Logical Unit window is the right-half of the ACE utility display (see Figure 4-15). You are in the Logical Unit window when the *Logical Unit window* label in the middle of the window is highlighted. Switch between this window and the Drive window using the Tab key.

Use the Logical Unit window to view the LUN status and the device name of each LUN. Also use this window to create, delete, or format LUNs, and to modify /display the LUN parameters.

The upper-half of the Logical Unit window, the Logical Unit Display, displays information about the drive groups and LUNs attached to the controller you are accessing. For more information on the Logical Unit display, see Section 4.4.4.1.

Also shown in Figure 4-15 are the operations you can perform while the Logical Unit window is active. For more information on the Logical Unit window options, see Section 4.4.4.3.

4.4.4.1 Logical Unit Display

The Logical Unit Display is located in the upper-half of the Logical Unit window. It shows the following information:

- The drive groups of the LUNs (there may be more than one LUN per drive group)
- The LUNs (and drive groups) configured on the array controller you are accessing
- LUN status (see Section 4.4.4.2 for the possible statuses)
- RAID levels of LUNs
- NetWare device numbers (NetWare only)
- LUN device names
- LUN size

- Disk space remaining in spare drives (this is drive group 0)
- Disk space remaining in drive groups with configured LUNs

For example, in Figure 4–15, LUN 0 has been configured as RAID level 5. It has an *Optimal* status, the device name of *ASPI00040* and 1620 MB of space. LUN 0 is in drive group 1.

4.4.4.2 ACE LUN Status

LUN Status means Logical Unit Status (LUN stands for logical unit number and is also referred to as *logical unit*). LUNs can have 5 possible statuses as listed in Table 4–6.

Table 4–6 LUN Status from the ACE Window

Status	Description
Creating	You are presently creating this LUN.
Dead	The LUN is no longer functioning. Either: <ul style="list-style-type: none"> • You have changed LUN parameters and have not yet reformatted the unit. • Two or more drives have failed. • The wrong drive was replaced. • Some component connected to the LUN failed or returned an unexpected value (wrong SCSI ID, wrong channel number, and so on). <p>Check the status of the LUN's drives in the Drive Matrix to help determine the cause of the error.</p>
Degraded ¹	The LUN is operating in degraded mode (that is, it is still functioning, but data must be reconstructed using data and parity from the good drives). Either: <ul style="list-style-type: none"> • A single drive has failed. • The replaced drive is still being formatted. • Some component connected to the LUN failed or returned an unexpected value (wrong SCSI ID, wrong channel number, and so on). <p>Check the status of the LUN's drives in the Drive Matrix to help determine the cause of the error.</p>
Optimal ²	The array is operating at an optimal level.
Reconstructing	The array controller is currently reconstructing the LUN.

¹If this is a RAID 1 LUN, it may not be accessible even if the LUN Status is Degraded. To determine if a RAID 1 array has drive problems, check the status of the drives in the LUN in the Drive Matrix.

²If this is a RAID 0 LUN, it may not be accessible even if the LUN Status is Optimal. To determine if a RAID 0 array has drive problems, check the status of the drives in the LUN in the Drive Matrix.

Note

Not all statuses are displayed for all RAID levels. For example, a RAID 0 LUN will never have a status of *Degraded*.

4.4.4.3 Logical Unit Window Options

The following options are given in the Logical Unit window. For more information on how to perform these options, see Section 5.3.

- **Create a Logical Unit**—Use the *Create* LUN option to create a LUN in two ways: by configuring *Spare* drives, or by configuring available drives in an existing drive group.
- **Delete a Logical Unit**—Use the *Delete* LUN option to delete a LUN to create a new LUN, or to change the RAID level or drives in an existing LUN.

Caution

Copy the data on a LUN to back-up media before deleting the LUN. Once the LUN is deleted, you can not access the LUN—the data is lost.

- **Modify/Display Logical Unit Parameters**—Use the *Modify/Display Logical Unit Parameters* option to modify and display the current values for the following LUN parameters:
 - Logical Unit Status (display only)
 - Logical Unit Drives (display only)
 - RAID Level (display only)
 - Logical Block Size in bytes (display only)
 - Logical Unit Size in megabytes
 - Segment Size in blocks
 - Segment Zero Size in blocks
 - Reconstruction Delay Interval in tenths of a second
 - Reconstruction Blocks Per Delay Interval

Note that the first three items are also displayed in the Drive Matrix and Logical Unit Display.
- **Format A Logical Unit**—Use the *Format* option to format a LUN and restore it to a working state. Also, use the *Format* option to restore any LUN with multiple drive failures. Note that the RAID Manager automatically formats any newly-created LUN.
- **Update LUN Status (NetWare only)**—The ACE display is updated in NetWare every 10 seconds. This option allows you to rescan the devices and update the LUN status.
- **Help**—Help menus for LUN options.
- **Quit**—Exit the ACE utility when the Logical Unit window is selected.

4.5 Array Parity Check (APC) and Array Parity Repair (APR) Utilities for OpenVMS VAX

This section contains information on the following:

- Overview
- APC command
- APR command

4.5.1 Overview

The Array Parity Check (APC) and the Array Parity Repair (APR) utilities run a parity check and/or repair on a specified LUN. Parity check/repair applies only to RAID 1 and RAID 5 LUNs. RAID 0 does not have array parity, and therefore cannot be checked and repaired. RAID 1 does not have parity either, but the parity check compares data on the mirrored drives.

Parity check/repair cannot be run on a degraded RAID 1 or RAID 5 LUN.

Parity Check/Repair performs the following functions:

- Scans the LUN and checks the array parity for each block in the LUN. On a RAID 1 LUN, parity check compares the data on each mirrored pair, block by block.
- Repairs any array parity errors found during the parity check. On a RAID 1 LUN, the array controller changes the data on the mirror disk to make it match the data on the data disk. On a RAID 5 LUN, the controller changes the parity segment so that it is consistent with the data segments.

Note that if the array parity errors resulted from corrupted data, the data is not repaired, only the array parity. Also, note that you may still lose some data as a result of a power failure or abnormal shutdown, especially if you do not have an uninterruptible power supply (UPS). Data cached in buffers are lost and cannot be reconstructed if you do not have a UPS. This is one of the reasons you should always maintain back-up files, even with a redundant array.

These utilities can be invoked either interactively or in a batch job. Although the parity check can be run at any time, it should be run on a periodic basis, and immediately after a reboot. It is recommended that a batch job be set up to run the parity check and repair on a daily basis.

The Array Parity Check utility does parity checking on the LUN, and outputs a list of corrupted parity blocks to either SYS\$OUTPUT or to a file you specify. The output file is used as the input source for the Array Parity Repair utility, which expects to see a list of blocks to be repaired.

Parity information is maintained on a logical block number basis. During a parity check operation, informational messages are sent to SYS\$OUTPUT indicating the range of logical blocks currently being checked. Any corrupted parity blocks found are displayed (or written to a file if the /OUTPUT qualifier is specified).

APC

Parity check on a LUN.

Format

APC device-name

Command Qualifiers

/OUTPUT=file_spec

Restrictions

None.

Parameters

device-name

Specifies the device name of the LUN of the device. This device name is of the form DKwxyz where:

- DK is the OpenVMS VAX device prefix.
- w is the OpenVMS VAX device bus adapter.
- x is the OpenVMS VAX device SCSI ID (0 to 7).
- y is a reserved field and always set to 0.
- z is the OpenVMS VAX device LUN (0 to 7).

Description

The APC command runs a parity check on a given LUN indicated by device-name.

Command Qualifiers

/OUTPUT=file_spec

Used to specify an output file containing a list of corrupted parity blocks found by APC. If the /OUTPUT qualifier is specified, a value for file-spec is required. If no /OUTPUT qualifier is specified, then the corrupted block list is sent to SYS\$OUTPUT.

Example

```
$ APC DKA203 /OUTPUT=PARITY_CHECK.LOG
```

Runs a parity check on LUN DKA203 and generates a file on corrupted parity blocks called PARITY_CHECK.LOG.

APR

Parity repair on a LUN.

Format

APR device-name

Command Qualifiers

/INPUT=file_spec

Restrictions

None.

Parameters**device-name**

Specifies the device name of the LUN of the device. This device name is of the form DKwxyz where:

- DK is the OpenVMS VAX device prefix.
- w is the OpenVMS VAX device bus adapter.
- x is the OpenVMS VAX device SCSI ID (0 to 7).
- y is a reserved field and always set to 0.
- z is the OpenVMS VAX device LUN (0 to 7).

Description

The APR command issues a parity repair on a given LUN indicated by device-name.

Command Qualifiers**/INPUT=file_spec**

Used to specify an input file to contain the list of corrupted parity blocks. When the /INPUT qualifier is specified, a value for file-spec is required. If no /INPUT qualifier is specified, then the corrupted block list is read from SYS\$INPUT.

Example

```
S APR DKA203 /INPUT=PARITY_CHECK.LOG
```

Runs a parity repair on LUN DKA203 using the corrupted parity block file PARITY_CHECK.LOG.

4.6 Array Status Monitor (ASM) Utility for OpenVMS VAX

This section contains information on the following:

- Overview
- ASM command

4.6.1 Overview

The Array Status Monitor (ASM) utility checks the array controller for events such as Degraded Logical Units or Drive Warning conditions. Any change in a LUN or drive status from the optimal state will be flagged by the Array Status Monitor. The operator is notified by means of OPCOM and the event is time-stamped and recorded in a log file. Any LUN and/or drive that is in any state other than optimal continues to set off alarms to OPCOM until the problem is rectified. This insures that corrective action will be taken to return the array to an optimal condition.

The Array Status Monitor automatically runs at startup through a call to the HSZ10\$STARTUP.COM file to start a given LUN. The system notifies you by means of an OPCOM message when the monitor begins. It runs as a background process, checking LUN and drive statuses at the specified time interval.

ASM

Invokes the Array Status Monitor (ASM).

Format

ASM device-name

Command Qualifiers

/INTERVAL=time

/LOG_FILE=file_name

/STATUS

/SYNCHRONIZE

Restrictions

None.

Parameters

device-name

Specifies the device name of the LUN of the device. This device name is of the form DKwxyz where:

DK is the OpenVMS VAX device prefix.

w is the OpenVMS VAX device bus adapter.

x is the OpenVMS VAX device SCSI ID (0 to 7).

y is a reserved field and always set to 0.

z is the OpenVMS VAX device LUN (0 to 7).

Description

The ASM command invokes the Array Status Monitor utility to track changes in the status of LUNs and drives. This utility is automatically invoked by the HSZ10\$STARTUP.COM file when a START command is issued for each LUN. See the *DEC RAID OpenVMS VAX Utility Release Notes and Installation Guide* for further information on invoking the HSZ10\$STARUTP.COM file at startup.

Use this command when you need to manually invoke the utility, or to run the monitor with options other than those specified at startup.

Command Qualifiers

/LOG_FILE=file-spec

Specifies an output file containing a time-stamped log of all events found by the Array Status Monitor. If the /LOG_FILE qualifier is specified, a value for file-spec is required. If no /LOG_FILE qualifier is specified, then a default log file is created in the directory specified by the logical HSZ10\$LOGS with the name 'device-name'_ASM.LOG.

/INTERVAL [=number-minutes]

Enables array controller polling and specifies the polling interval in minutes.

If the **/INTERVAL** qualifier is specified with no value, the Array Status Monitor polls the disk array at the specified LUN every 5 minutes. If the **/INTERVAL** qualifier is specified with a value, the Array Status Monitor polls the disk array at the specified LUN every specified number of minutes.

If no **/INTERVAL** qualifier is specified the Array Status Monitor does not poll, open a log file, or attempt to send any message to the operator.

/SYNCHRONIZE

Used to synchronize the system time with the array controller's internal time.

/STATUS =(item,...)]

Displays LUN or drive status information.

Valid items:

- **LUN**—display LUN status information
- **DRIVE**—display a map of drive status information similar to the information reported by the Array Configuration Editor utility.

Examples

1. \$ **SM DKA402 /INTERVAL=10**

Invoke the Array Status Monitor with an interval of 10 minutes.

2. \$ **ASM DKA500 /STATUS**

Display status for LUN DKA203. An example of this format follows:

LUN STATUS - DKA500: Optimal: Optimal condition

RAID_Level	: 5
LUN_Type	: 1
LUN_Block_Size	: 512
Drive_Sector_Size	: 512
LUN_Number_of_Blocks	: 8179712
LUN_Segment_Size	: 512
Segment_Zero_Size	: 0
Type_of_Mode_Page_to_Sense	: 0
Auto_Detect_Replace_Disabled	: FALSE
Degraded_Modes_Writes_Disabled	: FALSE
Array_Assurance_Disabled	: FALSE
AEN_Polling_Enabled	: TRUE
Parity_Verification_Enabled	: TRUE
Wr_Parity_Verification_Enabled	: TRUE
Block_Reconstructed	: 0
Reconstruction_Frequency	: .1 sec
Reconstruction_Amount	: 256
Lun_Number	: 0
Drives (Chan, ID)	: (1,0) (2,1) (3,2) (4,3) (5,4)

This chapter contains information on the following:

- StorageWorks RAID Array 110 Subsystem operations
- Adding logical units (LUNs) for the StorageWorks RAID Array 110 Utility for Novell NetWare
- Configuring LUNs

5.1 StorageWorks RAID Array 110 Subsystem Operations

This section discusses the following:

- StorageWorks RAID Array 110 Subsystem monitoring features
- User monitoring methods

5.1.1 StorageWorks RAID Array 110 Subsystem Monitoring Features

The StorageWorks RAID Array 110 Subsystem, once installed and configured, runs with little user intervention. User intervention is required only if a disk drive fails, if the array configuration needs to be changed, or if any other subsystem component needs to be replaced.

The StorageWorks RAID Array 110 Subsystem provides self-monitoring through two components:

- HSZ10-AA controller
- StorageWorks shelves (the BA350-EA shelf and/or the BA350-SA shelf)

5.1.1.1 Monitoring Through the HSZ10-AA Controller

The HSZ10-AA controller monitors drive operation and LUN status. If certain errors occur on a drive, the controller changes the status of a drive or LUN. A change in drive or LUN status is a mechanism to alert the user to the current condition of the array or drives and that maintenance steps may need to be taken. LUN status is discussed in more detail in LUN Status; drive status is discussed in more detail in Drive Status.

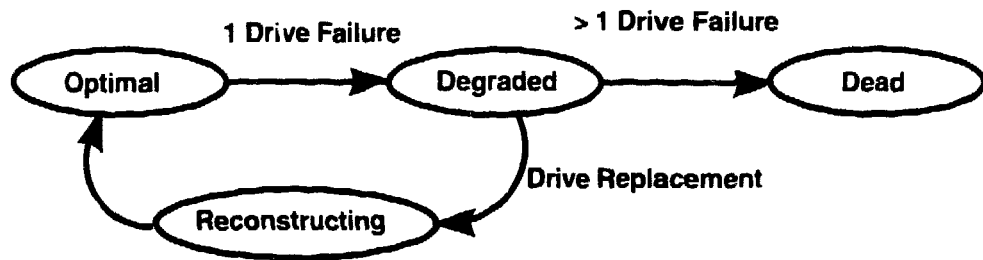
In addition, the HSZ10-AA controller is responsible for setting the fault indicators on the disk drive system building blocks (SBBs). This is done when the controller determines a drive has failed or the user has actively failed a drive through the StorageWorks RAID Array 110 Utilities.

For more information, refer to the *DEC RAID Utilities User's Guide*.

Operations

5.1 StorageWorks RAID Array 110 Subsystem Operations

Figure 5–1 LUN State Diagram



SHR-XR3015-GRA

LUN Status

The LUN or grouping of the drives can be in one of four possible states, as shown in Figure 5–1.

Table 5–1 describes each LUN state or status.

Table 5–1 LUN States

Status	Description
Optimal	The array is operating at an optimal level. This is the condition during normal operation.
Degraded	The LUN is operating in degraded mode. The array is still functioning, but a single drive could have failed. This state is only valid for RAID levels that provide redundancy (RAID 0+1, 1, 3, and 5). In order to return the LUN to optimal, reconstruction of the data must be done.
Dead	The LUN is no longer functioning. This is typical when two or more drives have failed.
Reconstructing	The array controller is currently reconstructing the LUN using good data and parity information. This state is valid only for RAID 0+1, 1, 3, and 5, which provide redundancy.

As detailed in Table 5–1, an *optimal* status is the desired condition for normal operation.

If a drive failure does occur in a redundant configuration (RAID 0+1, 1, 3, or 5) the StorageWorks RAID Array 110 Subsystem is operating in a *degraded* mode. When the array is in degraded mode, you can still continue to use the array.

In the RAID 1 and RAID 0+1 cases, the array controller retrieves the failed drive's data from its mirrored drive whenever you read or write to a LUN in the degraded mode. In the RAID 3 and 5 cases, the array controller determines if the I/O would be directed to the failed drive as data or parity. For the I/O that must be written to a data block on the failed drive, the array controller writes the new data to the operational drives in data blocks. For the I/O that must be written to a parity block on the failed drive, the array controller writes the data to the drives that contain the data blocks for the I/O, and recalculates the parity for the parity block of the I/O.

Whenever you issue an I/O that requires the array controller to read from the failed drive, the array controller recalculates the data and parity blocks from the data and parity on the operational drives.

5.1 StorageWorks RAID Array 110 Subsystem Operations

If a second drive failure occurs, the array is in a *dead* state. The data now is no longer valid. For this reason, it is crucial to replace a drive that has failed right away before a second drive failure occurs.

The process of bringing a degraded LUN back to optimal condition is known as *reconstruction*. Reconstruction is a valid state only for redundant RAID configurations. Data is recreated on the replaced drive using data and parity from the other drives.

Drive Status

Similar to LUN status, drives states are also classified according to status.

Table 5–2 lists the status classifications for drives.

Table 5–2 Drive Status

Status	Description
Optimal	The drive is operating at optimal level.
Warning	The drive has been put into a warning status by the controller as the result of a read or write error. The severity of this status depends to some extent on the RAID level of the LUN. In all cases, the drive is still usable but should be replaced as soon as possible.
Failed	The drive was failed by the array controller or by using a StorageWorks RAID Array 110 utility. The drive must be replaced.
Spare	The drive is connected to an array but not configured into a LUN.
Mismatch	The array controller sensed that the drive has different parameters or configuration information than it expected. This is typical if the user attempted to replace a drive of different capacity into a LUN.

An *optimal* drive state is the desired condition for normal array operation.

A drive is in a *warning* condition when an error occurs that may require drive replacement, but does not affect the reliability of the data on the drive. Conditions that cause the drive to be put into a warning state are:

- Unrecoverable read errors
- Read failures due to the drive being powered-off or transient error

A drive in *warning* state is still used by the controller since it may prove to be available later. In the case of a read error, the data is obtained using redundant information on the other drives. The *warning* condition provides an early warning to the user that drive degradation may be occurring. A drive in a warning state should be replaced with the same urgency as a drive failure.

A drive is marked as *failed* by the array controller when an error occurs that leaves the consistency of the user data and redundant information in a questionable state after using all available recovery actions. When the array determines that a drive has failed, it will not be used again. Conditions that can cause a drive to be marked as failed include the following:

- Unrecoverable errors were made during a write on a data or redundant information drive.
- An error restoring data to a drive after the automatic reallocation of a logical block occurred.
- Unrecoverable errors were made during the read portion of a read/modify/write operation (RAID 5).

5.1 StorageWorks RAID Array 110 Subsystem Operations

- The user marked a drive as failed using a SCSI command or StorageWorks RAID Array 110 utility.
- An error was reported during disk array format sequence.

5.1.1.2 Monitoring Through the StorageWorks Shelf

The StorageWorks storage shelves monitor blowers and power. Monitoring is done by means of a signal in the backplane. Two LED indicators are available on the power SBBs for status. A fault condition indicates one or more of the following problems:

- Power supply failure
- Blower problem
- Input power problem

Refer to Section 8.4 for instructions in determining a power supply or blower fault condition.

5.1.2 User Monitoring Methods

There are two methods for monitoring the operation of the StorageWorks RAID Array 110 Subsystem:

- StorageWorks RAID Array 110 utilities
- LED indicators

5.1.2.1 Monitoring Operation Using the StorageWorks RAID Array 110 Utilities

The StorageWorks RAID Array 110 utilities are key tools to maintaining a disk array. They provide the user with the ability to monitor, configure, repair, and maintain the StorageWorks RAID Array 110 Subsystem. Specific utility functionality is dependent on each operating system since each environment has different needs or capabilities. However, there are some common concepts and functions that are needed across all environments. The following is a generic list of functionality available:

- Check LUN and drive status.
- Restore the LUN (RAID 0+1, 1, 3, and 5) after a drive failure.
- Check and repair array parity on LUNs.
- Configure, reconfigure, and modify LUNs.
- Change default array parameters.
- Change the array configuration parameters (scheduled parity time, parity file name, and so forth).
- Format a logical array.
- Download controller software.
- Check error logs.

A comprehensive description of the StorageWorks RAID Array 110 utilities is beyond the scope of this document; therefore, only an overview of operations is provided here. For more detailed information, refer to the *DEC RAID Utilities User's Guide*.

5.1 StorageWorks RAID Array 110 Subsystem Operations

5.1.2.2 Monitoring Operation Using LED Indicators

The StorageWorks RAID Array 110 Subsystem also provides visual indicators for maintaining array operation. The HSZ10-AA controller, shelf, power supply, and disk status can be monitored using the LEDs.

There are three sets of LED indicators on the subsystem:


- HSZ10-AA controller LEDs
- Power supply SBB LEDs
- Drive SBB LEDs

HSZ10-AA Controller LED Indicators

There is a single bank of eight LEDs on the HSZ10-AA controller. These LEDs convey status information about the boot phase and possible errors. See Figure 1-4 for the location of the LED display.

Once all the controller diagnostics have passed and the controller is ready to begin uploading array software from the disk, the **Heartbeat LED** begins beating. See Figure 5-2.

Figure 5-2 LED Indicators



The Heartbeat LED, the bottommost LED, beats once per second. During the upload process, the upper-half of the LEDs displays a binary 5 while the Heartbeat LED beats rhythmically. After the array software has been uploaded and is running, the firmware clears the upper LEDs. Thus, as far as the firmware is concerned, the normal operating mode of any uploaded software has the Heartbeat LED beating and all other LEDs are off.

Table 5-3 shows the LED codes displayed by the HSZ10-AA controller. Solid codes indicate a boot phase and are considered normal conditions. Cyclic or flashing codes indicate an error condition. Figure 5-2 shows the layout of the LED indicators.

Operations

5.1 StorageWorks RAID Array 110 Subsystem Operations

Table 5–3 Summary of HSZ10-AA Controller LED Codes

Code MSB/LSB	S/F/C	State	Description
00	Solid	None	LEDs are all off, used to flash a value.
01	Solid	Normal	In between boot phases.
02	Solid	Normal	Boot scratch pad memory is being tested.
03	Solid	Normal	ROM is being searched for partitions.
04	Solid	Normal	ROM partitions are being validated.
05	Solid	Normal	RAM is being searched for partitions.
06	Solid	Normal	ROM application partition has been called.
0F	Solid	Debug	Debug mode has been entered.
22	Solid	Normal	Microprocessor (68020) diagnostics in progress.
22	Cyclic	Error	Microprocessor (68020) diagnostics failed.
31	Solid	Normal	EPROM diagnostics in progress.
31	Cyclic	Error	EPROM diagnostics failed.
33	Solid	Normal	EEPROM (58C65) diagnostics in progress.
33	Cyclic	Error	EEPROM (58C65) diagnostics failed.
34	Solid	Normal	Processor SRAM diagnostics in progress.
34	Cyclic	Error	Processor SRAM diagnostics failed.
36	Solid	Normal	Processor DRAM diagnostics in progress.
36	Cyclic	Error	Processor DRAM diagnostics failed.
53	Solid	Normal	Multi-Func Peripheral chip (68901) diagnostics in progress.
53	Cyclic	Error	Multi-Func Peripheral chip (68901) diagnostics failed.
65	Solid	Normal	Host SCSI channel (53C916) diagnostics in progress.
65	Cyclic	Error	Host SCSI channel (53C916) diagnostics failed.
68	Solid	Normal	SCSI Data Path chip (53C920) diagnostics in progress.
68	Cyclic	Error	SCSI Data Path chip (53C920) diagnostics failed.
6A	Solid	Normal	Drive SCSI channel (53C96) diagnostics in progress.
6A 0X	Cyclic	Error	Drive SCSI channel (53C96) diagnostics failed.
80	Flashing	Normal	Heartbeat LED.
88	Solid	Normal	SCSI channel data turn-around diagnostics in progress.
88 XX	Cyclic	Error	SCSI channel data turn-around diagnostics failed.

(continued on next page)

Table 5-3 (Cont.) Summary of HSZ10-AA Controller LED Codes


Code MSB/LSB	S/F/C	State	Description
89	Solid	Normal	Subsystem turn-around diagnostics in progress.
89	Cyclic	Error	Subsystem turn-around diagnostics failed.
A0 XX	Cyclic	Error	General fatal error.
F8 XX	Cyclic	Error	General exception error.
F9 XX	Cyclic	Error	Address error.
FA XX	Cyclic	Error	Instruction error.
FB XX	Cyclic	Error	Arithmetic error.
FC XX	Cyclic	Error	Privilege error.
FE	Solid	Normal	Board is in passive mode.
FF	Solid	Normal	Board is held in a hardware reset state failed.

Power Supply LEDs

Shelf and power supply status are displayed on the power supply LEDs shown in Figure 5-3. The upper LED displays the status of the shelf and the lower LED displays the status of the supply.

Figure 5-3 Power Supply Status LEDs

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



CXO-3613B-PH

- When the upper LED is on, both the blowers and the power supplies are functioning properly.
- When the upper LED is off, a fault condition exists.

Operations

5.1 StorageWorks RAID Array 110 Subsystem Operations

For a detailed explanation of the power supply LED codes, refer to Tables 5-4 and 5-5.

Table 5-4 Shelf and Single Power Supply Status LEDs

Status LED	Status	Indication
Shelf (upper) PS (lower)	On On	NORMAL Status System is operating normally. There are no shelf or power supply faults.
Shelf (upper) PS (lower)	Off On	FAULT Status There is a shelf fault; there is no power supply fault. Replace the blower.
Shelf (upper) PS (lower)	Off Off	FAULT Status Shelf and power supply fault. Refer to Section 8.4.

Table 5-5 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. There are no shelf or power supply faults.
Shelf (upper) PS (lower)	Off On	Off On	FAULT Status There is a shelf fault. There is no power supply fault. Replace the blower.
Shelf (upper) PS (lower)	Off On	Off Off	FAULT Status PS1 operational. Replace PS2.
Shelf fault PS (lower)	Off Off	Off On	FAULT Status PS2 operational. Replace PS1.
Shelf (upper) PS (lower)	Off Off	Off Off	FAULT Status Possible PS1 and PS2 fault or possible input power problem.


5.1 StorageWorks RAID Array 110 Subsystem Operations

Drive SBB LEDs

Each drive SBB has two LED indicators, as shown in Figure 5-4, which display the SBB's status. These LEDs have three states as follows:

- On
- Off
- Flashing

Figure 5-4 Shelf Status LEDs



CXO-3614A-PH

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

The actual orientation of the LEDs depends upon your configuration. The LEDs shown in Figure 5–4 are described as follows:

- The top-left LED (green) is the device activity LED and is On or Flashing when the SBB is active.
- The bottom LED (amber) is the drive SBB fault LED and defines the error condition, as indicated by the state of the LED – either On or Flashing.

Note

During initial power-on of the StorageWorks RAID Array 110 Subsystem, the bottom LED (amber) may be lit. This does not indicate a fault condition. A bus reset is required from the host to clear these lights.

Table 5–6 Drive SBB Status LEDs

LED	Status	Indication
SBB activity	On	SBB is operating normally.
SBB fault	Off	
SBB activity	On	Fault state. SBB is probably hung up. It is recommended that you replace the SBB.
SBB fault	On	
SBB activity	Off	Fault state. SBB is inactive and spun down. It is recommended that you replace the SBB.
SBB fault	On	
SBB activity	On	Fault state. SBB is active and drive is being spun down because of fault.
SBB fault	Flashing	

If a fault state exists, refer to Section 8.3 for instructions in replacing a drive SBB.

5.2 Adding LUNs to NetWare

This section contains information on the following:

- Using the Install command
- Adding LUNs to existing volumes
- Performance notes

5.2.1 Using the Install Command

After you configure a LUN, you use the NetWare Install command to create partitions and volumes so that your system can access it. You use this command in the same way as for a standard hard disk, using the instructions given in the NetWare documentation. For best results, select 32K or 64K as the block size of any volumes on the LUN.

5.2.2 Adding LUNs to Existing Volumes

You can add an array LUN to an existing volume just as you can add any other disk storage device. However, remember that the data reliability of a combined volume is only as great as the reliability of its weakest part.

That is, if you combine a RAID 5 LUN and a RAID 0 LUN into the same volume, the resulting volume has the data redundancy protection of a RAID 0 LUN (that is, none). Even though the RAID 5 portion of the volume has data redundancy, if the RAID 0 portion fails, the whole volume may become unusable and all data on the volume lost. A RAID 0+1, 1, and RAID 5 volume would, however, offer data redundancy.

In general you should not combine redundant and nonredundant disk storage systems in the same volume.

5.2.3 Performance Notes

Array performance depends on many different factors, including RAID level, block size, I/O block size, and so on. Here are a few suggestions that may improve performance on your array:

- Make your segment size at least the same size as your average I/O size. For example, if your average I/O size is 16 Kbytes (32 blocks), change your segment size (using the Configuration utility) to 32 blocks as well.
- Always use a segment 0 size of 0.
- If data redundancy is not important, try using RAID 0 instead of RAID 1, 0+1, or 5. In some circumstances this improves performance. However, all data is lost if a single drive fails.

5.3 Configuring LUNs

5.3.1 Array Configuration Summary

The drives in the StorageWorks RAID Array 110 Subsystem are preconfigured into LUNs at the factory. If this configuration does not meet your needs, you can change that configuration by doing one or more of the following:

- Modify LUN parameters other than RAID level and Drive Map (segment size, segment 0 size, and so on); see Section 5.3.4.
- Modify RAID Level/Drive Map parameters; see Section 5.3.7.
- Create a new LUN from spare drives; see Section 5.3.5.
- Create a new LUN from space on an existing drive group; see Section 5.3.6.
- Delete a LUN; see Section 8.5.

In addition, at any time after the initial installation of your array, you can reconfigure a LUN. However, any time you reconfigure, you must first back up any data on the LUN to avoid destroying the existing data on the unit.

If your operating system is MS-DOS or SCO UNIX, use the StorageWorks RAID Array 110 Utility for MS-DOS for configuring your LUNs. Refer to Table 5-7 for a summary of the steps involved in configuring your LUN. In NetWare, use the StorageWorks RAID 110 Utility for Novell NetWare (refer to Table 5-8).

The specific steps you must take depend on two factors:

- Whether you want to use the factory-set configuration (if any)
- Whether you want to put the operating system on the LUN (for NetWare, this refers to whether the LUN will contain the SYS volume).

If your array is preconfigured, and you want to change the configuration, treat the LUNs as being NOT configured when using the tables.

Table 5-7 StorageWorks RAID Array 110 Utility for MS-DOS LUN Configuration Tasks

LUN Configured?¹	LUN Contains the Operating System Volume?	Steps To Complete
Yes ¹	Yes	1. Install the operating system on the LUN.
Yes ¹	No	1. No action required.
No	Yes	1. If any data is on the LUN, back up the data. 2. Boot the StorageWorks RAID Array 110 Utility for MS-DOS diskette (see Section 4.1). 3. Configure or reconfigure the LUN (See the procedures in this section). 4. Install the operating system on the LUN. 5. If necessary, copy the data from the backup media to the LUN.
No	No	1. If any data is on the LUN, back up the data. 2. Shut down the system. 3. Boot the StorageWorks RAID Array 110 Utility for MS-DOS diskette (see Section 4.1). 4. Configure or reconfigure the LUN. (See the procedures in this section.) 5. Reboot the system. 6. If necessary, recopy the data back onto the LUN.

¹If the LUN is already configured, but you want to change that configuration, treat the LUN as if it is not already configured when using this table.

Table 5-8 StorageWorks RAID Array 110 Utility for Novell NetWare LUN Configuration Tasks

LUN Configured?¹	LUN Contains the Operating System Volume?	Steps To Complete
Yes ¹	Yes	1. Create the SYS volume on the LUN (Section 5.2). 2. Install the operating system on the LUN (with the NetWare Server and Install commands). 3. Install RAID Manager (see Section 3.2).
Yes ¹	No	1. If necessary, install RAID Manager (see Section 3.2). 2. Create partitions and volumes on the LUN with the NetWare Install command (see Section 5.2).

¹If the LUN is already configured, but you want to change that configuration, treat the LUN as if it is not already configured when using this table.

(continued on next page)

Table 5-8 (Cont.) StorageWorks RAID Array 110 Utility for Novell NetWare LUN Configuration Tasks

LUN Configured?¹	LUN Contains the Operating System Volume?	Steps To Complete
No	Yes	<ol style="list-style-type: none"> 1. If there is any data on the LUN, back up the data (using file-by-file backup). 2. Boot the StorageWorks RAID Array 110 Utility for MS-DOS diskette (see Section 4.1). 3. Configure or reconfigure the LUN (see the procedures in this section). 4. Install the operating system on the LUN (with the NetWare Server and Install commands). 5. Install RAID Manager (see Section 3.2). 6. If necessary, copy the data from the backup media to the LUN.
No	No	<ol style="list-style-type: none"> 1. If necessary, install RAID Manager (see Section 3.2). 2. If there is any data on the LUN, back up the data (using file-by-file backup). 3. If any volumes on the LUN are mounted, dismount them (with the NetWare Dismount command). 4. Configure or reconfigure the LUN using the StorageWorks RAID Array 110 Utility for Novell NetWare. 5. Create partitions and volumes on the LUN (see Section 5.2). 6. If necessary, recopy the data back to the LUN.

¹If the LUN is already configured, but you want to change that configuration, treat the LUN as if it is not already configured when using this table.

Table 5-9 StorageWorks RAID Array 110 Utility for OpenVMS VAX LUN Configuration Tasks

LUN Configured? ¹	LUN Contains the Operating System Volume?	Steps To Complete
Yes ¹	Yes	<ol style="list-style-type: none"> 1. Install the OpenVMS VAX operating system Version 5.5-2 on the LUN. 2. Install the DEC SCSI Tagged Command Queuing (TCQ) Drivers Update Kit (refer to the <i>DEC SCSI Tagged Command Queuing (TCQ) Driver for OpenVMS VAX Release Notes and Installation Guide</i>). 3. Install the StorageWorks RAID Array 110 Utility for OpenVMS VAX (see the <i>DEC RAID OpenVMS VAX Utility Release Notes and Installation Guide</i>).
Yes ¹	No	<ol style="list-style-type: none"> 1. If necessary, install the StorageWorks RAID Array 110 Utility for OpenVMS VAX (see the <i>DEC RAID OpenVMS VAX Utility Release Notes and Installation Guide</i>).
No	Yes	<ol style="list-style-type: none"> 1. If there is any data on the LUN, back up the data. 2. Boot a system disk that has OpenVMS VAX 5.5-2 with the DEC SCSI TCQ Drivers update, and the StorageWorks RAID Array 110 Utility for OpenVMS VAX installed on it. 3. Configure the LUN using the Array Configuration Editor (ACE) (see Section 4.4). 4. Install OpenVMS VAX Version 5.5-2 on the LUN. 5. Install the DEC SCSI TCQ Drivers update. 6. Install the StorageWorks RAID Array 110 Utility for OpenVMS VAX.
No	No	<ol style="list-style-type: none"> 1. If there is any data on the LUN, back up the data. 2. Boot a system disk that has OpenVMS VAX 5.5-2 with the DEC SCSI TCQ Drivers update, and the StorageWorks RAID Array 110 Utility for OpenVMS VAX installed on it. 3. Configure the LUN using the Array Configuration Editor (ACE) (see Section 4.4). 4. If necessary, restore the data back to the LUN.

¹If the LUN is already configured, but you want to change that configuration, treat the LUN as if it is not already configured when using this table

5.3.2 LUN Parameters

Each LUN has a set of parameters that determine how data is stored on it. Each LUN can have different parameters, with the following exception: All LUN in the same drive group must have the same RAID Level and Drive Map (that is, they must contain the same drives, with no overlap).

Table 5–10 table summarizes the parameters. The parameters are explained in greater detail in the *StorageWorks RAID Array 110 Subsystem Service Guide*. Sections 5.3.2.1 to 5.3.2.8 explain the parameters in more detail.

Table 5–10 Logical Unit Parameters

Parameter	Meaning	Destroys Data When Changed?	See Section:
RAID Level	Determines how data is stored on the LUN, and if there is data redundancy.	Yes ¹	5.3.2.1
Drive Map	Determines what drives make up the LUN.	Yes ¹	5.3.2.2
Logical Block Size	The logical block size (in bytes) used by this LUN.	N/A	5.3.2.3
Logical Unit Size	Sets the size of the LUN.	Yes	5.3.2.4
Segment Size	Determines the amount of data written to a single drive in the LUN before the controller writes data on the next drive.	Yes	5.3.2.5
Segment Zero Size	Sets the size of the first segment in the array.	Yes	5.3.2.6
Delay Interval	The amount of time between reconstruction operations.	No	5.3.2.7
Blocks Per Delay Interval	The number of blocks reconstructed in one reconstruction operation.	No	5.3.2.8

¹Changing RAID Level or Drive Map forces you to change those parameters for all the LUNs in that drive group.

5.3.2.1 RAID Level

The RAID Level parameter determines how data is stored on the LUN. Data is striped or mirrored. The RAID Level parameter also determines if data redundancy has occurred. RAID 0+1, 1, and 5 offer data redundancy; RAID 0 does not.

Caution

Changing the RAID Level parameter deletes any data on the LUN. Use this parameter only after doing a file backup.

The RAID Level also determines the number of drives that can be included in the LUN and the maximum size of that unit. Refer to Table 5–11.

If you want to change the RAID level for an existing LUN, you must first delete the unit and then recreate it. In addition, if there are other LUNs in the same drive group, you must delete all of them and change their RAID level as well. Similarly, if you want to create a new LUN from space in an existing drive group, that LUN must have the same RAID level as the other LUNs in the drive group.

The RAID levels offered by the StorageWorks RAID Array 110 Subsystem are 0, 1, and 5. The levels you select depend on your storage and performance needs.

5.3.2.2 Drive Map

The Drive Map parameter defines the drives included in the LUN. Individual drives are identified by Channel number and SCSI ID, which is determined by the physical location of the drive in the array subsystem. Each drive rank in an array has the same SCSI ID, and each rank contains channels 1 through 5.

Caution

Changing the Drive Map parameter deletes any data on the LUN. Use this parameter only after doing a file backup.

In ACE, the drive matrix shows the drives in a logical arrangement, with SCSI ID 0 first, then SCSI ID 1, and so on. This logical arrangement is not the same as the physical arrangement of the drives in the StorageWorks RAID Array 110 Subsystem. Refer to the *StorageWorks RAID Array 110 Subsystem User's Guide* for further details about the StorageWorks RAID Array 110 Subsystem physical configuration.

The RAID level of the LUN sets some restrictions on drive selection, as shown in Table 5-11.

Table 5-11 RAID Level/Drive Selection

RAID Level	Drive Map Restrictions
0	Number of drives allowed per LUN = 1 to 10.
1	Number of drives allowed per LUN = 2 to 10. Must specify an even number of drives. The mirrored pair is created by grouping the first and second drive you enter, third and fourth, and so forth. Drives in a mirrored pair cannot be on the same channel.
3	Number of drives allowed per LUN = 3 or 5.
5	Number of drives allowed per LUN = 3 to 5. Each drive must be on a separate channel.

Note

It is recommended that you use the maximum number of drives when you create a LUN. Remember that a single-rank array has only five drives.

From the table, you can see that defining a RAID 5 LUN as containing drives (5,3), (4,3), (3,3), (2,3), and (2,0) is not legal, as the last two drives are on the same channel (channel 2). Selecting (5,3), (4,3), (3,3), (2,3), and (1,3) is legal, because although the drives have the same SCSI ID, they are on different channels. Similarly, defining a RAID 1 LUN as (5,3), (5,0), (4,3), and (4,0) is not legal, as the first and second drives entered must be a mirrored pair and mirrored pairs cannot be on the same channel. The third and fourth drives have the same problem. Entering (5,3), (4,0), (4,3), and (5,0) is legal because, even though it contains the same drives as the previous example, the mirrored pairs are on different channels.

Operations

5.3 Configuring LUNs

If you want to change the Drive Map for an existing LUN, you must first delete the unit and then recreate it with the new drives. In addition, if there are other LUNs on the same drive group, you must delete all of them and recreate them as well. Similarly, if you want to create a new LUN from space in an existing drive group, that LUN must include the same drives as the other LUN in the drive group.

Within the restrictions above, the drives you select for a LUN depend on your storage and performance needs. In general, always use the maximum number of drives in a LUN.

5.3.2.3 Logical Block Size

The logical block size parameter defines the block size in bytes as seen by the LUN. Currently, the logical block size is fixed at 512 bytes and cannot be changed.

5.3.2.4 Logical Unit Size

This parameter sets the size of the LUN. Size is determined by the RAID level, number of drives, and the amount of space allocated on those drives. Unallocated space on a set of drives (a drive group) can be used to create other LUNs. In general, however, you should allocate all available space for a LUN, and create one LUN per drive group (multiple LUNs per group are useful for operating systems with size restrictions on disk storage).

Note

You may create any size LUN—NetWare does not have a size limitation. MS-DOS, however, does have a limit. It does not display a LUN size of greater than 2 GB. This means that if you create a LUN with a size greater than 2 GB, and then uses an MS-DOS command to display the size, MS-DOS only indicates that you have 2 GB space available.

Caution

Changing the LUN size parameter deletes any data on the LUN. Use this parameter only after doing a file backup.

Assign all available space to the LUN.

5.3.2.5 Segment Size

A segment is the amount of data written on a single drive in the LUN before the controller continues writing the data on the next drive in the LUN. For example, if the segment size of a RAID 0 LUN is 64 blocks, the controller writes 64 blocks of data on drive 1, the next 64 blocks of data on drive 2, the next on drive 3, and so on.

Caution

Changing the segment size parameter deletes any data on the LUN. Use this parameter only after doing a file backup.

The recommended segment size is 512 blocks. ACE accepts a segment size as small as 3 blocks and as large as 65,535 blocks (however, do not use an odd number of blocks if you want to enhance array performance).

5.3.2.6 Segment Zero Size

The first segment in a LUN is segment 0. Segment 0 is reserved for future development. The recommended segment 0 size is 0. Any other value may degrade array performance.

Caution

Changing the segment size parameter deletes any data on the LUN. Use this parameter only after doing a file-by-file backup.

5.3.2.7 Delay Interval

Note

The Delay Interval parameter only applies to RAID 0+1, 1, and 5 LUNs. The data on a RAID 0 array cannot be reconstructed.

During data reconstruction on a RAID 0+1, 1, or 5 LUN (after you replace a drive in a degraded LUN), the array controller divides its time between data reconstruction and regular I/O operations. In this way, you can continue using the LUN while it is being reconstructed. The Delay Interval is the amount of time, in tenths of a second, between reconstruction operations. During this time, the array controller is able to perform normal I/O operations. As the delay interval increases, system I/O performance also increases, but the reconstruction takes longer.

You can change this parameter during reconstruction to affect the reconstruction rate. The Delay Interval value interacts with the Blocks per Delay Interval value to determine the overall rate of reconstruction. See Section 8.7.7 for further information.

5.3.2.8 Blocks Per Delay Interval

Note

The Blocks Per Delay Interval parameter only applies to RAID 0+1, 1, and 5 LUNs. The data on a RAID 0 array cannot be reconstructed.

The Blocks Per Delay Interval parameter affects the amount of data, in blocks, that the array controller reconstructs at a time. The more blocks, the longer the time necessary to reconstruct them. This is time that cannot be used to perform system I/O. Therefore, the larger the Blocks Per Interval value, the more system performance degradation increases.

You can change the Blocks Per Interval value after reconstruction begins to adjust system performance.

The Blocks Per Delay Interval value is a decimal number between 1 and 32768. However, the value interacts with the Delay Interval value to determine the over-all rate of reconstruction. See Section 8.7.7 for further information.

5.3.3 Displaying and Modifying Logical Unit Parameters

This section describes how to display and modify LUN parameters for an existing LUN. The LUN parameters are described in Section 5.3.2.

Caution

Note that changing any parameter on an existing LUN except the Reconstruction Rate parameters destroys all data on that unit. Make sure you back up all data on the LUN before modifying any parameters other than Delay Interval or Blocks Per Delay Interval.

Use the Display/Modify Logical Unit Parameter function when:

- You want to change LUN parameters on an existing LUN (except the RAID level and Drive Map parameters, which use a separate procedure—see Section 5.3.7.)

Note that changing parameters on an existing LUN will usually destroy all data on that unit.

- You want to change the Reconstruction Rate parameters (Delay Interval and Blocks Per Delay Interval). Note that you can change these parameters without destroying data on the LUN.
- You are creating a new LUN (in which case this function is started automatically during the creation procedure).
- You want to display the values for the Logical Block Size, Segment Size, Segment 0 size, and Reconstruction Rate parameters.

5.3.4 Modify/Display Procedure

Perform the following steps to display or modify LUN parameters.

Caution

Changing any LUN parameter except the Delay Interval or Blocks Per Delay Interval parameters destroys all data in the LUN. Make sure you back up your data file-by-file before modifying any of these LUN parameters.

1. Back up any data on the LUN you want to change.
2. In NetWare, unmount any volumes contained on the LUN you want to change if any are mounted (if you are changing the Reconstruction Rate parameters, you do not need to dismount any volumes).
3. If you are going to be modifying parameters in MS-DOS or SCO UNIX, or in NetWare if you are going to be modifying parameters on the LUN that contains any part of your SYS volume, you must use the StorageWorks RAID Array 110 Utility for MS-DOS. Reboot the system using the StorageWorks RAID Array 110 Utility for MS-DOS diskette and continue with this procedure.
4. Start RAID Manager from the console by entering one of the following commands:

In the StorageWorks RAID Array 110 Utility for MS-DOS:

raidmgr

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

load raidmgr

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

\$ ACE device-name

5. Select the *Array Controller Configuration* option on the *Disk Array Services* menu. Check the Disk Array Devices window on the right of the screen to make sure it is displaying the LUN you want to modify.
6. Select the *Array Configuration Editor* option from the menu. Note that you are in the Logical Unit window of the Array Configuration Editor (ACE). If you are not in this window, press the Tab key to switch to the Logical Unit window. Use the Arrow keys to highlight the LUN you want to modify.
7. Press the M key to modify or display the LUN parameters. The following screen is displayed:

LUN 1 Status = Optimal
LUN 1 Drives = (1,2) (2,2) (3,2) (4,2) (5,2)

PARAMETER	VALUE
RAID Level	5
Logical Block Size (bytes)	512
Logical Unit Size (MB)	1600
Segment Size (blocks)	512
Segment Zero Size (blocks)	0
Delay Interval (tenths of a second)	1
Blocks Per Delay Interval	256

Are the values shown OK (y/n/q)?

Note

In NetWare, if you created a very large LUN, and you hold down the Ctrl key and press the Esc key to do an MS-DOS operation while the format is occurring (for example, you try to load Install), it is possible that the screen will lock up until the array controller completes the format. This situation occurs because of the way the NetWare driver handles outstanding I/O commands. Wait until the array controller completes the format before doing an MS-DOS operation.

8. Respond to the *Are the values shown OK (y/n/q)?* prompt as follows:

To do this:	Press this key:
Save the parameters. ¹	Y
Modify the values. ²	N
After you enter N, do the following:	
1. Use the arrow keys to move the cursor to select a value you want to change and enter a new value.	
2. Press the Enter key.	
3. After you have set all the parameters you want to set, use the Down-arrow key or the Enter key to move to the <i>Are the values shown OK</i> prompt, then press:	
— the Y key to save the current LUN parameter values.	
— the Q key to quit without changing the LUN parameters.	
Quit without saving the parameters.	Q

¹The array controller automatically formats the LUN if you modified the LUN size, segment size, or segment zero size. The utility returns to the ACE menu when done.

²**NOTE:** You cannot change the RAID level parameter, or the Logical Block Size. To determine the values you can use for the LUN parameters, see Section 5.3.2.

9. Press the Q key to quit the ACE utility.
10. In the StorageWorks RAID Array 110 Utility for MS-DOS, remove the diskette and reboot the system. In the StorageWorks RAID Array 110 Utility for Novell NetWare, use the NetWare Install command to create partitions and volumes on the modified LUN (see Section 5.2). In the StorageWorks RAID Array 110 Utility for OpenVMS VAX, reboot the operating system before using the device.

You are done with the Display/Modify LUN Parameter procedure.

5.3.5 Creating a LUN from Spare Drives

Use the following procedure to create a LUN from *Spare* drives.

- To create a LUN in MS-DOS or SCO UNIX, or in NetWare to create the LUN that contains your SYS volume, you must use the StorageWorks RAID Array 110 Utility for MS-DOS. Reboot the system using the StorageWorks RAID Array 110 Utility for MS-DOS diskette and continue with this procedure.
- Start RAID Manager from the console by entering one of the following commands:

In the StorageWorks RAID Array 110 Utility for MS-DOS:

```
raidmgr
```

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
load raidmgr
```

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

```
$ ACE device-name
```

3. Select the *Array Controller Configuration* option on the menu. Check the Disk Array Devices window on the right of the screen to make sure it is displaying LUNs attached to the controller you want to create the new LUN on.
4. Select the *Array Configuration Editor* option from the menu. Note that you are in the Logical Unit window of the Array Configuration Editor (ACE). If you are not in this window, press the Tab key to switch to the Logical Unit window. Use the Arrow keys to highlight Drive Group 0 (the drive group containing the spare drives attached to the current array controller).
5. Press the C key to create a LUN. When prompted, specify the RAID level choice. Enter the Channel and ID of each drive you want in the LUN using the guidelines given below. The drives you can configure into a LUN are marked with the status 2 in the Drive Matrix (see the Drive window). The cursor moves as you enter each selection.

Refer to Table 5-11 for considerations when assigning drives.

6. Press the Enter key after entering the drive ID number for the last drive in the LUN. *Is the data OK* prompt appears. Press the Y key to respond. The following screen is displayed:

LUN 1 Status = Optimal

LUN 1 Drives = (1,2) (2,2) (3,2) (4,2) (5,2)

PARAMETER	VALUE
RAID Level	5
Logical Block Size (bytes)	512
Logical Unit Size (MB)	1600
Segment Size (blocks)	512
Segment Zero Size (blocks)	0
Delay Interval (tenths of a second)	1
Blocks Per Delay Interval	256

Are the values shown OK (y/n/q)?

7. Respond to the *Are the values shown OK (y/n/q)?* prompt as follows:

To do this:	Press this key:
Save the parameters. ¹	Y
Modify the values. ²	N

After you enter N, do the following:

1. Use the arrow keys to move the cursor to select a value you want to change and enter a new value.
2. Press the Enter key.
3. After you have set all the parameters you want to set, use the Down-arrow key or the Enter key to move to the *Are the values shown OK* prompt, then press:
 - the Y key to save the current LUN parameter values.
 - the Q key to quit without changing the LUN parameters. This returns you to the ACE menu.

Quit without saving the parameters. Q

¹The array controller automatically formats the LUN. The utility returns to the ACE menu when done.

²**NOTE:** You cannot change the RAID level parameter, or the Logical Block Size. To determine the values you can use for the LUN parameters, see Section 5.3.2.

8. Press the Q key to quit the ACE utility.
9. In the StorageWorks RAID Array 110 Utility for MS-DOS, remove the diskette and reboot the system. In the StorageWorks RAID Array 110 Utility for Novell NetWare, use the NetWare Install command to create partitions and volumes on the modified LUN (see Section 5.2). In the StorageWorks RAID Array 110 Utility for OpenVMS VAX, reboot the operating system before using the device.

You are done with this procedure.

5.3.6 Creating a LUN from an Existing Drive Group

Use the following procedure to create a LUN from space in an existing drive group. Note that the new LUN must use the same drives and the same RAID level as the other drives in the group.

1. To create a LUN in MS-DOS or SCO UNIX, or in NetWare to create the LUN that contains your SYS volume, you must use the StorageWorks RAID Array 110 Utility for MS-DOS. Reboot the system using the StorageWorks RAID Array 110 Utility for MS-DOS diskette and continue with this procedure.

2. Start RAID Manager from the console by entering one of the following commands:

In the StorageWorks RAID Array 110 Utility for MS-DOS:

```
raidmgr
```

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
load raidmgr
```

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

```
$ ACE device-name
```

3. Select the *Array Controller Configuration* option on the menu. Check the *Disk Array Devices* window on the right of the screen to make sure it is displaying the LUNs in the drive group in which you want to create the new LUN.
4. Select the *Array Configuration Editor* option from the menu. Note that you are in the Logical Unit window of ACE. If you are not in this window, press the Tab key to switch to the Logical Unit window. Use the arrow keys to highlight the space remaining in the drive group in which you want to create the new LUN.

Note that when the desired drive group is highlighted, all the drives that are in that drive group are highlighted in the Drive Matrix, and will automatically be configured into the LUN you create. The new LUN will also have the same RAID level as the other LUNs in the group.

5. Press the C key to create a LUN. The following screen (or one like it—parameter values may vary) is displayed:

```
LUN 1 Status = Optimal
LUN 1 Drives = (1,2) (2,2) (3,2) (4,2) (5,2)

PARAMETER                                VALUE
RAID Level                               5
Logical Block Size (bytes)               512
Logical Unit Size (MB)                   1600
Segment Size (blocks)                     512
Segment Zero Size (blocks)                0
Delay Interval (tenths of a second)       1
Blocks Per Delay Interval                256

Are the values shown OK (y/n/q)?
```

Operations

5.3 Configuring LUNs

6. Respond to the *Are the values shown OK (y/n/q)?* prompt as follows:

To do this:	Press this key:
Save the parameters. ¹	Y
Modify the values. ²	N
After you enter N, do the following:	
1. Use the arrow keys to move the cursor to select a value you want to change and enter a new value.	
2. Press the Enter key.	
3. After you have set all the parameters you want to set, use the Down-arrow key or the Enter key to move to the <i>Are the values shown OK</i> prompt, then press:	
— the Y key to save the current LUN parameter values.	
— the Q key to quit without changing the LUN parameters. This utility returns you to the ACE menu.	
Quit without saving the parameters.	Q

¹The array controller automatically formats the LUN. The utility returns to the ACE menu when done.

²**NOTE:** You cannot change the RAID level parameter, or the Logical Block Size. To determine the values you can use for the LUN parameters, see Section 5.3.2.

7. Press the Q key to quit the ACE utility.
8. In the StorageWorks RAID Array 110 Utility for MS-DOS, remove the diskette and reboot the system. In StorageWorks RAID Array 110 Utility for Novell NetWare, use the NetWare Install command to create partitions and volumes on the modified LUN (see Section 5.2). In the StorageWorks RAID Array 110 Utility for OpenVMS VAX, reboot the operating system before using the device.

You are done with this procedure.

5.3.7 Changing RAID Level/Drives on an Existing LUN

If you want to change the RAID level or drive grouping of an existing LUN, you must first delete the current LUN and recreate it with a new RAID level or new drives. If the LUN is part of a drive group, change the RAID level/drives on all the LUNs in that group. This means that you must delete all the LUNs in the group.

Use the following procedure to change the RAID level or drive group of an existing LUN.

1. Back up the data on all the LUNs in the drive group.
2. In NetWare, unmount any volumes contained on the LUNs in the drive group you want to change if any are mounted.
3. To create a LUN in MS-DOS or SCO UNIX, or in NetWare to create the LUN that contains your SYS volume, you must use the StorageWorks RAID Array 110 Utility for MS-DOS. Reboot the system using the StorageWorks RAID Array 110 Utility for MS-DOS diskette and continue with this procedure.
4. Start RAID Manager from the console by entering one of the following commands:

In the StorageWorks RAID Array 110 Utility for MS-DOS:

```
raidmgr
```

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
load raidmgr
```

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

```
$ ACE device-name
```

5. Select the *Array Controller Configuration* option on the menu. Check the Disk Array Devices window on the right of the screen to make sure it is displaying the LUNs you want to modify.
6. Select the *Array Configuration Editor* option from the menu. Note that you are in the Logical Unit window of ACE. If you are not in this window, press the Tab key to switch to the Logical Unit window.
7. Use the Arrow keys to highlight the LUN you want to modify.
Press:
 - The D key to delete the LUN.
 - The C key to confirm that you want to delete the LUN.
 - The Q key if you do not want to delete the unit.

Repeat this step for each LUN in the drive group. When you have deleted all the LUNs in the group, the status of the highlighted drives in the Drive Matrix change to *Spare (2)*.

8. Refer to Section 5.3.6 to create the LUN(s) you want with the new RAID levels and/or drive groupings.

Troubleshooting

This chapter contains troubleshooting information to correct problems that may be easy to fix. It also directs you to the appropriate documentation for additional troubleshooting information if needed.

This chapter discusses the following:

- Before you begin troubleshooting
- Using the troubleshooting table
- If you have expanded your StorageWorks RAID Array 110 Subsystem

6.1 Before You Begin Troubleshooting

To determine where the problem with your StorageWorks RAID Array 110 Subsystem exists, follow these steps:

1. Turn off the StorageWorks RAID Array 110 Subsystem.
2. Turn off the host system.
3. Check to see that the cables are correctly connected including the following:
 - 68-pin SCSI “P” Cable connected to the HSZ10-AA controller
 - Terminator
 - SCSI cable connected to the host system
 - 50-pin SCSI cables connected to BA350-SA shelves (if applicable)
4. Turn the StorageWorks RAID Array 110 Subsystem back on.
5. Verify that all the drive SBB LED indicators flash on initial power on.
6. Verify that the two LED indicators on the power supplies are lit.
7. Verify that the HSZ10-AA controller LED indicators flash and the controller begins its diagnostics.

6.2 Using the Troubleshooting Table

When the StorageWorks RAID Array 110 Subsystem does not operate correctly, use the information in this section to help diagnose the problem.

The troubleshooting techniques described do not identify all possible problems with the subsystem, nor do the corrective actions suggested remedy all problems.

To use Table 6–1, follow these steps:

1. Note the symptoms of the problem displayed by your StorageWorks RAID Array 110 Subsystem.
2. Check the *Symptom* column in Table 6–1 for a match.

Troubleshooting

6.2 Using the Troubleshooting Table

3. Check the conditions for that symptom in the *Possible Cause* column. If more than one possible cause is given, check all of the possible causes in the order listed.
4. Follow the advice in the *Corrective Action* column.

Table 6–1 Troubleshooting Systems Problems

Symptom	Possible Cause	Corrective Action
Drive SBB Fault Light is on.	Drive has failed.	Replace drive using instructions in Chapter 5, Section 8.3.
Drive SBB Fault and Activity Lights are on.	Drive has failed or is hung.	Replace drive using instructions in Chapter 5, Section 8.3.
Drive SBB Fault Light is Flashing.	Drive has been failed and is spinning down.	Replace drive using instructions in Chapter 5, Section 8.3.
Replaced drive has not spun up.	Drive not seen by HSZ10-AA controller.	Remove drive SBB, wait 10 seconds, reinsert drive SBB.
Power Supply SBB shelf status light is off.	Shelf Fault.	Refer to Chapter 5, Section 5.1.2 and Section 8.4, for description of fault condition and resolution.
Both lights are off on Power Supply.	Input power problem.	Check for proper connection of input power.
	Shelf and Power Supply Fault.	Refer to Chapter 5, Section 5.1.2 and Section 8.4, for description of fault condition and resolution.
HSZ10-AA controller LED indicators are off.	Controller is not properly installed.	Remove and reseal controller using the guidelines in the <i>HSZ10-AA Controller Site Preparation Guide</i> .
	HSZ10-AA controller has failed.	Replace controller using guidelines in the <i>HSZ10-AA Controller Site Preparation Guide</i> .
StorageWorks RAID Array 110 Subsystem is not seen by host.	SCSI cable is not connected.	Check the SCSI cable both at the host and StorageWorks RAID Array 110 Subsystem ends for proper connection.

(continued on next page)

Table 6-1 (Cont.) Troubleshooting Systems Problems

Symptom	Possible Cause	Corrective Action
	Incorrect termination.	Check that both ends of the SCSI bus are terminated correctly. Verify that proper termination scheme is being used (see Chapter 3).
	Duplicate SCSI IDs on bus.	Check SCSI ID settings on all devices connected to the bus for duplication. Run bus scan console diagnostics if they are available on your host system.
	Defective HSZ10-AA controller.	Refer to Chapter 5, Section 5.1.2, for a description of the possible LED status. Determine whether the controller has failed. Replace the controller if needed.
Software does not boot from the StorageWorks RAID Array 110 Subsystem.	A problem exists with the StorageWorks RAID Array 110 Subsystem	Use any available system console diagnostics or the StorageWorks RAID 110 Array utilities to test StorageWorks RAID Array 110 Subsystem.
	A problem exists with the software installed on the StorageWorks RAID Array 110 Subsystem.	Refer to the <i>DEC RAID Utilities User's Guide</i> for help.

6.3 If You Have Expanded Your StorageWorks RAID Array 110 Subsystem

If you have attempted an expansion of your StorageWorks RAID Array 110 Subsystem, follow these steps to determine the problem:

1. Turn off the StorageWorks RAID Array 110 Subsystem.
2. Turn off the host system.
3. Check that the reconfiguration of the shelves has been done correctly according to the shelf reconfiguration diagrams in Chapter 10:
 - Check that jumpers are placed in proper locations.
 - Check that terminators are placed in proper locations.
4. Check to see that the cables are correctly connected, including the following:
 - 68-pin SCSI "P" Cable connected to the HSZ10-AA controller
 - Terminator
 - SCSI cable connected to the host system

6.3 If You Have Expanded Your StorageWorks RAID Array 110 Subsystem

- 50-pin SCSI cables connected to BA350-SA shelves (if applicable)
5. Check that the drive movement and drive addition have been done correctly. Refer to Chapter 10.
 6. Turn the StorageWorks RAID Array 110 Subsystem back on.
 7. Verify that all the drive SBB LED indicators flash on initial power on.
 8. Verify that the two LED indicators on the power supplies are lit.
 9. Verify that the HSZ10-AA controller LED indicators flash and the controller begins its diagnostics.
 10. Use the StorageWorks RAID Array 110 utilities to determine the status of the LUN and drive.

Error Information and Symptoms Tables

This chapter contains information on the following:

- System message tables
- Symptoms and status tables
- Error reporting under OpenVMS VAX

7.1 System Message Tables

This section lists the tables of system messages available in this guide. Table 7-1 describes the contents of the tables, and indicates where to locate them.

Table 7-1 System Messages and How to Find Them

Table	Description of Contents	Location
Table 3-4, Installation Error Messages	Lists the error messages, their probable cause, and action(s) to take.	Chapter 3
Table 5-3, Summary of HSZ10-AA Controller LED Codes	Lists and describes the codes, indicates whether a code is solid, cyclic, or flashing, and the state of the device.	Chapter 5
Table 12-1, Array Component Errors	Lists the error messages, their probable cause, and action(s) to take.	Chapter 12
Table 12-2, Array Device Errors	Lists the error messages, their probable cause, and action(s) to take.	Chapter 12
Table 12-3, Other Status Change Errors	Lists the error messages, their probable cause, and action(s) to take.	Chapter 12
Table 12-4, RAID Manager Messages	Lists the error messages, their probable cause, and action(s) to take.	Chapter 12

(continued on next page)

Error Information and Symptoms Tables

7.1 System Message Tables

Table 7-1 (Cont.) System Messages and How to Find Them

Table	Description of Contents	Location
Table A-1, SCSI-2 Status Codes for the Array Controller	Lists the status description and corresponding code.	Appendix A
Table A-2, SCSI Sense Keys	Lists and describes the sense keys.	Appendix A
Table A-3, SCSI Error Codes	Lists the Additional Sense Code, the Additional Sense Code Qualifier, the sense key and its description.	Appendix A
Table B-1, SCSI Status Codes for the KZESA Host Adapter	Lists and describes the KZESA host adapter error codes.	Appendix B

7.2 Symptoms and Status Tables

This sections lists the tables of symptoms and status codes available in this guide. Table 7-2 describes the contents of the tables, and indicates where to locate them.

Table 7-2 Symptoms and Status Codes and How to Find Them

Table	Description of Contents	Location
Table 4-2, LUN Status from the Disk Array Devices Window	Lists and describes the logical unit (LUN) status as it appears on the Disk Array Devices window.	Chapter 4
Table 4-3, Drive Status	Lists and describes drive status as it appears on the Disk Drive Information window.	Chapter 4
Table 4-6, LUN Status from the ACE Window	Lists and describes the LUN status as it appears on the ACE window.	Chapter 4
Table 5-1, LUN States	Lists and describes the LUN state.	Chapter 5
Table 5-4, Shelf and Single Power Supply Status LEDs	Lists the position of the shelf and single power supply LED, its status, and describes the meaning of the status.	Chapter 5
Table 5-5, Shelf and Dual Power Supply Status LEDs	Lists the position of the shelf and dual power supply status LED, its status, and describes the meaning of the status.	Chapter 5

(continued on next page)

Table 7-2 (Cont.) Symptoms and Status Codes and How to Find Them

Table	Description of Contents	Location
Table 5-6, Drive SBB Status LEDs	Lists the position of the drive status LED, its status, and describes the meaning of the status.	Chapter 5
Table 6-1, Troubleshooting Systems Problems	Describes symptoms, gives possible causes, and provides corrective actions.	Chapter 6

7.3 Error Reporting under OpenVMS VAX

This section discusses the methods for reporting errors used by the utilities and by the OpenVMS VAX operating system including the following:

- Errors Reported by the StorageWorks RAID Array 110 Utility for OpenVMS VAX
- Errors Reported by OpenVMS VAX

7.3.1 Errors Reported by the StorageWorks RAID Array 110 Utility for OpenVMS VAX

One purpose of the StorageWorks RAID Array 110 Utility for OpenVMS VAX is to provide information to you on the status of the LUNs and drives. This information is available in a variety of formats based on which utility you use as follows:

Utility/Command	Description
Array Configuration Editor (ACE) Utility	Provides a concise display of LUNS, drives and the status of each LUN or drive. For more information, see Section 4.4.
Array Status Monitor (ASM) STATUS command	Displays all LUN parameters and the drives associated with a particular LUN. For more information, see Section 4.6.

The Array Status Monitor also runs a process that periodically checks the status of the LUN or drive and reports queues through OPCOM any changes from the optimal state.

These utilities are used as the primary tool to determine an error condition with either a LUN or drive. LUN and drive statuses are described in detail in Section 4.4.

Errors reported by the utilities may or may not contain actual SCSI bus level error information. A SCSI-2 type error is the combination of three different codes returned by the Request Sense SCSI command. These are the Additional Sense Code (ASC), Additional Sense Code Qualifier (ASCQ), and the Sense Key. These codes combined describe a particular error condition with the LUN or the drive. A complete listing of these codes and their definitions is contained in Section A.2.

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

There are other error codes associated with the actual running of the utilities. These errors are reported when an attempted utility operation fails. Table 7-3 lists the error codes and their definitions.

Table 7-3 StorageWorks RAID Array 110 Utility for OpenVMS VAX Error Codes

Error Code	Description
E-ARG_RANGE_ERROR	Argument range error
E-BAD_DEV_NAME	Invalid device name on command line
E-BAD_OS_REV_LEVEL	Error, Incorrect operating system revision level
E-BAD_RELID_OPEN	Error, Could not determine the operating system revision level
E-CAN_FILE_OPEN_ERROR	Error, cannot open the file of pre-defined configurations
E-CMDLINE_SYNTAX_ERROR	Command line syntax error
E-CONV_MAJ_MIN_ERROR	Error, Could not convert hw-addr to device major/minor number
E-DAC_SBUS_RANGE_ERROR	Argument range error, Disk array controller SCSI device ID (p delimiter) should be in the range 0-15
E-DISPLAY_TIME_NI_ERR	dac_time.display not implemented
E-DOWNLOAD_TIMEOUT_MSG	Error, timed out waiting for controller to come ready after download
E-EXCLUSIVE_ACCESS_ERROR	Error, Cannot obtain exclusive access to the device
E-FILE_LOCK_ERROR	Error, could not lock the amd message file
E-FILE_OPEN_ERROR	Error, could not open the amd message file
E-FORMAT_ERROR	Error, cannot perform format operation
E-GENERAL_DEV_OPEN_ERROR_MSG	Error, could not open device
E-GENERAL_FILE_NOT_FOUND_MSG	Error, file does not exist
E-GENERAL_FILE_OPEN_ERROR_MSG	Error, could not open file
E-INAPPROPRIATE_RAIDLVL_MSG	Device RAID level is inappropriate for this function
E-INVALID_HWADDR	Invalid hardware address
E-INVALID_PF_OPTS	Invalid combination of PUN CUTOFF and FIRST PUN options
E-INVALID_SELECT	Invalid selection
E-IOB_RANGE_ERROR	Argument range error, System I/O bus (m delimiter) should be in the range 0-7
E-IOCTL_FAILURE_MSG	Error, I/O control function failed
E-MISCOMP_RANGE_ERROR	Error, Miscompared block not in expected range
E-MODE_SELECT_ERROR	Error, MODE SELECT failure
E-MODE_SENSE_ERROR	Error, MODE SENSE failure
E-NEW_SCSI_DEV_FAIL	Error, could not open SCSI device

(continued on next page)

Table 7-3 (Cont.) StorageWorks RAID Array 110 Utility for OpenVMS VAX Error Codes

Error Code	Description
E-NON_LUN	Non-existent LUN
E-NOT_OFFLINE_MSG	Cannot repair parity on an open volume
E-NOT_SINGLE_USER_MSG	Must be in single-user mode to repair parity on root
E-NO_ERROR_MESSAGE	Error, and message does not exist
E-NO_INQUIRY_DATA	SCSI-related error, could not obtain inquiry data
E-OUT_OF_MEMORY_MSG	Error, could not allocate memory
E-READ_AAD_ERR	Error in reading Disable Array Assurance Flag
E-READ_ADR_ERR	Error in reading Auto Detect Replacement Flag
E-READ_ARCMDE_ERR	Error in setting alternate redundant controller mode
E-READ_ARRE_ERR	Error in reading ARRE Flag
E-READ_AWRE_ERR	Error in reading AWRE Flag
E-READ_BLKSIK_ERR	Error in reading Block Size
E-READ_CAPACITY_FAIL	Error, READ CAPACITY failure
E-READ_DMW_ERR	Error in reading Degrade Mode Writes Flag
E-READ_DRVMAP_ERR	Error in reading Drive Map
E-READ_DRVSTS_ERR	Error in reading Drive Status
E-READ_LUNSTS_ERR	Error in reading LUN Status
E-READ_MDEPGE_ERR	Error in reading Type of Mode Page to Sense Flag
E-READ_NBLKS_ERR	Error in reading LUN Capacity
E-READ_PER_ERR	Error in reading PER Flag
E-READ_RAIDLK_ERR	Error in reading Raid Level
E-READ_RCAMT_ERR	Error in reading Reconstruction Amount
E-READ_RCFREQ_ERR	Error in reading Reconstruction Frequency
E-READ_RCMDE_ERR	Error in reading redundant controller mode
E-READ_RCPROG_ERR	Error in reading reconstruction progress
E-READ_RRETRY_ERR	Error in reading Read Retry Count
E-READ_SECSIZ_ERR	Error in reading Sector Size
E-READ_SEGSIZ_ERR	Error in reading Segment Size
E-READ_SG0SIZ_ERR	Error in reading Segment 0 Size
E-READ_TIME_NI_ERR	dac_time.read not implemented
E-READ_WRETRY_ERR	Error in reading Write Retry Count

(continued on next page)

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Table 7-3 (Cont.) StorageWorks RAID Array 110 Utility for OpenVMS VAX Error Codes

Error Code	Description
E-RECONSTRUCT_RAID_ZERO	Reconstruction of RAID 0 occurred - Restore data from backup media
E-SBUS_RANGE_ERROR	Argument range error, System SCSI bus (b delimiter) should be in the range 0-7
E-SCSI_LOCK_FAIL	Error, could not lock SCSI device
E-SCSI_NULL_ERROR	Invalid selection
E-SCSI_READ_ERROR	Error, SCSI error encountered during READ operation
E-SCSI_REPPARITY_ERROR	Error, SCSI error encountered during REPAIR_PARITY operation
E-SCSI_VERIFY_ERROR	SCSI-related error during Block_Verify
E-SCSI_WRTVERIFY_ERROR	Error, SCSI error encountered during WRITE_VERIFY operation
E-SET_AAD_ERR	Error in setting Disable Array Assurance Flag
E-SET_ADR_ERR	Error in setting Auto Detect Replacement Flag
E-SET_ARCMDE_ERR	Error in setting alternate redundant controller mode
E-SET_ARRE_ERR	Error in setting ARRE Flag
E-SET_AWRE_ERR	Error in setting AWRE Flag
E-SET_BLKRECON_ERR	Error in setting highest block number reconstructed
E-SET_BLKSIKZ_ERR	Error in setting Block Size
E-SET_DMW_ERR	Error in setting Degrade Mode Writes Flag
E-SET_DRVMAP_ERR	Error in setting Drive Map
E-SET_DRVSTS_ERR	Error in setting Drive Status
E-SET_LUNSTS_ERR	Error in setting LUN Status
E-SET_MDEPGE_ERR	Error in setting Type of Mode Page to Sense Flag
E-SET_NBLKS_ERR	Error in setting LUN Capacity
E-SET_PER_ERR	Error in setting PER Flag
E-SET_RAIDLK_ERR	Error in setting Raid Level
E-SET_RCAMT_ERR	Error in setting Reconstruction Amount
E-SET_RCFREQ_ERR	Error in setting Reconstruction Frequency
E-SET_RCMDE_ERR	Error in setting redundant controller mode
E-SET_RCPROG_ERR	Error, illegal to set reconstruction progress
E-SET_RRETRY_ERR	Error in setting Read Retry Count
E-SET_SECSIZ_ERR	Error in setting Sector Size
E-SET_SEGSIZ_ERR	Error in setting Segment Size

(continued on next page)

Table 7-3 (Cont.) StorageWorks RAID Array 110 Utility for OpenVMS VAX Error Codes

Error Code	Description
E-SET_SG0SIZ_ERR	Error in setting Segment 0 Size
E-SET_TIME_ERR	Error in setting DAC Time
E-SET_WRETRY_ERR	Error in setting Write Retry Count
E-SLOT_RANGE_ERROR	Argument range error, System SCSI slot (c delimiter) should be in the range 0-15
E-SLUN_RANGE_ERROR	Argument range error, System SCSI LUN (l delimiter) should be in the range 0-7
E-SPUN_RANGE_ERROR	Argument range error, System SCSI device ID (p delimiter) should be in the range 0-15
E-SYSTEM_CALL_ERROR	Error, cannot execute pre-defined configuration script
E-VERIFY_FAIL_NOSENSE	Error, SCSI verify failed, but no sense data is available
I-INVALID_RANGE	Error, entry out of proper range

7.3.2 Errors Reported by OpenVMS VAX

OpenVMS VAX V5.5-2 views the array controller as a generic SCSI device. Any system level errors, which increment the error count for that device, have entries in the error logs. These logs can be viewed using the ANALYZE/ERROR utility. If the error resulted in a Request Sense SCSI command issued to the device, then the sense data displays in the log, if it is present. Example 7-1 shows a typical log entry.

Example 7-1 Error Log Example

```

Error Log Report Generator
***** ENTRY 873. *****
ERROR SEQUENCE 96.
DATE/TIME 16-MAR-1993 06:12:58.64
SYSTEM UPTIME: 0 DAYS 14:18:05
SCS NODE:
DEVICE ERROR KA46 CPU FW REV# 3. CONSOLE FW REV# 0.1
GENERIC DK SUB-SYSTEM, UNIT _DKA500: 1

```

Version V5.5
LOGGED ON: SID 12000003
SYS_TYPE 04010002
VAX/VMS V5.5-2

(continued on next page)

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Example 7-1 (Cont.) Error Log Example

HW REVISION	34303330	HW REVISION = 0304 ②
ERROR TYPE	05	EXTENDED SENSE DATA RECEIVED ③
SCSI ID	05	SCSI ID = 5. ④
SCSI LUN	00	SCSI LUN = 0.
SCSI SUBLUN	00	SCSI SUBLUN = 0. ⑤
PORT STATUS	00000001	%SYSTEM-S-NORMAL, NORMAL SUCCESSFUL COMPLETION
SCSI CMD	A93B050A 0001	WRITE ⑥
SCSI STATUS	02	CHECK CONDITION
EXTENDED SENSE DATA ⑦		
EXTENDED SENSE	000100F0 98A93B05 00000000 00100009 00000000 00000001 00000000 00000000 00000000 01000000 3B050A01 000001A9 00000000 54310F00 31313332 34353030 00000000 33300000 00013430 54100000 52363833 2036325A 20202020 20294328 00434544 00100000 000100F0 0AA93B05 00000000 00020009 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000	

(continued on next page)

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Example 7-1 (Cont.) Error Log Example

UCBSB_ERTCNT	88	RECOVERED ERROR ⑧
UCBSB_ERTMAX	00	TRACK FOLLOWING ERROR
ORBSL_OWNER	00800088	136. RETRIES REMAINING
UCBSL_CHAR	1DCD4008	0. RETRIES ALLOWABLE
		OWNER UIC [200,210]
		DIRECTORY STRUCTURED
		FILE ORIENTED
		SHARABLE
		AVAILABLE
		MOUNTED
		ERROR LOGGING
		ALLOCATED
		FOREIGN
		CAPABLE OF INPUT
		CAPABLE OF OUTPUT
		RANDOM ACCESS
UCBSW_STS	0000	
UCBSL_OPCNT	0013089B	1247387. QIO'S THIS UNIT
UCBSW_ERRCNT	0001	1. ERRORS THIS UNIT
IRPSW_BCNT	0200	TRANSFER SIZE 512. BYTE(S)
IRPSW_BOFF	0010	16. BYTE PAGE OFFSET
IRPSL_PID	0001002B	REQUESTOR "PID"
IRPSQ_IOSB	00000080	
	000A0000	IOSB, 0. BYTE(S) TRANSFERRED

- ① Device name
- ② HSZ10 software revision number
- ③ Indicates that additional REQUEST SENSE information follows. See ⑦.
- ④ SCSI ID number
- ⑤ LUN number
- ⑥ SCSI command
- ⑦ Extended sense data (REQUEST SENSE data)—256 bytes
- ⑧ Top-level problem definition

As seen in Example 7-1, all of the Request Sense data is presented to the user as a 256 byte dump. A complete discussion of this data dump is presented in Section 7.3.3.

7.3.3 Request Sense Data Format

Example 7-2 shows the format of the sense data returned by the Request Sense command. The information in the example is described in the following sections.

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Example 7-2 Request Sense Data Format

	7	6	5	4	3	2	1	0

0-17	SCSI Standard Sense Data							

0	Valid Error Code (0x70 or 0x71)							

1	00							

2	0	0	ILI	0	Sense Key			

3-6	Information							

7	Additional Sense Length; 152 (0x98) or 10 (0x0A)							

8-11	Command Specific Information							

12	Array Additional Sense Code (ASC)							

13	Array Additional Sense Code Qualifier (ASCQ)							

14	Field Replaceable Unit (FRU)							

15-17	SKSV Sense Key Specific Bytes							

18-40	Error Recovery Information							

18-19	Recovery Action(s)							

20	Total Number of Errors (for this operation)							

21	Total Retry Count (for this operation)							

22-25	ASC/ASCQ Stack (for multiple errors)							

26-33	Additional FRU Information							

34-36	Error Specific Information							

37-40	Error Detection Point							

41-108	Environment Information							

41-50	Original CDB							

51	Host ID							

52-53	Host Descriptor							

54-69	Array Board Serial Number							

70-73	Array Application Software Revision Level							

74	Data Transfer Operation							

(continued on next page)

Example 7-2 (Cont.) Request Sense Data Format

75	LUN Number
76	LUN Status
77	Drive Identifier (for non-zero LUN status)
78	Transfer Start Drive ID
79-82	Drive Software Revision Level
83-98	Drive Product ID (from drive Inquiry)
99-100	Array Power-Up Status
101	RAID Level
102-135	Drive Sense Data
102-103	Drive Sense ID
104-135	Drive Sense Data
136-159	Expansion Area
136-159	Reserved (zero filled)

Field Descriptions

Incorrect Length Indicator (ILI) Bit (Offset 2) This bit informs the host system that the requested non-zero byte transfer length for a Read or Write Long command does not exactly match the available data length. The information field in the sense data is set to the difference (residue) of the requested length minus the actual length in bytes. Negative values are indicated by two's complement notation. Because the controller does not support Read or Write Long, this bit is always zero.

Sense Key The following table shows the possible sense keys returned.

Sense Key	Description
0x00	No Sense
0x01	Recovered Error
0x02	Not Ready
0x03	Medium Error
0x04	Hardware Error
0x05	Illegal Request

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Sense Key	Description
0x06	Unit Attention
0x07	Data Protect
0x0B	Aborted Command
0x0E	Miscompare

Information Bytes (Offset 3—6) This field is implemented as defined in the SCSI standard for direct access devices. The information is either of the following types:

- The unsigned logical block address indicating the location of the error being reported.
- The first invalid logical block address if the sense key indicates an Illegal Request.

Additional Sense Length - 152 (0x98) or 10 (0x0A) Bytes (Offset 7) This value indicates the number of additional sense bytes to follow. Some errors cannot return valid data in all of the defined fields. For these errors, invalid fields are zero-filled unless the SCSI-2 standard specifies them as containing 0xFF if invalid.

The value in this field is 152 (0x98) in most cases. However, in some situations the controller returns only the standard sense data. These situations include, but are not limited to the following:

- Requests for sense data when the controller has no sense stored
- Requests for sense data before the controller has loaded application software

For these sense blocks, the additional sense length is 10 (0x0A).

Command-Specific Information (Offset 8—11) This field is valid only for sense data returned after an unsuccessful Reassign Blocks command. The logical block address of the first defect descriptor not reassigned is returned in this field. These bytes are 0xFFFFFFFF if information about the first defect descriptor not reassigned is not available or if all the defects have been reassigned.

The command-specific field is always zero-filled for sense data returned for commands other than Reassign Blocks.

Additional Sense Code (ASC) And Qualifier (ASCQ) (Offset 12, 13) See Appendix A for the supported sense codes and qualifiers returned in these fields.

Field Replaceable Unit (FRU) Code (Offset 14) A non-zero value in this byte identifies a field replaceable unit that has failed or a group of field replaceable modules that includes one or more failed devices.

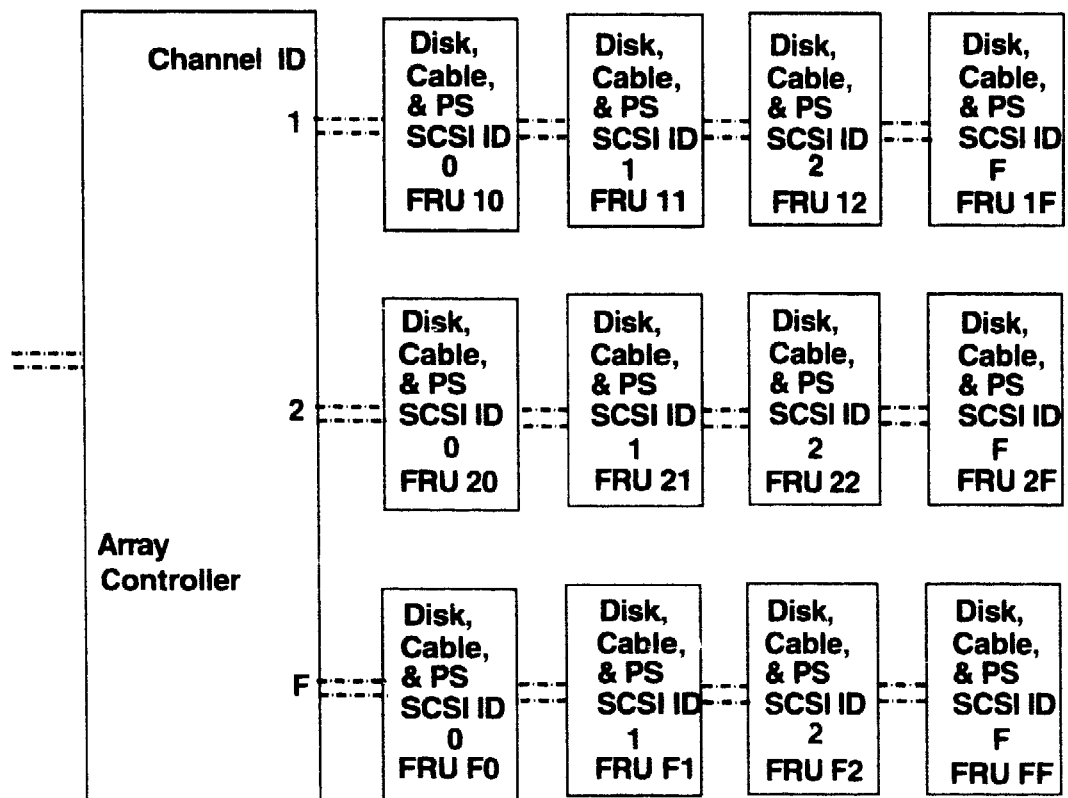
Because of the large number of drives, cables, and power supplies possible in an array, a single byte is not sufficient to report a unique identifier for each individual field replaceable unit. To provide meaningful information to help decrease field troubleshooting and problem resolution time, FRUs have been grouped. The following table lists the defined FRU groups:

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

FRU Code	Description
0x01	Host Channel Group—consists of the host SCSI bus, its SCSI interface chip, and all initiators and other targets connected to the bus.
0x02	Controller Drive Interface Group—consists of the SCSI interface chips on the controller that connect to the drive buses.
0x03	Controller Buffer Group—consists of the controller logic used to implement the on-board data buffer.
0x04	Controller Array ASICs Group—consists of the ASICs on the controller that are associated with the array functions.
0x05	Controller Other Group—consists of all controller-related hardware not associated with another group. This group includes the controller power supply.
0x06	Subsystem Group—consists of subsystem components (such as power supplies, fans, thermal sensors, and AC power monitors) that are monitored by the controller.
0x07-0x0F	Reserved
0x10-0xFF	Drive Group—consists of a drive (embedded controller, drive electronics, and head disk assembly), its power supply, and the SCSI cable that connects it to the previous drive in the chain. Figure 7-1 illustrates the drive groups for a theoretical 15-channel array with 15 drives on each channel. The hexadecimal FRU group codes are in parentheses.

Figure 7-1 Drive Groups



----- Indicates SCSI Cable

SHR-XR3028-GRA

The FRU codes designate the channel ID in the most significant nibble and the SCSI ID of the drive in the least significant nibble.

NOTE: Channel ID 0 is not used because a failure of drive ID 0 on this channel causes an FRU code of 0x00, which the SCSI-2 standard defines as no specific unit has been identified to have failed or the data is not available.

You can obtain additional information about the specific failure within the FRU group from the additional FRU bytes field of the array sense.

For some Additional Sense Codes, you must use the FRU code to determine where the error occurs. For example, the Additional Sense Code for SCSI bus parity error is returned for a parity error detected either on the host bus or on one of the drive buses. In this case, you must evaluate the FRU field to determine if the error occurred on the host channel or on a drive channel.

Sense-Key-Specific Bytes (Offset 15—17) This field is valid for sense keys of Illegal Request, Recovered Error, Hardware Error, and Medium Error when the Sense-Key-Specific Valid (SKSV) bit is on.

When the sense key is Illegal Request and the SKSV bit is set to one, the Sense-Key-Specific field contains the data defined in the following table:

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

	7	6	5	4	3	2	1	0
15	SKSV	C/ D1	Reserved		BPV 2	Bit Pointer		
16	Field Pointer (MSB)							
17	Field Pointer (LSB)							

- 1 If C/D = 0, the illegal parameter is in the parameters sent during a Data Out phase. If C/D = 1, the illegal parameter is in the CDB.
- 2 If BPV = 0, the value in the Bit Pointer field is not valid. If BPV = 1, the Bit Pointer field specifies which bit of the byte designated by the Field Pointer field is in error. When a multiple-bit error exists, the Bit Pointer field points to the most-significant (left-most) bit of the field.

The Field Pointer field indicates which byte of the CDB or the parameter was in error. Bytes are numbered from zero. When a multiple-byte field is in error, the pointer points to the most-significant byte.

If the sense key is Recovered Error, Hardware Error, or Medium Error (and if the SKSV bit is set to one), the Sense-Key-Specific field is defined as shown in the following table:

	7	6	5	4	3	2	1	0
15	SKSV	Reserved						
16	Actual Retry Count (MSB)							
17	Actual Retry Count (MSB)							

The Actual Retry Count field in the previous table indicates the number of times the controller executed its recovery algorithm in attempting to recover the error or exception condition being reported. Retries performed by the drive are indicated in the drive sense data portion of the array sense. A count of the total retries performed for the I/O is returned in byte 21 of the array sense. The Total Retry Count field can be greater than this field, because it includes retries of operations within the recovery algorithm.

Recovery Actions (Offset 18, 19) This bit-significant field indicates the recovery actions that the controller performs. The following table describes the function of the bits:

Byte/Bit	Description
Byte 19, bit 0	Parity used
Byte 19, bit 1	Error manager reconstruction
Byte 19, bit 2	Issued Rezero Unit command
Byte 19, bit 3	Issued Reassign Block command
Byte 19, bit 4	Issued Start/Stop Unit command
Byte 19, bit 5	Issued Abort or Bus Device Reset message
Byte 19, bit 6	Asserted Reset signal on drive channel

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Byte/Bit	Description
Byte 19, bit 7	Retried SCSI operation without chip sequences
Byte 18, bit 0	Transferred data asynchronously
Byte 18, bit 1	Renegotiated synchronous parameters
Byte 18, bit 2	Retried drive busy status
Byte 18, bit 3	Drive marked in Warning state
Byte 18, bit 4	Drive command retried
Byte 18, bit 5	Downed drive
Byte 18, bit 6	Downed LUN
Byte 18, bit 7	Information logged

Total Number Of Errors (For This Operation) (Offset 20) This field contains a count of the total number of errors encountered during execution of the command. The ASC and ASCQ for the last two errors encountered are in the ASC/ASCQ stack field.

Total Retry Count (For This Operation) (Offset 21) The total retry count is for all errors seen during command execution. If only one error was encountered, the value in this field can be the same as that in the Sense-Key-Specific bytes.

ASC/ASCQ Stack (For Multiple Errors) (Offset 22—25) These fields store information when multiple errors are encountered during execution of a command. The ASC/ASCQ pairs are presented in order of the most recent to the least recent error detected.

Additional FRU Information (Offset 26—33) These bytes provide additional information about the FRU identified in byte 14. The first two bytes are qualifier bytes that provide details about the FRU in byte 14. Byte 28 is an additional FRU code that identifies a second FRU. Interpret the value in byte 28 by using the description for byte 14. Bytes 29 and 30 provide qualifiers for byte 28, just as bytes 26 and 27 provide qualifiers for byte 14.

The following table shows the layout of this field. Following the table is a description of the FRU group code qualifiers:

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

	7	6	5	4	3	2	1	0
26	FRU Group Qualifier for byte 14 (MSB)							
27	FRU Group Qualifier for byte 14 (LSB)							
28	Additional FRU Group Code							
29	Additional FRU Group Code Qualifier (MSB)							
30	Additional FRU Group Code Qualifier (LSB)							
31	Second Additional FRU Group Code							
32	Second Additional FRU Group Code Qualifier (MSB)							
33	Second Additional FRU Group Code Qualifier (LSB)							

FRU Group Qualifiers For The Drive Group Codes 0x10-0xff If the FRU Group Code identifies a drive group (0x10-0xff), refer to the following table for a description of the group qualifier:

Bit	Description
MSB bits 0-7	Reserved
LSB bit 0	Drive HDA
LSB bit 1	Drive controller electronics
LSB bit 2	Drive power supply
LSB bit 3	SCSI cable (drive)
LSB bits 4-7	Reserved

FRU Group Qualifiers For The Host Channel Group Code 0x01 If the FRU Group Code identifies the host channel (0x01), refer to the following table for a description of the group qualifier:

Bit	Description
MSB bits 0-7	Reserved
LSB bit 0	SCSI bus cable
LSB bit 1	SCSI interface chip
LSB bit 2-7	Reserved

FRU Group Qualifiers For The Controller Drive Interface Group Code 0x02 If the FRU Group Code identifies the controller drive interface (0x02), refer to the following table for a description of the group qualifier:

Bit	Description
MSB bit 0	Channel 9 SCSI chip
MSB bit 1	Channel 10 SCSI chip
MSB bit 2	Channel 11 SCSI chip
MSB bit 3	Channel 12 SCSI chip

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Bit	Description
MSB bit 4	Channel 13 SCSI chip
MSB bit 5	Channel 14 SCSI chip
MSB bit 6	Channel 15 SCSI chip
MSB bit 7	Reserved
LSB bit 0	Channel 1 SCSI chip
LSB bit 1	Channel 2 SCSI chip
LSB bit 2	Channel 3 SCSI chip
LSB bit 3	Channel 4 SCSI chip
LSB bit 4	Channel 5 SCSI chip
LSB bit 5	Channel 6 SCSI chip
LSB bit 6	Channel 7 SCSI chip
LSB bit 7	Channel 8 SCSI chip

FRU Group Qualifiers For The Controller Buffer Group Code 0x03 If the FRU Group Code identifies the controller buffer group (0x03), refer to the following table for a description of the group qualifier:

Bit	Description
MSB bits 0-7	Reserved
LSB bits 0-7	Failing memory device identifier

FRU Group Qualifiers For The Controller Array ASIC Group Code 0x04 If the FRU Group Code identifies the controller array ASIC group (0x04), refer to the following table for a description of the group qualifier:

Bit	Description
MSB bits 0-7	Reserved
LSB bit 0	SDP Chip
LSB bits 1-7	Reserved

FRU Group Qualifiers For The Controller Other Group Code 0x05 If the FRU Group Code identifies controller other logic group (0x05), refer to the following table for a definition of the group qualifier:

Bit	Description
MSB bits 0-7	Reserved
LSB bit 0	Microprocessor
LSB bit 1	Array controller power supply
LSB bit 2	Nonvolatile memory
LSB bits 3-7	Reserved

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

FRU Group Qualifiers For The Subsystem Group Code 0x06 If the FRU Group Code identifies the subsystem group (0x06), refer to the following table for a definition of the group qualifier:

Bit	Description
MSB bit 0	Value in LSB determined from discrete input signals.
MSB bits 1-255 (0xFF)	Reserved
LSB bits 0-1	Reserved
LSB bit 2	Fault indicated at hardware address 0x1C6423, bit 5
LSB bits 3-7	Reserved
LSB bit 8	Fault indicated at hardware address 0x1C6425, bit 0
LSB bit 9	Fault indicated at hardware address 0x1C6425, bit 1
LSB bit 10 (0x0A)	Fault indicated at hardware address 0x1C6425, bit 2
LSB bit 11 (0x0B)	Fault indicated at hardware address 0x1C6425, bit 3
LSB bit 12 (0x0C)	Fault indicated at hardware address 0x1C6425, bit 4
LSB bit 13 (0x0D)	Fault indicated at hardware address 0x1C6425, bit 5
LSB bit 14 (0x0E)	Fault indicated at hardware address 0x1C6425, bit 6
LSB bit 15 (0x0F)	Fault indicated at hardware address 0x1C6425, bit 7
LSB bits 16-255 (0x10-0xFF)	Reserved

NOTE: Refer to the *StorageWorks RAID Array 110 Subsystem User's Guide* for the specific board being used.

If none of the bits in a FRU qualifier byte are on, the controller is unable to isolate the error to a lower level than the FRU group.

The additional FRU group is loaded when errors on multiple FRU groups are detected. For example, if two drives report Not Ready sense keys, both drives are identified: one by the FRU group in byte 14 (and its qualifier), and the other by the Additional FRU group code (and its qualifier).

The Additional FRU group code and its qualifier are also loaded if an error can not be isolated to a single group. For example, when a bus error occurs that could have been caused by either a SCSI chip problem on the array or a drive problem, both FRU group codes and qualifiers are loaded.

Error-Specific Information (Offset 34—36) This field provides information read from the controller VLSI chips and other sources. It is intended primarily for development testing and the contents are not specified.

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Error Detection Point (Offset 37—40) The Error Detection Point field indicates where in the software the error was detected. It is intended primarily for development testing and the contents are not specified.

Original CDB (Offset 41—50) This field contains the original CDB received from the host.

Host ID (Offset 51) The host ID is the SCSI ID of the host that selected the controller to execute a command.

Host Descriptor (Offset 52, 53) This bit-position field provides information about the host as described in the following table:

Bit	Description
MSB bit 0	Data is being transferred 16-bit wide
MSB bit 1	Data is being transferred 32-bit wide
MSB bit 2	Wide negotiation completed successfully (may have negotiated for 8-bit wide)
MSB bits 3-7	Reserved
LSB bit 0	Message-using host
LSB bit 1	Reselectable host
LSB bit 2	Data is being transferred synchronously (bit off means asynchronous data transfer)
LSB bit 3	Synchronous negotiation successful (may have negotiated for asynchronous)
LSB bit 4	Reserved
LSB bit 5	AEN supported
LSB bit 6	Polled AEN supported
LSB bit 7	Reserved

Array Board Serial Number (Offset 54—69) This sixteen-byte field contains the manufacturing identification of the array hardware. Bytes of this field are identical to the information returned by the Unit Serial Number page in the Inquiry Vital Product Data.

Array Application Software Revision Level (Offset 70—73) This revision level matches the revision level returned by an Inquiry command.

Data Transfer Operation (Offset 74) This byte indicates the controller's data transfer mode at the time of the error. By using this byte along with the RAID level of the LUN, you can determine the logical step being executed when the error occurred. If the most significant bit (bit 7) is set, it indicates that the first error encountered occurred on the parity drive.

Value in Bits 0-6	Description
0x00	Read operation to the drive(s)
0x01	Read portion of a read/modify/write operation for RAID 5 Write operation to the drive(s) for RAID 0, 1, or 3
0x02	Write portion of a read/modify/write operation for RAID 5
0x03	Read recovery operation

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Value in Bits 0-6	Description
0x04	Read recovery for the read portion of a read/modify/write operation in RAID 5 Write recovery in RAID 1 or 3
0x05	Write recovery for the write portion of a read/modify/write operation in RAID 5
0x06	Reconstruction
0x07	Processor data transfer to source
0x08	Processor data transfer from source
0x09	Processor data transfer to destination
0x0A	Processor data transfer from destination
0x0B	Transfer zeros to destination
0x0C	SRAM transfer to source
0x0D	SRAM transfer from source
0x0E	Verify parity operation
0x0F	Processor to source transfer followed by save pointers
0x10	Processor from source transfer followed by save pointers
0x11	Write data and generate parity by reading all other data drives
0x12	SRAM read from destination
0x13	SRAM write to destination
0x14	Processor from source with no disconnect
0x15	No data transfer
0x16	Copy data channel to channel
0x17	Completing the write operation of a read/modify/write after an abort condition
0x18	Continuous write operation
0x19	SDP data transfer abort
0x1A	Reserved
0x1B	Two-drive reconstruct with SRAM
0x1C-0x7F	Reserved

LUN Number (Offset 75) The LUN Number field is the logical unit number in the Identify message received from the host after selection.

LUN Status (Offset 76) This field indicates the status of the LUN. It is defined in the Logical Array page description except for the value of 0xFF, which is unique to this field. A value of 0xFF returned in this byte indicates that the LUN is undefined.

Drive Identifier (For Non-Zero LUN Status) (Offset 77) The drive identifier indicates which drive is responsible for a non-zero LUN status. The drive ID assignments are the same as those used for FRU groups.

Transfer Start ID (Offset 78) This field identifies the drive that was accessed first when performing the host operation data transfer. For RAID 3 operations, this field is always 00. Refer to the FRU group designations for ID assignments. For operations that do not involve disk data transfers, this field is 00.

Error Information and Symptoms Tables

7.3 Error Reporting under OpenVMS VAX

Drive Software Revision Level (Offset 79—82) This field specifies the software revision level of the drive involved in the error (if the error was a drive error). This field also specifies the software revision level of the first drive of a rank (if the error is nondrive related; for example, Unit Attention).

Drive Product ID (Offset 83—98) This field identifies the type of drives that make up the LUN. This information is obtained from the Inquiry - Unit Serial Number - Vital Product Data page.

Array Power-Up Status (Offset 99—100) These bytes indicate the status of the level-0 diagnostics run on power-up/reset. Values returned in these two bytes are defined in the Firmware Manual. In this release of the software, these bytes are always set to zero.

RAID Level (Offset 101) This byte indicates the configured RAID level for the LUN returning the sense data. The values that can be returned are 0, 1, 3, or 5.

Drive Sense ID (Offset 102, 103) These bytes identify the source of the sense block returned in the next field. Refer to the FRU group codes for physical drive ID assignments. Byte 103 is reserved for identification of a drive LUN in future implementations. It is always set to zero.

Drive Sense Data (Offset 104—135) For drive detected errors, these bytes contain the data that the drive returns in response to the Request Sense command from the controller. If multiple drive errors occur during the transfer, the sense data from the last error is returned.

Removal and Replacement

This chapter contains information on the following:

- How to replace an HSZ10-AA controller
- How to add an additional controller
- When and how to replace a drive
- When and how to replace power supplies and blowers
- Deleting a logical unit (LUN)
- Modifying and displaying a drive status
- Reconstruction
- Restoring a LUN

8.1 How to Replace an HSZ10-AA Controller

Do the following steps to replace an HSZ10-AA controller board.

Caution

To avoid unnecessary static damage, follow adequate antistatic procedures when handling HSZ10-AA controllers.

If the unit is mounted as a part of a vertical mounting kit (BA35X-VA), then the cabinet door and front bezel should be removed as shown in Figure 8-1. The procedure is described in the *BA35X-VA Vertical Mounting Kit User's Guide*. After you open the cabinet door, press down on the plastic clips at the base of the bezel and pull the bezel toward you.


Note that Figure 8-1 shows the removal of the front drive bezel. Remove the front shelf bezel in the same way that you would remove the drive bezel.

Removal and Replacement

8.1 How to Replace an HSZ10-AA Controller

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


Figure 8-1 Removing the Front Bezel



Replace the HSZ10-AA controller as follows:

1. Power down the subsystem.
2. Loosen the captive slotted screws shown in Figure 8-2 on the host connector cable and remove the cable from the HSZ10-AA controller.

Figure 8-2 HSZ10-AA in a BA350-EA Shelf



3. Use a screwdriver to loosen the large captive screws at the top and bottom of the panel on the BA350-EA shelf to free the HSZ10-AA controller board. You do not have to remove the terminator by unscrewing its thumb screws. That is optional.

Removal and Replacement

8.1 How to Replace an HSZ10-AA Controller


4. Carefully remove the new controller board from the antistatic bag.

Warning

The surfaces of the board are rough and the edges are sharp. Place the board down gently.

5. The existing controller board was loosened in Step 2. Remove the HSZ10-AA controller board by pulling the handles straight toward you and sliding the board along the guide inside the cabinet as shown in Figure 8-3.

Figure 8-3 Removing the HSZ10-AA Controller




SHR-XR3004-VIF

Note that the controller board fits tightly within the shelf. When the board is free of the cabinet, make note of the SCSI ID of the original controller board. The location and setting of the SCSI ID is shown in Figure 8-4.

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

Removal and Replacement
8.1 How to Replace an HSZ10-AA Controller

Figure 8-4 SW1—SCSI ID Settings



SHR-XR3020-GRA

Note


By default, the switch is set to SCSI ID 0. The SCSI ID can be set to any value between 0 and 15. To change the default SCSI ID, you must remove the HSZ10-AA controller from the BA350-EA shelf and then set the SCSI ID.

Figure 8-5 shows examples of SCSI ID 4, 6, and 12.

Removal and Replacement

8.1 How to Replace an HSZ10-AA Controller

Figure 8-5 SCSI ID Setting Examples



Note

If the HSZ10-AA controller is to be used with 8-bit devices on the host SCSI bus, then the SCSI IDs of the controller should be restricted to between 0 and 7.

You can put the original board in the empty antistatic bag that came with the new controller board. If you are replacing the HSZ10-AA controller board because of a defect, place the board in the shipping box and send it to your Digital Customer Service Center for repair or credit.



6. Set the SCSI ID of the replacement controller board to that of the original controller. Slide the new controller into the slot that the original controller came out of using the guides provided. The board must be properly seated. It should snap into place and sit flush within the cabinet before you start to secure the controller by tightening the captive screws.
7. Reattach the host SCSI cable and terminator as necessary.
8. Power on the subsystem.

Removal and Replacement

8.1 How to Replace an HSZ10-AA Controller

9. When the power is on, the controller shows the green power light. The "Heartbeat" LED next to it is lit, and at the other end of the LED indicator has the LED value of 5. See Figure 8-6.



Figure 8-6 LED Indicators



SHR-XR3022-GRA

Figure 8-7 shows the new LED indicators.

Figure 8-7 LED Indicators for the Fully Functional State



SHR-XR3023-GRA

If the LED does not appear as shown in Figure 8-7, then refer to Chapter 5 and Chapter 6.


10. If the unit is contained in a vertical mounting kit, then replace the bezel. Refer to the *BA35X-VA Vertical Mounting Kit User's Guide* for details.

8.2 How to Add an Additional Controller

Follow these steps to add a controller:

1. If two controllers are to be daisy chained together on the same SCSI bus, then set the SCSI ID of the new controller so that it does not conflict with the existing controller.
2. Slide the controller into the empty slot shown as "Controller 2" in Figure 8-8 and secure it in place using the thumb screws.
3. Connect to the existing controller and/or to the host using the appropriate cable provided.

Figure 8-8 Diagram of the BA350-EA Shelf



CXO-3589A-MC

Note

An additional controller will be in a *passive* state.

The additional controller can only be taken out of this state by using a host adapter/operating system environment that can use a redundant controller configuration. Such a configuration is not yet supported.

8.3 When and How to Replace a Drive

You need to replace a drive when the following occurs:


- StorageWorks RAID Array 110 utilities provide information stating that a drive is in a *warning* state.
- StorageWorks RAID Array 110 utilities provide information stating that a drive is in a *failed* state.
- The drive SBB fault indicator is On.

The HSZ10-AA controller and StorageWorks products allow for hot swapping a drive SBB. This means the user can remove and replace a drive SBB without interrupting host operation or removing power to the subsystem.

To replace a drive, perform the following steps:

1. If the drive is not yet in a failed state, fail the drive using the StorageWorks RAID Array 110 utilities. Section 8.6.3 describes the procedure for failing a drive.
2. Wait until the right LED (amber) light is solid (the flashing has stopped).
3. Press the two mounting tabs to release the unit, and slide the unit out of the shelf, as shown in Figure 8-9.

Figure 8-9 Replacing a Drive



CXO-3611B-PH

4. Wait 10 seconds before inserting a new drive SBB.
5. Insert the replacement drive unit into the guide slots and push it in until the tabs lock into place.
6. Wait for the HSZ10-AA controller to spin up the new drive.
7. Perform the reconstruction process, if needed. Refer to Section 8.6.4 for more detailed information.

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

Note

It is important that you wait at least 10 seconds before inserting a new drive SBB. The HSZ10-AA controller requires this period to scan the SCSI buses and to become informed of a drive removal.

8.4 When and How to Replace Power Supplies and Blowers

You need to replace a power supply or blower when the LEDs indicate a fault condition as outlined in Power Supply LEDs.

The input power for each ac power supply is controlled by a switch on the cabinet power controller. Turning this switch off removes power from *all* power supplies in the cabinet. To remove power from a single power supply, disconnect the power cable from that power supply.

There are three swapping methods for replacing power supplies:

- **Hot swapping**—The hot swapping method is used when there are two power supplies in a shelf. This method allows you to remove the defective power supply while the other power supply furnishes the power.
- **Warm swapping**—The warm swapping method is used when there is no operational power supply, but all the other shelf power supplies in a cabinet are functioning properly. In this case, none of the shelf devices are operational until the replacement power supply is installed.

Note

Whenever operational requirements permit, it is recommended that you use the warm swapping method.

- **Cold swapping**—The cold swapping method is used when the input power is removed from all shelves in a cabinet. This normally only occurs during initial installation. None of the shelves are operational until the input power is restored.

8.4.1 Replacing a Shelf Power Supply

Warning

Always use both hands when removing or replacing an SBB, to fully support its weight.

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

Removal and Replacement

8.4 When and How to Replace Power Supplies and Blowers

To replace either a shelf primary or redundant power supply, complete the following procedure:

Step	Procedure
1.	Remove the input power cable from the power supply.
2.	Press the two mounting tabs to release the unit, and slide the unit out of the shelf. This is similar to replacing a drive as shown in Figure 8-9.
3.	Insert the replacement unit into the guide slots and push it in until the tabs lock into place.
4.	Connect the input power cord.
5.	Observe the LEDs and ensure the power supply is functioning properly (refer to Table 5-4 or to Table 5-5).
6.	Note: Use this step only when you have a single power supply. Sequentially place the storage devices online. Observe the LEDs on both the power supply and the storage devices for normal operation indications.

8.4.2 Replacing a Blower

Each shelf has two blowers mounted on the rear. Connectors on the backplane provide the +12 V dc to operate the blowers. As long as one blower is operational on each shelf, there is sufficient airflow to prevent an overtemperature condition. When either blower fails, the upper (shelf status) LED on the power SBB turns on.

Warning

Service procedures described in this manual involving blower removal or access to the rear of the shelf must be performed only by qualified service personnel.

To reduce the risk of electrical energy hazard, disconnect the power cables from the shelf power SBBs before removing shelf blower assemblies or performing service in the backplane area, such as modifying the SCSI bus.

Caution

When a blower is removed, the change in the airflow pattern reduces the cooling to the point that the shelf can overheat within 60 seconds.

Removal and Replacement

8.4 When and How to Replace Power Supplies and Blowers

Note


Replacing a blower requires access to the rear of the shelf. When you cannot access the rear of the shelf you must turn off the power, remove the shelf from the cabinet, and perform Steps 1 through 6 listed in Table 8-1. Replace the shelf in the cabinet and apply power.

To replace a blower, follow the steps in Table 8-1:

Table 8-1 Blower Replacement

Step	Procedure
1.	Disconnect all power cables to shelf power SBBs.
2.	Use a Phillips screwdriver to remove the safety screw in the upper right corner of the blower.
3.	As shown in Figure 8-10, press the upper and lower blower mounting tabs in to release the blower.
4.	Pull the blower straight out to disconnect it from the shelf power connector.
5.	Align the replacement blower connector and insert the modules, straight in, making sure that both mounting tabs are firmly seated in the shelf.
6.	Replace the safety screw in the upper right corner of the blower.
7.	Connect the shelf power cables and verify that the shelf and all SBBs are operating properly.

Figure 8-10 Replacing BA35X-MA Blowers



CXO-3515A-PH

8.5 Deleting a LUN

You delete a LUN to complete one of the following tasks:

- Changing the RAID level of the LUN
- Changing the drive grouping containing a LUN

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

Removal and Replacement

8.5 Deleting a LUN

- Deleting a drive from the array

The following procedure provides the steps to delete a LUN. Normally, there is no reason to use the procedure (deleting a LUN to change RAID level/drive group is described in the procedure in Section 5.3.7) as there is usually no reason to remove a drive.

1. Back up the data on all the LUNs in the drive group.
2. In NetWare, unmount any volumes contained on the LUNs in the drive group you want to change if any are mounted.
3. To create a LUN in MS-DOS or SCO UNIX, or in NetWare to create the LUN that contains your SYS volume, you must use the StorageWorks RAID Array 110 Utility for Novell NetWare. Reboot the system using the StorageWorks RAID Array 110 Utility for Novell NetWare diskette and continue with this procedure.
4. Start ACE from the console by entering one of the following commands:
In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
raidmgr
```


In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
load raidmgr
```


In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

```
$ ACE device-name
```
5. In MS-DOS and NetWare, select the *Array Controller Configuration* option on the menu. Check the Disk Array Devices window on the right of the screen to make sure it is displaying the LUNs you want to modify.
6. In MS-DOS and NetWare, select the *Array Configuration Editor* option from the menu.
7. For all operating systems, note that you are in the Logical Unit window of the Array Configuration Editor (ACE). If you are not in this window, press the Tab key to switch to the Logical Unit window.
8. Use the arrow keys to highlight the LUN you want to modify.
Press:
 - The D key to delete the LUN.
 - The C key to confirm that you want to delete the LUN.
 - The Q key if you do not want to delete the LUN.
9. Repeat this step for each LUN in the drive group. When you have deleted all the LUNs in the group, the status of the highlighted drives in the Drive Matrix changes to *Spare* (2).
10. Press the Q key to quit the Array Configuration Editor.
In MS-DOS and NetWare, you return to the Array Controller Configuration menu. Press the Esc key to quit RAID Manager. Then press the Y key and the Enter key.
11. In the StorageWorks RAID Array 110 Utility for Novell NetWare, remove the diskette and reboot the system.

12. In the StorageWorks RAID Array 110 Utility for OpenVMS VAX, edit the system startup file (SYS\$STARTUP:SYSTARTUP_V5.COM) as described in Section 4.2.1 with the name of the LUN. Then reboot the operating system to prevent further access to this device.

You are done with this procedure.

8.6 Modifying and Displaying a Drive Status

You can add, delete, fail, and reconstruct a drive.

8.6.1 Adding a Drive

Adding a drive through software causes the array controller to change the drive status from *Non-existent* (1) to *Spare* (2), thus making the drive available for you to create a LUN.

Note that the *Add* option in ACE may not be used to revive a drive in a LUN that has a *Failed* (3) status, but should not have been failed (such as in the circumstance of a multiple-drive failure where the LUN status is *Dead - replaced wrong drive*).

Perform the following steps to add a drive:

1. Start ACE from the console by entering one of the following commands:

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
raidmgr
```

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
load raidmgr
```

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

```
$ ACE device-name
```

2. For MS-DOS and NetWare only, select the *Array Controller Configuration* option on the menu. Check the Disk Array Devices window on the right of the screen to make sure it is displaying LUNs attached to the controller connected to the drive you want to add. If not, use the *Select Controller* option to select the correct controller.
3. For MS-DOS and NetWare only, select the *Array Configuration Editor* option from the menu.

In MS-DOS and NetWare, note that you are in the Logical Unit window of the ACE utility. Press the Tab key to switch to the Drive window. Use the spacebar to move the cursor to select the ID and channel of the drive to be added.

Note

To add a drive, you must first physically connect the drive to the array while it is operating.

4. Press the A key to add the drive. The status of the drive changes from *Non-existent* (1) to *Spare* (2).

You have added a drive with the selected ID and channel.

You can now create a LUN using the drive. (See Section 5.3.5.)

8.6.2 Deleting a Drive

Delete a drive when you plan to remove the drive from the StorageWorks RAID Array 110 Subsystem, and are not planning to replace it or reboot the StorageWorks RAID Array 110 Subsystem.

Deleting a drive through software causes the array controller to change the drive status from *Spare* (2) to *Non-existent* (1). This means you can never delete a drive that is part of a LUN.

1. Delete all the LUNs contained on the drive (refer to Section 8.5 to delete the LUN).
2. Delete the drive.
3. Physically remove the drive from the subsystem. Section 8.3 explains how to remove the drive.

Use the following procedure to delete a drive:

1. Start ACE from the console by entering one of the following commands:

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

raidmgr

In StorageWorks RAID Array 110 Utility for Novell NetWare:

load raidmgr

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

\$ ACE device-name

2. For MS-DOS and NetWare only, select the *Array Controller Configuration* option on the menu. Check the Disk Array Devices window on the right of the screen to make sure it is displaying the LUNs attached to the controller connected to the drive you want to delete. If not, use the *Select Controller* option to select the correct controller.
3. For MS-DOS and NetWare only, select the *Array Configuration Editor* option from the menu.

Note that you are in the Logical Unit window of the Array Configuration Editor (ACE). Press the Tab key to switch to the Drive window. Use the Arrow keys to highlight the drive you want to delete.
4. Press the D key to delete the drive. The status of the drive changes from *Spare* (2) to *Non-existent* (1).

You have deleted a drive.

Physically remove the drive(s) you deleted from the array.

8.6.3 Failing a Drive

Fail a drive when you want to replace a *Warning* and restore a LUN.

Failing a drive through software causes the array controller to change the drive status from *Warning* (5) to *Failed* (3). Once a drive is failed, the array controller cannot access that drive's data until you reconstruct the drive data.

Important Note

Do NOT fail a drive in a RAID 0 LUN because there is no parity.

Removal and Replacement

8.6 Modifying and Displaying a Drive Status

Do NOT fail a drive in a RAID 1 or 5 LUN if the LUN is already degraded.

Use the following procedure to fail a drive:

1. Start ACE from the console by entering one of the following commands:
In the StorageWorks RAID Array 110 Utility for Novell NetWare:
`raidmgr`
In the StorageWorks RAID Array 110 Utility for Novell NetWare:
`load raidmgr`
In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
`$ ACE device-name`
2. For MS-DOS and NetWare only, select the *Array Controller Configuration* option on the menu. Check the Disk Array Devices window on the right of the screen to make sure it is displaying the LUNs containing the *Warning* drive you want to fail.
3. For MS-DOS and NetWare only, select the *Array Configuration Editor* option from the menu.
Note that you are in the Logical Unit window of the ACE utility. Press the Tab key to switch to the Drive window. Use the Arrow keys to highlight the *Warning* drive you want to fail.
4. Press the F key to fail the drive. The following message is displayed:

```
If you fail a drive, you may lose
data redundancy or data.
```

```
Press 'c' to continue.
Press 'q' to quit.
```

Press the C key to confirm that you want to continue with the operation you selected, or press the Q key to quit.

If you press the C key, the status of the failed drive changes from *Warning* (5) to *Failed* (3).

You have failed a drive.

Reconstruct the drive data, or reformat the LUN.

8.6.4 Reconstructing a Drive

By default, the HSZ10-AA controller automatically initiates the reconstruction process after you replace a drive in a degraded RAID 1 or 5 LUN. Through the StorageWorks RAID Array 110 utilities, you can change the default configuration and schedule reconstruction or manually start reconstruction on the replaced drive.

Reconstruction is a process used to restore a degraded RAID 1 or 5 LUN to its original state after a single drive has been replaced. During reconstruction, the HSZ10-AA controller recalculates the data on the drive that was replaced, using data and parity from the other drives in the LUN. The controller then writes this data to the replaced drive. Note that although RAID 1 does not have parity, the

Removal and Replacement

8.6 Modifying and Displaying a Drive Status

array controller can reconstruct data on a RAID 1 LUN by copying data from the mirrored disk.

Note

Reconstruction applies only to a degraded RAID 1 or 5 LUN with a single-drive failure.

Remember, a RAID 0 LUN does not have parity so you cannot recover the data by reconstructing it. You must restore the RAID 0 LUN by formatting the LUN, then copying the data from the backup media file-by-file (see Section 8.7 for more information).

A drive can be manually reconstructed if, for some reason, autoreconstruction is not initiated on drive replacement.

Once reconstruction is initiated (either by you or by the array controller), the HSZ10-AA controller completes the following actions:

- Senses that there is a new drive in the array.
- Changes the drive status to *Replace* (4).
- Formats the new drive (if the array controller determines it is necessary).
- Copies special array software files to the new drive.
- Recalculates the data and parity from the data and parity on the other LUN drives.
- Writes the recalculated data and parity to the new drive.

Once reconstruction is started, it can take place while the LUN is in use. You do not need to shut down the StorageWorks RAID Array 110 Subsystem. You can see the access to all the drives in the LUN for a period of time. The length of time it takes for reconstruction to complete depends upon the drive capacity and the reconstruction settings.

Important

Do not remove a second drive during the reconstruction process. Doing so may result in data loss.

If the drive you want to reconstruct contains more than one LUN, all the LUNs are reconstructed.

Use the following procedure to reconstruct a drive:

1. Physically replace the failed drive.
2. Start ACE from the console by entering one of the following commands:

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
raidmgr
```

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
load raidmgr
```

Removal and Replacement

8.6 Modifying and Displaying a Drive Status

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

\$ ACE *device-name*

3. In the StorageWorks RAID Array 110 Utility for Novell NetWare and the StorageWorks RAID Array 110 Utility for Novell NetWare, select the *Array Controller Configuration* option on the menu. Check the Disk Array Devices window on the right of the screen to make sure it is displaying the LUNs containing the drive you want to reconstruct. Select the *Array Configuration Editor* option from the menu.
4. For all operating systems, note that you are in the Logical Unit window of the ACE utility. Press the Tab key to switch to the Drive window. Check the status of the drive you want to reconstruct. If the status is currently *Replaced* (4) or *Formatting* (8) and the LUN status is *Reconstructing*, your system has hot swapping and you do not need to continue with this procedure.
5. Use the arrow keys to highlight the *Failed* drive you want to reconstruct. Press the R key to start reconstruction.

You have started the reconstruction of drive data.

8.7 Restoring a LUN

This section provides information on the following:

- Restoring a LUN task summary
- When to replace a drive
- Drive failures on RAID 0
- Drive failures on RAID 1
- Drive failures on RAID 5
- What is reconstruction?
- Reconstruction rate
- Restoring a RAID 0 LUN
- Restoring a RAID 1 LUN
- Restoring a RAID 1 LUNs after multiple failures
- Restoring a RAID 5 LUN
- Restoring a RAID 5 LUNs after multiple failures

8.7.1 Restoring a LUN Task Summary

Restoring a LUN is the process of returning the LUN to a working state after one or more drive failures. Restoring may or may not involve the ability of the array to reconstruct data from the parity information it has stored. The tasks involved depend on the RAID level of the LUN and on how many drives have failed.

Table 8–2 describes the procedure for restoring LUNs.

Note

Restore refers to copying a copy. *Reconstruct* refers to a rebuilding process.

Removal and Replacement

8.7 Restoring a LUN

Table 8–2 Procedure for Restoring LUNs

RAID Level	LUN/Drive Status	Recovery Steps:
0	OK - Warning (1 or more drives in warning)	<ol style="list-style-type: none"> 1. <i>Stop using the LUN immediately.</i> 2. Back up the data on the LUN. 3. Fail all warning drives. 4. Replace all failed drives. 5. Reformat the LUN. 6. In NetWare, recreate the NetWare partitions and volumes on the restored LUN. 7. Copy backed-up data to the restored LUN.
1	OK - Warning (1 or more drives in warning)	<ol style="list-style-type: none"> 1. If two or more drives are in warning, <i>stop using the LUN immediately and back up the LUN before continuing.</i> 2. <i>One at a time</i>, fail a warning drive, replace the failed drive, and reconstruct the data on it. Then proceed as described above for each drive in a warning state.
1	Degraded (1 or more failed drives)	<ol style="list-style-type: none"> 1. Replace the failed drives (all drives can be replaced at the same time). 2. <i>One at a time</i>, reconstruct the data on the replaced drives. Then proceed as described above for each drive in a warning state.
1	Degraded - Warning (1 or more failed drives and 1 or more drives in warning)	<ol style="list-style-type: none"> 1. <i>Stop using the LUN immediately.</i> 2. Back up the data on the LUN. <ul style="list-style-type: none"> • If the backup is successful, then <i>one at a time</i>, replace and reconstruct each failed drive, then fail, replace, and reconstruct each warning drive. • If the backup is not successful, then fail all warning drives, replace all failed drives, and reformat the LUN. In NetWare, recreate the NetWare partitions and volumes. Copy backed-up data back to the restored LUN.

(continued on next page)

Table 8-2 (Cont.) Procedure for Restoring LUNs

RAID Level	LUN/Drive Status	Recovery Steps:
5	OK - Warning (1 or more drives in warning)	<ol style="list-style-type: none"> 1. If two or more drives are in warning, <i>stop using the LUN immediately and back up the LUN before continuing.</i> 2. Fail a warning drive, replace the failed drive, and reconstruct the data on it. Then proceed as described above for each drive in a warning state.
5	Degraded (1 failed drive)	<ol style="list-style-type: none"> 1. Replace the failed drive. 2. Reconstruct the data on the replaced drive.
5	Degraded - Warning (1 failed drive and 1 or more drives in warning)	<ol style="list-style-type: none"> 1. <i>Stop using the LUN immediately and back up the LUN before continuing.</i> 2. Replace the failed drive and reconstruct the data on it. 3. Fail a warning drive, replace the failed drive, and reconstruct the data on it. Then proceed as described above for each drive in a warning state.
5	Drive Failures (2 or more failed drives and any drives in warning)	<ol style="list-style-type: none"> 1. Fail all warning drives. 2. Replace the failed drives. 3. Reformat the LUN. Data cannot be reconstructed. 4. In NetWare, recreate the NetWare partitions and volumes on the restored LUN. 5. Copy backed-up data to the restored LUN.

8.7.2 When to Replace a Drive

You need to replace a drive when the following occurs:

- For NetWare, the Array Monitor Daemon (AMD) sends a message indicating that a drive has failed or been put into a warning state. Or, you see a LUN status other than *Ok - Optimal* displayed in the RAID Manager Disk Array Devices window.
- For SCO UNIX, the Array Monitor sends a message indicating a change in status to *Failed* or *Warning*.
- For OpenVMS VAX, the Array Status Monitor sends a message to OPCOM indicating a drive has *Failed* or has been put into a warning state.

Removal and Replacement

8.7 Restoring a LUN

Table 8–3 shows the LUN status displayed for the varying degrees of LUN failures.

Table 8–3 LUN Status

RAID Level	LUN Status	Meaning
0	Ok - Warning	One or more drives has a read or write error. The array controller never assigns a <i>Failed</i> status to a drive in a RAID 0 LUN, no matter how severe the error, so this is the most serious status possible for a RAID 0 LUN. Stop using the LUN immediately (if it is still accessible), as further errors may result in loss of all data.
1	Ok - Warning	One or more drives has read or write errors that were not serious enough to fail the drives. The LUN remains accessible. Replace the drive as soon as possible. If two or more drives are in warning, stop using the LUN immediately.
1	Degraded	One or more drives have failed, but not two drives in a mirrored pair. Data on the LUN can be recovered from mirrored disks. The failed drives should be replaced as soon as possible. If they cannot be replaced immediately, back up the data on the LUN to prevent data loss.
1	Degraded - Warning	Either: <ul style="list-style-type: none"> One or more drives have failed <i>and</i> the mirrored pair of a failed drive has also failed. The array controller never assigns a <i>Failed</i> status to the second failed drive in a mirrored pair, no matter how severe the error (it assigns a status of <i>Warning</i>), so this is the most serious status possible for RAID 1 LUN. The LUN is no longer accessible, and all data on the unit is lost. One or more drives have failed <i>and</i> a minor error has occurred on another disk (the disk may or may not be the mirrored pair of a failed disk). The LUN remains accessible.
5	Ok-Warning	One or more of the drives have read or write errors that were not serious enough to fail the drives. The warning drives should be replaced as soon as possible, before they fail. The LUN remains accessible. If two or more drives are in warning, stop using the LUN immediately, as the failure of these drives results in the loss of data.

(continued on next page)

Table 8-3 (Cont.) LUN Status

RAID Level	LUN Status	Meaning
5	Degraded	A single drive has failed, but the data on that drive can be reconstructed using data and parity from the other drives. The failed drive should be replaced as soon as possible. If they cannot be replaced immediately, back up the data on the LUN to prevent future data loss.
5	Degraded - Warning	A single drive has failed, <i>and</i> one or more drives have been put in warning because of minor errors. The LUN is still accessible. Stop using the LUN immediately, as the failure of a warning drive results in data loss.
5	Drive Failures	Two or more drives have failed and the LUN is inaccessible. All data in the unit is lost.

Note

Refer to the following sections for a further description of drive failures and RAID levels, and the actions to take when they occur.

8.7.3 Drive Failures on RAID 0

This section describes drive failures on a RAID 0 LUN and the LUN and drive status that results from such failures.

8.7.3.1 Single Drive Failure on RAID 0

When a single drive in a RAID 0 LUN experiences unrecoverable read or write errors, the array controller puts the drive in a *Warning* state and changes the LUN status to *Ok - Warning*. This occurs even if the drive is completely inaccessible (for example, if power to the drive is lost). All the data on the LUN may be lost, depending on the severity of the error that generated the warning.

Whenever one or more drives are put in *Warning*, back up the data on the LUN file-by-file (if possible), fail the drives, replace the drives, and then reformat the LUN. For NetWare, you must then add the LUN to your system by creating partitions and volumes (using the NetWare Install command). Copy data back to the restored LUN from your backup media.

8.7.4 Drive Failures on RAID 1

This section describes drive failures on a RAID 1 LUN and the LUN and drive status that result from such failures.

8.7.4.1 Single Drive Failures on RAID 1

Whenever a single drive in an optimal RAID 1 LUN experiences read or write errors, the array controller changes the status of the drive to *Warning* if the errors are minor (read errors on a few sectors) or *Failed* if the errors are serious (loss of drive power, drive component failure, and so on). The LUN status changes to *Ok - Warning* or *Degraded*, depending on whether the drive status is *Warning* or *Failed*. In either case, you can continue to access the LUN, because the data on the failed drive can be read from its mirrored drive.

Whenever a single drive is marked as *Failed* (or *Warning*), you should replace it as soon as possible. The array controller then reconstructs the data on the drive automatically. If you cannot replace the drive immediately, back up the data file-by-file, then replace it when you can.

8.7.4.2 Multiple Drive Failures on RAID 1

Whenever the second drive in a RAID 1 LUN experiences read or write errors, the array controller takes the following actions:

- If the drive is the mirrored pair of a drive that is already failed, the drive status changes to *Warning*, even if the drive is completely inaccessible. The LUN status changes to *Degraded - Warning*, but the LUN may not be accessible.
- If the drive is not the mirrored pair of a failed drive, the drive status changes to *Warning* or *Failed*, depending on the severity of the error. The LUN status changes to *Ok - Warning*, *Degraded*, or *Degraded - Warning*, depending on the previous status of the LUN.

A RAID 1 LUN can have more than one failed drive and still remain degraded as long as none of the failed drives are a mirrored pair. For example, in a six drive RAID 1 LUN, it is possible for as many as three drives to fail while the LUN remains accessible (in degraded mode). You should not, however, continue to operate a RAID 1 LUN with any number of failed or warning drives longer than necessary.

The most serious status in RAID 1 LUN can have is *Degraded - Warning*. This status means that either a mirrored pair of drives has failed, and all data on the LUN has been lost, or that a minor read error has occurred on a drive in an already degraded LUN. To determine the action to take to restore the unit, you should attempt to back up the data.

- If the backup succeeds, the LUN does not contain a failed mirrored pair. You can probably reconstruct the data on the failed and warning drives.
- If the backup fails, a mirrored pair has failed and all data has been lost. You must restore the LUN by reformatting it.

8.7.5 Drive Failures on RAID 5

This section describes drive failures on a RAID 5 LUN and the LUN and drive status that result from such failures.

8.7.5.1 Single Drive Failure on RAID 5

Whenever a single drive in an optimal RAID 5 LUN experiences read or write errors, the array controller changes the status of the drive to *Warning* if the errors are minor (read errors on a few sectors) or *Failed* if the errors are serious (loss of drive power, drive component failure, and so on). The LUN status changes to *Ok - Warning*, or *Degraded*, depending on whether the drive status is *Warning* or *Failed*. In either case, you can continue to access the LUN, as the data on the failed drive can be reconstructed using parity and data on the other drives.

Whenever a drive is marked as *Failed* (or *Warning*), you should replace it as soon as possible. The array controller then reconstructs the data on the drive automatically. If you cannot replace the drive immediately, back up the data file-by-file, then replace it when you can.

8.7.5.2 Multiple Drive Failures on RAID 5

If a second drive in a RAID 5 LUN experiences read or write errors, the array controller marks the drive as *Warning* or *Failed*, depending on the severity of the error. The LUN status changes to *Optimal - Warning*, *Degraded - Warning*, or *Drive Failures*, depending on the previous status of the LUN. If the status changes to *Drive Failures* (because of the failure of a second drive), all data in the LUN is lost.

If the LUN status is *Optimal - Warning*, or *Degraded - Warning*, you probably can reconstruct the data on the failed and warning drives by replacing the drives, one by one, beginning with the failed drive. If the LUN status is *Drive Failures*, you have to replace the failed drives and reformat the LUN.

8.7.6 What is Reconstruction?

Reconstruction is the process that restores a degraded RAID 1 or 5 LUN to its original state after a single drive has been replaced. During reconstruction, the array controller recalculates the data on the drive that was replaced, using data and parity from the other drives in the LUN. The controller then writes this data to the new drive. Although RAID 1 does not have parity, the array controller can reconstruct data on a RAID 1 LUN by copying data from the mirrored disk.

Note

Reconstruction only applies to a degraded RAID 1 or 5 LUN with a single-drive failure.

The array controller automatically initiates the reconstruction process after you replace a drive in a degraded RAID 1 or 5 LUN. Once reconstruction is initiated the array controller completes the following actions:

- Formats the new drive (if the array controller determines it is necessary)
- Copies special array software files to the new drive
- Recalculates the data and parity from the data and parity on the other LUN drives
- Writes the recalculated data and parity to the new drive

Removal and Replacement

8.7 Restoring a LUN

Reconstruction can take place while the LUN is in use. You do not need to shut it down. While the utility is formatting the new drive, there is no additional effect on the system I/O performance (the LUN continues to operate in the degraded mode). Once data reconstruction begins, you can adjust the rate of reconstruction to increase system performance.

If there are multiple-drive failures in a RAID 1 or 5 LUN (that is, the status is *Degraded - Warning* or *Drive Failures*), you cannot reconstruct the drive data. To restore a RAID 1 or 5 LUN to a working state, replace the drives and reformat the LUN. See Section 8.7.4.2 and Section 8.7.5.2 for more information.

8.7.7 Reconstruction Rate

The rate of reconstruction depends on two parameters, the Delay Interval parameter and the Blocks Per Delay Interval parameter.

The Delay Interval is the time between reconstruction operations. When the array controller reconstructs data on a drive, it divides its time between reconstruction and normal I/O operations. System I/O operations take place during this delay interval. As the delay interval increases, system I/O performance also increases, but so does reconstruction time.

The Blocks Per Delay Interval is the number of blocks the array controller reconstructs during each reconstruction operation. The more blocks, the longer the time necessary to reconstruct them. This is time that cannot be used to perform system I/O. Therefore, the larger the reconstruction amount, the more system performance decreases.

These parameters control only the rate of data reconstruction, not total reconstruction time. The first step in reconstruction, formatting the replaced drive, can take 12-15 minutes. During this time, however, you can continue to access the array with no decrease in current system performance (the array remains in the degraded mode).

Some suggested rates, as well as sample timings, are given in Table 8-4.

Table 8-4 Reconstruction Rates

Rate (relative)	Reconstruction Frequency in tenths of a second	Reconstruction Amount in blocks	Time for Completion in minutes/GB ¹
Slow	6	250	20
Moderate	1	256	7
Fast	1	1024	3

¹Timings are approximate and apply only to data reconstruction; formatting time (typically 15-20 minutes) is not included. Times are slower if other processes are accessing the array during reconstruction.

In general, reconstruction times tend to level off after block sizes of 20K. Although higher block sizes (up to 64K) are possible, there is no great reduction in time at the higher levels.

8.7.8 Restoring a RAID 0 LUN

Restore a RAID 0 LUN when one or more of the following occurs:

- A drive has a status of *Warning*.
- The LUN has a status of *Ok - Warning*
- In NetWare, the AMD notifies you of a change in disk or LUN status or NetWare notifies you that it can no longer access the LUN.
- In SCO UNIX, the Array Monitor notifies you of a change in LUN or drive status.
- In OpenVMS VAX, the Array Status Monitor notifies you of a change in LUN or drive status by means of OPCOM.

Use the following procedure to restore a RAID 0 LUN.

1. Back up the data on the LUN, if possible.
2. In NetWare, unmount any volumes on the unit, if they are still mounted.
3. If you modified parameters in MS-DOS or SCO UNIX, or in NetWare and if these parameters are on the LUN that contains any part of your SYS volume, you must use the StorageWorks RAID Array 110 Utility for Novell NetWare. Reboot the system using the StorageWorks RAID Array 110 Utility for Novell NetWare diskette and continue with this procedure.
4. Start ACE by entering one of the following commands from the console:
In the StorageWorks RAID Array 110 Utility for Novell NetWare:
`raidmgr`
In the StorageWorks RAID Array 110 Utility for Novell NetWare:
`load raidmgr`
In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
`$ ACE device-name`
5. In MS-DOS and NetWare only, from the Disk Array Service menu, select *Device Selection*. Scroll through the Disk Array Devices window until you find a LUN with a status of *Ok - Warning*. This should be the RAID 0 LUN you need to restore.
6. In MS-DOS and NetWare only, select that unit. Examine the Array Information and Disk Drive Information windows at the bottom of the screen to make sure that you have selected the correct LUN. The selected unit should be a RAID 0 LUN with one or more drives with a *Warning* status.
7. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
 - a. Replace the drive with the *Warning* status with a new drive, restart the system, and restart RAID Manager.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

Removal and Replacement

8.7 Restoring a LUN

- b. From the *Disk Array Services* menu, select the *Array Controller Configuration* option.
 - c. From the *Array Controller Configuration* menu, select the *Array Controller Editor* option.
 - d. From the Logical Unit window, select *Format* to format the LUN.
 - e. After the reformat is finished, press the Esc key to quit the RAID Manager. Then, press the Y key and the Enter key.
 - f. Remove the diskette and reboot the system.
 - g. Copy backed-up data back to the LUN.
8. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
- a. From the *Disk Array Services* menu, select *Device Service*.
 - b. From the Service Disk Array Device menu, select *Replace the disk drive*. Follow the instructions given to replace the *Warning* drive or drives. Press the Enter key as you complete each step to check the step off.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- c. From the Service Disk Array Device menu, select *Format disk array device*. Answer the prompts to reformat the LUN.
 - d. After the reformat is finished, press the Esc key to quit RAID Manager. Then, press the Y key and the Enter key.
 - e. In the StorageWorks RAID Array 110 Utility for Novell NetWare, use the NetWare Install command to create partitions and volumes on the restored LUN. (See Section 5.2.)
 - f. Copy backed-up data back to the LUN.
9. For the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
- a. Scroll through the LUNs until you find the unit with a status of *Warning*. This should be the RAID 0 LUN that you need to restore. Examining the Drive window, there should be one or more drives with a *Warning* status.
 - b. Replace the drive with the *Warning* status with a new drive.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- c. From the Logical Unit window, select *Format* to format the LUN.
 - d. After the reformat is finished, press the Q key to quit the ACE utility.
 - e. Copy backed-up data back to the LUN.
- You are done with this procedure.

8.7.9 Restoring a RAID 1 LUN

You need to restore a RAID 1 LUN when one or more of the following occurs:

- A drive has a status of *Warning* or *Failed*.
- The LUN has a status of *Ok - Warning*, *Degraded*, or *Degraded - Warning*.
- In NetWare, AMD notifies you of a change in disk or LUN status or NetWare notifies you that it can no longer access the LUN.
- In SCO UNIX, the AMD notifies you of a change in LUN or disk status.
- In OpenVMS VAX, the Array Status Monitor notifies you of a change in LUN or drive status by means of OPCOM.

Use the following procedure to restore a RAID 1 LUN:

1. If the status of the LUN is *Degraded - Warning*, see Section 8.7.10. If the status of the LUN is *Ok - Warning* or *Degraded*, continue with the next step.
2. If you cannot replace the disk right away, attempt to back up the data on the LUN, if possible.
3. To restore a LUN in MS-DOS or SCO UNIX, you must use the StorageWorks RAID Array 110 Utility for Novell NetWare. Reboot the system using the StorageWorks RAID Array 110 Utility for Novell NetWare diskette and continue with this procedure.
4. Start ACE by entering one of the following commands from the console:
In the StorageWorks RAID Array 110 Utility for Novell NetWare:
`raidmgr`
In the StorageWorks RAID Array 110 Utility for Novell NetWare:
`load raidmgr`
In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
`$ ACE device-name`
5. In MS-DOS and NetWare only, from the Disk Array Service menu, select *Device Selection*. Scroll through the Disk Array Devices window until you find a LUN with an *Ok - Warning* status or *Degraded* status. This should be the RAID 1 LUN you need to restore.
6. In MS-DOS and NetWare only, select that unit. Examine the Array Information and Disk Drive Information windows at the bottom of the screen to make sure that you have selected the correct LUN. The selected unit should be a RAID 1 LUN with one or more drives with a *Failed* status.
7. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
 - a. Replace the drive with the *Failed* status with a new drive.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- b. From the *Disk Array Services* menu, select the *Array Controller Configuration* option.

Removal and Replacement

8.7 Restoring a LUN

- c. From the *Array Controller Configuration* menu, select the *Array Configuration Editor* option.
 - d. From the Logical Unit window, select *Reconstruct* to initiate reconstruction of the LUN.
 - e. After the reconstruction is finished, press the Esc key to quit RAID Manager. Press the Y key and the Enter key.
 - f. Remove the diskette and reboot the system.
8. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
- a. From the *Disk Array Services* menu, select *Device Service*.
 - b. Check the Disk Drive Information window. The drive marked *Failed* should be highlighted. If it is not, select *Select Disk Drive* and select the *Failed* drive. If there is more than one drive in *Failed* drive, highlight the first one. If there are *Failed* and *Warning* drives, select a *Failed* one first.
 - c. From the Service Disk Array Device menu, select *Replace the disk drive*.
 - d. Follow the instructions given to replace the *Failed* drives. Replace the highlighted drive, not all the *Failed* or *Warning* drives. Press the Enter key as you complete each step to check the step off.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- e. After you have checked off all the tasks, watch the status of the LUN you just restored. It should change to *Reconstructing*.
If the status does not change to *Reconstructing*, make sure the drive you just replaced is highlighted, then select *Reconstruct disk drive* from the *Service Disk Array Device* menu to start data reconstruction.
 - f. Wait for the reconstruction to finish. After it is finished, if you have any *Failed* or *Warning* drives, repeat these steps to replace and reconstruct them.
 - g. After the LUN status changes to *Ok - Optimal*, press the Esc key to quit RAID Manager. Then, press the Y key and the Enter key.
9. For the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
- a. Scroll through the LUNs until you find the unit with a status of *Warning* or *Degraded*. This should be the RAID 1 LUN that you need to restore. In the Drive window, there should be one or more drives with a *Failed* status.
 - b. Replace the drive with the *Failed* status with a new drive.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- c. From the Logical Unit window, select *Reconstruct* to initiate reconstruction of the LUN.
- d. After the reconstruction is finished, press the **Q** key to quit the ACE utility.

8.7.10 Restoring RAID 1 LUNs after Multiple Failures

After a multiple drive failure, a RAID 1 LUN has a status of *Degraded - Warning*, at least one *Failed* drive, and a mirrored pair of a *Failed* drive with a status of *Warning*. Depending on the severity of the error that caused the *Warning*, the LUN may not be accessible and all the data in it may be lost.

Use the following procedure to restore a RAID 1 LUN after a multiple drive failure:

1. Stop using the LUN immediately if it is still working. Continuing to access the LUN may result in complete loss of data.
2. Attempt to back up the data on the LUN.
 - If the backup was successful, you may be able to reconstruct the data in the LUN. Go to Section 8.7.9, and complete the procedure. You must replace the failed drive first, reconstruct it, and then replace the *Warning* drive. If the reconstruction operation fails, continue with this procedure to restore the LUN.
 - If the backup was not successful, continue with this procedure to restore the LUN.
3. To restore LUN in MS-DOS or SCO, or in NetWare and if the LUN contains any part of your SYS volume, you must use the StorageWorks RAID Array 110 Utility for Novell NetWare. Reboot the system using the StorageWorks RAID Array 110 Utility for Novell NetWare diskette and continue with this procedure.
4. Start ACE by entering one of the following commands from the console:
In the StorageWorks RAID Array 110 Utility for Novell NetWare:
`raidmgr`
In the StorageWorks RAID Array 110 Utility for Novell NetWare:
`load raidmgr`
In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
`$ ACE device-name`
5. From the Disk Array Service menu, select *Device Selection*. Scroll through the Disk Array Devices window until you find a LUN with a status of *Degraded - Warning*. This should be the RAID 1 LUN you need to restore. Select that unit. Examine the Array Information and Disk Drive Information windows at the bottom of the screen to make sure that you have selected the correct LUN. The selected unit should be a RAID 1 LUN with one or more drives with a *Failed* or *Warning* status.
6. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
 - a. Replace the drive with the *Failed* or *Warning* status with a new drive, restart the system, and restart RAID Manager.

Removal and Replacement

8.7 Restoring a LUN

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- b. From the *Disk Array Services* menu, select the *Array Controller Configuration* option.
 - c. From the *Array Controller Configuration* menu, select the *Array Configuration Editor* option.
 - d. From the Logical Unit window, select *Format* to initiate reconstruction of the LUN.
 - e. After the reformat is finished, press the Esc key to quit RAID Manager. Then, press the Y key and the Enter key.
 - f. Remove the diskette and reboot the system.
Copy backed-up data back to the LUN.
7. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
- a. From the *Disk Array Services* menu, select *Device Service*.
 - b. Check the Disk Drive Information window. The drive marked *Warning* should be highlighted. If it is not, select *Select Disk Drive* and select the *Warning* drive. If there is more than one drive in *Warning*, highlight the first one.
 - c. From the Service Disk Array Device menu, select *Replace the disk drive*.
 - d. Follow the instructions given to replace the *Warning* drive or drives. Press the Enter key as you complete each step to check the step off.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- e. From the Service Disk Array Device menu, select *Format disk array device*. Answer the prompts to reformat the LUN.
 - f. After the reformat is finished, press the Esc key to quit RAID Manager. Then, press the Y key and the Enter key.
 - g. Use the NetWare Install command to create partitions and volumes on the restored LUN (see Section 5.2).
Copy backed-up data back to the LUN.
8. For the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
- a. Scroll through the LUNs until you find the unit with a status of *Warning*. This should be the RAID 1 LUN that you need to restore. Examining the Drive window, there should be one or more drives with a *Failed* or *Warning* status.

- b. Replace the drive with the *Failed* or *Warning* status with a new drive, restart the system, and restart the utility.

Caution

It takes 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- c. From the Logical Unit window, select *Format* to initiate reconstruction of the LUN.
- d. After the reformat finishes, press the Q key to quit the ACE utility.
Copy backed-up data back to the LUN.

You are done with this procedure.

8.7.11 Restoring a RAID 5 LUN

Restore a RAID 5 LUN when one or more of the following occurs:

- A drive has a status of *Warning* or *Failed*.
- The LUN has a status of *Ok - Warning*, *Degraded*, or *Degraded - Warning*.
- In NetWare, the Array Monitor Daemon (AMD) notifies you of a change in disk or LUN status or NetWare notifies you that it can no longer access the LUN.
- In SCO UNIX, the Array Monitor notifies you of a change in LUN or drive status.
- In OpenVMS VAX, the Array Status Monitor notifies you of a change in LUN or drive status by means of OPCOM.

Use the following procedure to restore a RAID 5 LUN:

1. If the status of the LUN is *Degraded - Warning*, complete the procedure in Section 8.7.12. If the status of the LUN is *Ok - Warning* or *Degraded*, continue with the next step.
2. If you cannot replace the disk right away, attempt to back up the data on the LUN.
3. To restore a LUN in MS-DOS or SCO UNIX, you must use the StorageWorks RAID Array 110 Utility for Novell NetWare. Reboot the system using the StorageWorks RAID Array 110 Utility for Novell NetWare diskette and continue with this procedure.

4. Start ACE by entering one of the following commands from the console:

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
raidmgr
```

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

```
load raidmgr
```

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

```
$ ACE device-name
```

Removal and Replacement

8.7 Restoring a LUN

5. In MS-DOS and NetWare only, from the Disk Array Service menu, select *Device Selection*. Scroll through the Disk Array Devices window until you find a LUN with an *Ok - Warning* status, *Degraded* or *Degraded - Warning*. This should be the RAID 5 LUN you need to restore.
6. In MS-DOS and NetWare only, select that unit. Examine the Array Information and Disk Drive Information windows at the bottom of the screen to make sure that you have selected the correct LUN. The selected unit should be a RAID 5 LUN with one *Failed* drive and/or one or more drives with a *Warning* status.
7. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
 - a. Replace the drive with the *Warning* status with a new drive.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- b. From the *Disk Array Services* menu, select the *Array Controller Configuration* option.
 - c. From the *Array Controller Configuration* menu, select the *Array Configuration Editor* option.
 - d. From the Logical Unit window, select *Reconstruct* to initiate reconstruction of the LUN.
 - e. After the reconstruction is finished, press the Esc key to quit RAID Manager. Then, press the Y key and the Enter key.
 - f. Remove the diskette and reboot the system.
8. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
 - a. From the *Disk Array Services* menu, select *Device Service*.
 - b. Check the Disk Drive Information window. The drive marked *Failed* should be highlighted. If it is not, select *Select Disk Drive* and select the *Failed* drive. If there is more than one drive in *Failed* drive, highlight the first one. If there are *Failed* and *Warning* drives, select a *Failed* one first.
 - c. From the Service Disk Array Device menu, select *Replace the disk drive*.
 - d. Follow the instructions given to replace the *Failed* drives. Replace the highlighted drive, not all the *Failed* or *Warning* drives. Press the Enter key as you complete each step to check the step off.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- e. After you have checked off all the tasks, watch the status of the LUN you just restored. It should change to *Reconstructing*.
If the status does not change to *Reconstructing*, make sure the drive you just replaced is highlighted, then select *Reconstruct disk drive* from the *Service Disk Array Device* menu to start data reconstruction.

- f. Wait for the reconstruction to finish. After it is finished, if you have any *Failed* or *Warning* drives, repeat these steps to replace and reconstruct them.
 - g. After the LUN status changes to *Ok - Optimal*, press the Esc key to quit RAID Manager. Then, press the Y key and the Enter key.
9. For the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
- a. Scroll through the LUNs until you find the unit with a status of *Ok - Warning*, *Degraded* or *Degraded-Warning*. This should be the RAID 5 LUN that you need to restore. Examining the Drive window, there should be one drive with a *Failed* status and/or one or more drives with a *Warning* status.
 - b. Replace the drive with the *Warning* status with a new drive.

Caution

It takes 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- c. From the Logical Unit window, select *Reconstruct* to initiate reconstruction of the LUN.
- d. After the reconstruction is finished, press the Q key to quit the ACE utility.

8.7.12 Restoring RAID 5 LUNs after Multiple Failures

After a multiple drive failure, a RAID 5 LUN has a status of *Drive Failures* and at least two *Failed* drives. The LUN is not accessible and all the data in it has been lost.

Use the following procedure to restore a RAID 5 LUN after a multiple drive failure:

1. To restore a LUN on MS-DOS or SCO, or in NetWare and the LUN contains any part of your SYS volume, you must use the StorageWorks RAID Array 110 Utility for MS-DOS. Reboot the system using the StorageWorks RAID Array 110 Utility for Novell NetWare diskette and continue with this procedure.
2. Start ACE by entering one of the following commands from the console:

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

`raidmgr`

In the StorageWorks RAID Array 110 Utility for Novell NetWare:

`load raidmgr`

In the StorageWorks RAID Array 110 Utility for OpenVMS VAX:

`$ ACE device-name`
3. For MS-DOS and NetWare only, from the Disk Array Service menu, select *Device Selection*. Scroll through the Disk Array Devices window until you find a LUN with a status of *Degraded - Warning* or *Drive Failures*. This should be the LUN you need to restore.

Removal and Replacement

8.7 Restoring a LUN

4. For MS-DOS and NetWare only, select that unit. Examine the Array Information and Disk Drive Information windows at the bottom of the screen to make sure that you have selected the correct LUN. The selected unit should be a RAID 5 LUN with two or more drives with a *Failed* status.
5. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
 - a. Replace the drives with *Failed* status with new drives.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- b. From the *Disk Array Services* menu, select the *Array Controller Configuration* option.
 - c. From the *Array Controller Configuration* menu, select the *Array Configuration Editor* option.
 - d. From the Logical Unit window, select *Format* to initiate reconstruction of the LUN.
 - e. After the reformat is finished, press the Esc key to quit RAID Manager. Then, press the Y key and the Enter key.
 - f. Remove the diskette and reboot the system.
Copy backed-up data back to the LUN.
6. For the StorageWorks RAID Array 110 Utility for Novell NetWare:
 - a. From the *Disk Array Services* menu, select *Device Service*.
 - b. Check the Disk Drive Information window. The drive marked *Failed* should be highlighted. If it is not, select *Select Disk Drive* and select the first *Failed* drive.
 - c. From the Service Disk Array Device menu, select *Replace the disk drive*.
 - d. Follow the instructions given to replace all the *Failed* drives. Press the Enter key as you complete each step to check the step off.

Caution

Wait 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- e. From the Service Disk Array Device menu, select *Format disk array device*. Answer the prompts to reformat the LUN.
 - f. After the reformat is finished, press the Esc key to quit RAID Manager. Then, press the Y key and the Enter key.
 - g. Use the NetWare Install command to create partitions and volumes on the restored LUN (see Section 5.2).
Copy backed-up data back to the LUN.

7. For the StorageWorks RAID Array 110 Utility for OpenVMS VAX:
 - a. Scroll through the LUNs until you find the unit with a status of *Degraded* - *Warning* or *Drive Failures*. This should be the RAID 5 LUN that you need to restore. Examining the Drive window, there should be two or more drives with a *Failed* status.
 - b. Replace the drives with *Failed* status with new drives.

Caution

It takes 10 seconds from the time you take out a drive to the time that you reinsert a new one for the bus to recognize the drive change.

- c. From the Logical Unit window, select *Format* to initiate reconstruction of the LUN.
 - d. After the reformat is finished, press the Q key to quit the ACE utility.
Copy backed-up data back to the LUN.
- You are done with this procedure.

Adjustments and Alignments

There are no adjustments or alignments for this product.

Configurations and Rules

To do advanced configurations, you need to use one of the StorageWorks RAID Array 110 utilities. The following sections outline the configuration limits and steps needed for reconfigurations. All information is given in a generic sense; specific details are described in the *DEC RAID Utilities User's Guide* for the appropriate operating system.

This chapter discusses the following:

- Modifying basic configurations
- Multiple rank configurations
- Configuration guidelines
- Recommended configurations
- Custom expansions

10.1 Modifying Basic Configurations

If the preconfigured base subsystem does not meet your needs, you can do one or more of the following:

- Modify the RAID level.
- Modify the Drive Mapping and LUN configurations.
- Modify the parameter characteristics such as segment size and reconstruction frequency.

Some restrictions exist with drive selection on different RAID levels:

RAID Level	Drive Map Restrictions
0	Number of drives allowed per LUN = 1 to 5.
1	Number of drives allowed per LUN = 2 to 4. Must specify an even number of drives. The mirrored pair is created by grouping the first and second drive you enter, third and fourth, and so on. Drives on a mirrored pair cannot be on the same channel.
3	Supports 3 or 5 drives.
5	Number of drives allowed per LUN = 3 to 5. Each drive must be on separate channel.
ALL	Drives within a RAID level must be of the same vendor type and capacity.

Configurations and Rules

10.1 Modifying Basic Configurations

Note

It is recommended that you use the maximum number of drives when you create a LUN.

10.2 Multiple Rank Configurations

The StorageWorks RAID Array 110 Subsystem can be expanded up to seven ranks. This means it can support up to 7 disks per channel, or a total of 35 drives. Expansion is done by adding StorageWorks BA350-SA shelves to the base BA350-EA shelf. In addition, RAID level types can be mixed within a configuration. For example, a 4-rank configuration can be built in which two of the ranks are RAID 5 configurations and the other two ranks are RAID 0 configurations. Both the expansion and configuration flexibility allow the user to create new configurations that meet their growing business needs.

Certain expansion paths may require reconfiguration of your array and rebuilding of your data. This is because, with certain expansion paths, you cannot maintain the same array channel and SCSI ID for drives when you migrate from one shelf to multiple shelves. The array channel and SCSI ID information is critical for maintaining the integrity of your array.

Before planning to expand your configuration, review Section 10.3 to understand the RAID configuration and expansion path restrictions.

10.3 Configuration Guidelines

A key feature of the StorageWorks RAID Array 110 Subsystem is that it provides the capability of mixing different RAID level types. However, there are some guidelines and restrictions you must follow when configuring multiple rank systems.

These restrictions are listed below:

RAID Level	Drive Map Restrictions
0	Number of drives allowed per LUN = 1 to 10.
1	Number of drives allowed per LUN = 2 to 10. Must specify an even number of drives. The mirrored pair is created by grouping the first and second drive you enter, third and fourth, and so on. Drives on a mirrored pair cannot be on the same channel.
5	Number of drives allowed per LUN = 3 to 5. Each drive must be on separate channel.
ALL	Drives within a RAID level must be of the same vendor type and capacity.

Recommendations

It is recommended that you configure a LUN with the maximum number of drives (for example, five drives) to prevent performance issues.

It is recommended, with multiple-rank configurations, that disk drives that support the SCSI-3 messages Target Transfer Disable (TTD) and Continue I/O Process (CIOP) are used. This gives the array controller the ability to sequence data transfers across a bus with multiple targets present.

The following are examples of legal and illegal configurations:

LEGAL: Since all LUNs are unique RAID levels.

LOGICAL UNIT	(CHANNEL, SCSI ID)	RAID LEVEL
0	(1,1) (2,1) (3,1) (4,1) (5,1)	5
1	(1,2) (2,2) (3,2) (4,2)	0
2	(1,3) (2,3) (3,3) (4,3) (5,2)	1
3	(1,4) (2,4) (3,4) (4,4) (5,4)	5

ILLEGAL: Since attempting to mix RAID 5 and 0 within LUN 0.

LOGICAL UNIT	(CHANNEL, SCSI ID)	RAID LEVEL
0	(1,1) (2,1) (3,1) (4,1) (5,1)	5,0
1	(1,2) (2,2) (3,2) (4,2)	0
2	(1,3) (2,3) (3,3) (4,3) (5,2)	1
3	(1,4) (2,4) (3,4) (4,4) (5,4)	5

10.3.1 Expansion Guidelines

There are three recommended packaging expansion paths as follows:

- **Single BA350-EA to Single BA350-EA/Dual BA350-SA**
Allows expansion from one rank to a maximum of three ranks of drives.
- **Single BA350-EA to Single BA350-EA/Quad BA350-SA**
Allows expansion from one rank to a maximum of six ranks of drives.
- **Single BA350-EA/Dual BA350-SA to Single BA350-EA/Quad BA350-SA**
Allows expansion from one ranks to a maximum of six ranks of drives.

These paths are recommended because they do not require reconfiguration of your array. For each path, information is provided on the steps required for the expansion. Each of the sections provides the following:

- Shelf requirements
- Instructions
- A shelf reconfiguration diagram
- A drive reconfiguration diagram

Configurations and Rules

10.3 Configuration Guidelines

The shelf reconfiguration diagram is to be used for cable, terminator, and jumper reconfiguration. The drive reconfiguration diagram is to be used for drive movement and recommended rank drive mapping.

Recommendation

It is strongly recommended that you select a path that does not require rebuilding of data.

Refer to Table 10–1 for a listing of the possible expansion paths available to you.

Table 10–1 Expansion Paths

Present	Future	BA350-EA	BA350-SA	Custom Reconfiguration Needed	Refer to:
1 Rank	2 Ranks	1	1	Yes	Section 10.5
1 Rank	2 Ranks	1	2	No	Section 10.4.1
1 Rank	3 Ranks	1	2	No	Section 10.4.1
1 Rank	4 Ranks	1	4	No	Section 10.4.2
1 Rank	5 Ranks	1	4	No	Section 10.4.2
1 Rank	6 Ranks	1	4	No	Section 10.4.2
2 Ranks	3 Ranks	1	3	Yes	Section 10.5
2 Ranks	4 Ranks	1	4	No	Section 10.4.3
2 Ranks	5 Ranks	1	4	No	Section 10.4.3
2 Ranks	6 Ranks	1	4	No	Section 10.4.3
3 Ranks	4 Ranks	1	4	No	Section 10.4.3
3 Ranks	5 Ranks	1	4	No	Section 10.4.3
3 Ranks	6 Ranks	1	4	No	Section 10.4.3
4 Ranks	5 Ranks	1	4	No	Section 10.4.3
4 Ranks	6 Ranks	1	4	No	Section 10.4.3
5 Ranks	6 Ranks	1	4	No	Section 10.4.3

10.4 Recommended Configurations

The following sections describe the recommended expansion paths. Each section describes the following:

- Shelf requirements
- Purpose of the expansion
- Prereconfiguration tasks
- Reconfiguration tasks
- Postreconfiguration tasks

10.4.1 Expansion from a Single BA350-EA to a Single BA350-EA/Dual BA350-SA Configuration

This section describes how to expand your one BA350-EA shelf configuration with two BA350-SA shelves, expanding to a maximum of three ranks of drives. Each rank contains five drives.

Shelf Requirements: One BA350-EA shelf, Two BA350-SA shelves

This expansion allows you to expand to either two ranks of drives (10 drives) or three ranks of drives (15 drives).

Purpose

The single BA350-EA configuration contains one rank of drives (five drives, one per channel). In expanding to a second rank, an additional five drives are added to the configuration (ten drives, two drives per channel). In expanding to a third rank, an additional ten drives are added to the configuration (fifteen drives, three drives per channel). This expansion path requires cabling two additional BA350-SA shelves, and relocating the initial five drives into slots at the same channel and SCSI IDs as in the original configuration.

To expand your configuration, perform the following sets of tasks:

- Prereconfiguration
- Shelf reconfiguration
- Postreconfiguration

10.4.1.1 Prereconfiguration


Before you begin to expand your configuration, perform the following steps:

1. Back up your data to other media.
2. Figure 10–1 shows your current configuration. Use Figure 10–1 as the guide for identification information. Label each of your current drives with array channel and SCSI ID information.

Configurations and Rules

10.4 Recommended Configurations

Figure 10–1 Single BA350-EA Shelf



CXO-3589B-MC

Note

You can label the SBB in pencil. You can erase the SBB later and relabel it as needed.

- Figure 10–2 shows the physical layout of the reconfiguration. Figure 10–3 represents the logical view of the reconfiguration. Study Figure 10–2 and Figure 10–3 to be certain that you understand the cabling and relocation of drives. Refer to Chapter 3 to review connection information.

- Power down the StorageWorks RAID Array 110 Subsystem.


You are ready to begin the shelf reconfiguration described in Section 10.4.1.2.

10.4.1.2 Shelf Reconfiguration

To expand your configuration, perform the following steps:

- Reconfigure your BA350-EA shelf and BA350-SA shelves using the diagram shown in Figure 10–2.

Figure 10-2 Single BA350-EA to Single BA350-EA/Dual BA350-SA Shelf Reconfiguration



CXO-3650B-MC

To begin, the BA350-EA shelf is reconfigured by replacing some of the terminators with jumpers, as described below. The jumpers and terminators are located at the back of the unit on the backplane.

- Add the jumpers at SCSI ID slots 2 and 3. This connects those SCSI IDs to channel 5.
- Add the jumper at SCSI ID slot 5. This connects SCSI IDs 4, 5, and 0 to channel 2. Note that SCSI IDs 5 and 0 are connected through the backplane and no jumper is required.

Configurations and Rules

10.4 Recommended Configurations

Refer to Figure 10-3. The channel and SCSI ID are designated as (x,y) or (channel, SCSI ID). For example, channel 5, SCSI ID 1 is designated as (5,1).

The BA350-EA shelf now has channel and SCSI IDs (5,1), (5,2), (5,3), (2,4), (2,5), (2,0). The jumpers at SCSI ID slots 2, 3, and 5 have freed up channels 4, 3, and 1 to be routed to the BA350-SA expansion shelves.

The BA350-SA expansion shelves contain eight slots and can hold up to seven drives. Starting at the top of the shelf, the SCSI IDs are 0 to 6. The BA350-SA shelf can be configured as either one SCSI bus with IDs 0 through 6, or as two separate SCSI buses with the following IDs on each bus:

- SCSI IDs 1, 3, 5, and 6
- SCSI IDs 0, 2, and 4

If a terminator is placed in slot 5 (SCSI ID 5), then the SCSI buses will be separate. Otherwise, if you put a jumper in slot 5, then the SCSI IDs are all connected together.

1. For purposes of this configuration, put a terminator in slot 5 as shown in Figure 10-2. As you continue with the reconfiguration, be certain to follow the pattern of jumpers and terminators shown in Figure 10-2.
2. Using Figure 10-3 as a guide, perform the following steps:

Figure 10-3 Single BA350-EA to Single BA350-EA/Dual BA350-SA Drive Reconfiguration

350-E	350-E	350-S	350-S
(5,1) 1	(5,1) 1	ADD 0	0
(4,2) 2	ADD 2	ADD 1	ADD 1
(3,3) 3	ADD 3	(4,2) 2	2
(2,4) 4	(2,4) 4	ADD 3	(3,3) 3
(1,5) 5	ADD 5	ADD 4	4
0	ADD 0	(1,5) 5	ADD 5
		6	6

Rank 1: (5,1) (4,2) (3,3) (2,4) (1,5)

Rank 2: (5,2) (4,4) (3,5) (2,5) (1,3)

Rank 3: (5,3) (4,0) (3,1) (2,0) (1,1)

(x,y) = (Channel, SCSI Id)

SHR-XR3016-GRA


- a. Connect the BA350-SA shelves to the BA350-EA shelf.
 - The section of the BA350-EA shelf that contains the HSZ10-AA controller contains a column of six 50-pin, high-density connectors on the right-hand side. These connectors are labeled as channels in Figure 10-4. The top connector is used exclusively for connecting to host adapter SSBs (for example, the DWZZA-VA) in slot 0; this connector is unnumbered in Figure 10-4. The other connectors are labeled as 5, 4, 3, 2, and 1.

Configurations and Rules

10.4 Recommended Configurations

- The BA350-SA shelf contains two 50-pin, high-density connectors on the backplane at the top of the shelf. If you look at the shelf from the front, the left-hand connector connects SCSI IDs 0, 2, and 4. The right-hand connector connects SCSI IDs 1, 3, and 5. As previously discussed, these SCSI buses are independent of each other unless you put a jumper in slot 5.
 - To connect a channel to the BA350-SA shelf, run a cable from the BA350-EA shelf to the BA350-SA shelf. Use a BN21H-Series cable, which has two 50-pin, high-density, male, straight connectors.
- b. For this configuration, connect channels 1, 3, and 4 from the BA350-EA shelf to the BA350-SA shelves as follows:
- Attach a BN21H cable from channel 1 of the BA350-EA shelf to the right-hand connector in the BA350-SA shelf. This connects channel 1 with the SCSI IDs 1, 3, and 5, as shown in Figure 10-4.

Figure 10-4 BA350-EA and BA350-SA Combination Storage Array



SHR-XR3029-GRA

Configurations and Rules

10.4 Recommended Configurations

- Attach a cable in channel 4 of the BA350-EA shelf to the left-hand connector in the BA350-SA shelf. This connects channel 4 to SCSI IDs 0, 2, and 4.
- Connect channel 3 to the second BA350-SA shelf by attaching a cable to channel 3 in the BA350-EA shelf to the right-hand connector of the second BA350-SA shelf (SCSI IDs 1, 3, and 5).

Refer to Figure 10-2.

- c. Insert the original five drives from the BA350-EA shelf into the BA350-SA shelf locations with the same channel and SCSI ID as originally assigned to the drive. This allows the operating system to recognize the original LUNs immediately. The original channel and SCSI ID assignments were (5,1), (4,2), (3,3), (2,4), and (1,5).

In the new configuration, drives (5,1) and (2,4) are still in the BA350-EA shelf, while drives (4,2), (3,3), and (1,5) are relocated to the BA350-SA shelves. Refer to Figure 10-3.

3. To expand to two ranks, add an additional five drives in the slot locations designated by the rank-2 list.

These drives must be inserted in the following locations: (5,2) and (2,5) in the BA350-EA shelf, and (4,4), (3,5), and (1,3) in the BA350-SA shelves. Refer to Figure 10-3.

4. To expand to three ranks, add an additional five drives in the slot locations designated by the rank-3 list.

These drives must be inserted in the following locations: (5,3) and (2,0) in the BA350-EA shelf, and (4,0), (3,1), and (1,1) in the BA350-SA shelves. Refer to Figure 10-3.

2. Power on the StorageWorks RAID Array 110 Subsystem.

You have completed the shelf reconfiguration. Verify that you have correctly reconfigured your BA350-EA with two BA350-SA shelves as described in Section 10.4.1.3.

10.4.1.3 Postreconfiguration

To verify your configuration, use StorageWorks RAID Array 110 utilities appropriate for your operating system to perform the following steps:

1. Verify that the original drives are all set to optimal condition.
2. Verify that the new drives are displayed as "spare" drives in the proper array channel/SCSI ID positions.
3. Configure and format the desired RAID level configurations.
4. Reboot the system.

Information on using the StorageWorks RAID Array 110 utilities is contained in Chapter 4.

10.4.2 Expansion from a Single BA350-EA to a Single BA350-EA/Quad BA350-SA Configuration

This section describes how to expand your one BA350-EA shelf configuration with four BA350-SA shelves, expanding to a maximum of six ranks of drives. Each rank contains five drives.

Shelf Requirements: One BA350-EA shelf, Four BA350-SA shelves

This expansion allows you to expand from two ranks of drives (10 drives) to up to six ranks of drives (30 drives).

Purpose

The single BA350-EA configuration contains one rank of drives (five drives, one per channel). Expansion involves adding additional ranks of five drives to the configuration. This expansion path requires cabling four additional BA350-SA shelves, and relocating the initial five drives into slots at the same channel and SCSI IDs as in the original configuration.

To expand your configuration, perform the following sets of tasks:

- **Prereconfiguration**
- **Shelf reconfiguration**
- **Postreconfiguration**

10.4.2.1 Prereconfiguration


Before you begin to expand your configuration, perform the following steps:

1. Back up your data to other media.
2. Figure 10–5 shows your current configuration. Use Figure 10–5 as the guide for identification information. Label each of your current drives with array channel and SCSI ID information.

Configurations and Rules

10.4 Recommended Configurations

Figure 10–5 Single BA350-EA Shelf



CXO-3589B-MC

Note

You can label the SBB in pencil. You can erase the SBB later and relabel it as needed.

3. Figure 10–6 shows the physical layout of the reconfiguration. Figure 10–7 represents the logical view of the reconfiguration. Study Figure 10–6 and Figure 10–7 to be certain that you understand the cabling and relocation of drives. Refer to Chapter 3 to review connection information.

4. Power down the StorageWorks RAID Array 110 Subsystem.

You are ready to begin the shelf reconfiguration described in Section 10.4.2.2.

10.4.2.2 Shelf Reconfiguration


To expand your configuration, perform the following steps:

1. Reconfigure your BA350-EA shelf and BA350-SA shelves using the diagram shown in Figure 10–6.

Configurations and Rules

10.4 Recommended Configurations

Figure 10-6 Single BA350-EA to Single BA350-EA/Quad BA350-SA Shelf Reconfiguration



To begin, the BA350-EA shelf is reconfigured by replacing some of the terminators with jumpers, as described below. The jumpers and terminators are located at the back of the unit on the backplane.

- Add the jumpers at SCSI ID slots 2, 3, 4, and 5. This connects SCSI IDs 1, 2, 3, 4, 5, and 0 to channel 5. Note that SCSI IDs 5 and 0 are connected through the backplane and no jumper is required.

Refer to Figure 10-7. The channel and SCSI ID are designated as (x,y) or (channel, SCSI ID). For example, channel 5, SCSI ID 1 is designated as (5,1).

Configurations and Rules

10.4 Recommended Configurations

The BA350-EA shelf now has channel and SCSI IDs (5,1), (5,2), (5,3), (5,4), (5,5), (5,0). The jumpers at SCSI ID slots 2, 3, 4, and 5 have freed up channels 4, 3, 2, and 1 to be routed to the BA350-SA expansion shelves.

The BA350-SA expansion shelves contain eight slots and can hold up to seven drives. Starting at the top of the shelf, the SCSI IDs are 0 to 6. The BA350-SA shelf can be configured as either one SCSI bus with IDs 0 through 6, or as two separate SCSI buses with the following IDs on each bus:

- SCSI IDs 1, 3, 5, and 6
- SCSI IDs 0, 2, and 4

If a terminator is placed in slot 5 (SCSI ID 5), then the SCSI buses will be separate. Otherwise, if you put a jumper in slot 5, then the SCSI IDs are all connected together.

1. For purposes of this configuration, put a terminator in slot 5 as shown in Figure 10-6. As you continue with the reconfiguration, be certain to follow the pattern of jumpers and terminators shown in Figure 10-6.
2. Using Figure 10-7 as a guide, perform the following steps:

Figure 10-7 Single BA350-EA to Single BA350-EA/Quad BA350-SA Drive Reconfiguration

350-E	350-E	350-S	350-S	350-S	350-S
(5,1) 1	(5,1) 1	0	0	0	0
(4,2) 2	ADD 2	ADD 1	ADD 1	ADD 1	ADD 1
(3,3) 3	ADD 3	(4,2) 2	ADD 2	ADD 2	ADD 2
(2,4) 4	ADD 4	ADD 3	(3,3) 3	ADD 3	ADD 3
(1,5) 5	ADD 5	ADD 4	ADD 4	(2,4) 4	ADD 4
0	0	ADD 5	ADD 5	ADD 5	(1,5) 5
		6	6	6	6

Rank 1: (5,1) (4,2) (3,3) (2,4) (1,5)

Rank 2: (5,2) (4,1) (3,1) (2,1) (1,2)

Rank 3: (5,3) (4,3) (3,2) (2,2) (1,3)

Rank 4: (5,4) (4,4) (3,4) (2,3) (1,4)

Rank 5: (5,5) (4,5) (3,5) (2,4) (1,1)

Rank 6: (5,0) (4,0) (3,0) (2,0) (1,0)

(x,y) = (Channel, SCSI Id)

SHR-XR3017-GRA

- Connect the BA350-SA shelves to the BA350-EA shelf.
 - The section of the BA350-EA shelf that contains the HSZ10-AA controller contains a column of six 50-pin, high-density connectors on the right-hand side. These connectors are labeled as channels in Figure 10-5. The top connector is used exclusively for connecting to host adapter SSBs (for example, the DWZZA-VA) in slot 0; this connector is unnumbered in Figure 10-5. The other connectors are labeled as 5, 4, 3, 2, and 1.

- The BA350-SA shelf contains two 50-pin, high-density connectors on the backplane at the top of the shelf. If you look at the shelf from the front, the left-hand connector connects SCSI IDs 0, 2, and 4. The right-hand connector connects SCSI IDs 1, 3, and 5. As previously discussed, these SCSI buses are independent of each other unless you put a jumper in slot 5.
 - To connect a channel to the BA350-SA shelf, run a cable from the BA350-EA shelf to the BA350-SA shelf. Use a BN21H-Series cable, which has two 50-pin, high-density, male, straight connectors.
- b. For this configuration, connect channels 1, 2, 3, and 4 from the BA350-EA shelf to the BA350-SA shelves as follows:
- Attach a BN21H cable from the channel in the BA350-EA shelf to the right-hand connector in the BA350-SA shelf. This connects that channel with the SCSI IDs 0, 1, 2, 3, 4, and 5. For example, to connect channel 1 in the BA350-EA shelf, attach the BN21H cable from channel 1 to the right-hand connector of the BA350-SA shelf.

Refer to Figure 10–6.

- c. Insert the original five drives from the BA350-EA shelf into the BA350-SA shelf locations with the same channel and SCSI ID as originally assigned to the drive. This allows the operating system to recognize the original LUNs immediately. The original channel and SCSI ID assignments were (5,1), (4,2), (3,3), (2,4), and (1,5).

In the new configuration, drive (5,1) is still in the BA350-EA shelf, while drives (4,2), (3,3), (2,4) and (1,5) are relocated each to a BA350-SA shelf. Refer to Figure 10–7.

- d. To expand to two ranks, add an additional five drives in the slot locations designated by the rank-2 list.
- These drives must be inserted in the following locations: (5,2) in the BA350-EA shelf, and (4,1), (3,1), (2,1), and (1,2) in each of the BA350-SA shelves. Refer to Figure 10–7.
- e. To expand to three ranks, add an additional five drives in the slot locations designated by the rank-3 list.
- These drives must be inserted in the following locations: (5,3) in the BA350-EA shelf, and (4,3), (3,2), (2,2), and (1,3) in each of the BA350-SA shelves. Refer to Figure 10–7.
- f. To expand to four ranks, add an additional five drives in the slot locations designated by the rank-4 list.
- These drives must be inserted in the following locations: (5,4) in the BA350-EA shelf, and (4,4), (3,4), (2,3), and (1,1) in each of the BA350-SA shelves. Refer to Figure 10–7.
- g. To expand to five ranks, add an additional five drives in the slot locations designated by the rank-5 list.
- These drives must be inserted in the following locations: (5,5) in the BA350-EA shelf, and (4,5), (3,5), (2,4), and (1,1) in each of the BA350-SA shelves. Refer to Figure 10–7.

Configurations and Rules

10.4 Recommended Configurations

- h. To expand to six ranks, add an additional five drives in the slot locations designated by the rank-6 list.

These drives must be inserted in the following locations: (5,0) in the BA350-EA shelf, and (4,0), (3,0), (2,0), and (1,0) in each of the BA350-SA shelves. Refer to Figure 10-7.

3. Power on the StorageWorks RAID Array 110 Subsystem.

You have completed the shelf reconfiguration. Verify that you have correctly reconfigured your BA350-EA with four BA350-SA shelves as described in Section 10.4.2.3.

10.4.2.3 Postreconfiguration

To verify your configuration, use the StorageWorks RAID Array 110 utilities appropriate for your operating system to perform the following steps:

1. Verify that the original drives are all set to optimal condition.
2. Verify that the new drives are displayed as “spare” drives in the proper array channel/SCSI ID positions.
3. Configure and format the desired RAID level configurations.
4. Reboot the system.

Information on using the StorageWorks RAID Array 110 utilities is contained in Chapter 4.

10.4.3 Expansion from a Single BA350-EA/Dual BA350-SA to a Single BA350-EA/Quad BA350-SA Configuration

This section describes how to expand your one BA350-EA shelf with two BA350-SA shelves to four BA350-SA shelves, expanding from three ranks to a maximum of six ranks of drives. Each rank contains five drives.

Shelf Requirements: One BA350-EA shelf, Four BA350-SA shelves

This expansion allows you to expand from three ranks of drives (15 drives) to up to six ranks of drives (30 drives).

Purpose

The single BA350-EA shelf with two BA350-SA shelves configuration contains three ranks of drives (15 drives). Expansion involves adding additional ranks of five drives each to the configuration. The single BA350-EA/dual BA350-SA configuration can contain up to three ranks of drives (15 drives). This expansion path requires cabling two additional BA350-SA shelves, and relocating the initial five drives into slots at the same channel and SCSI IDs as in the original configuration.

To expand your configuration, perform the following sets of tasks:

- Prereconfiguration
- Shelf reconfiguration
- Postreconfiguration

10.4.3.1 Prereconfiguration


Before you begin to expand your configuration, perform the following steps:

1. Back up your data to other media.
2. Figure 10–8 shows your current configuration. Use Figure 10–8 as the guide for identification information. Label each of your current drives with array channel and SCSI ID information.

Configurations and Rules

10.4 Recommended Configurations

Figure 10–8 Single BA350-EA Shelf



CXO-3589B-MC

Note

You can label the SBB in pencil. You can erase the SBB later and relabel it as needed.

- Figure 10–9 shows the physical layout of the reconfiguration. Figure 10–10 represents the logical view of the reconfiguration. Study Figure 10–9 and Figure 10–10 to be certain that you understand the cabling and relocation of drives. Refer to Chapter 3 to review connection information.

- Power down the StorageWorks RAID Array 110 Subsystem.


You are ready to begin the shelf reconfiguration described in Section 10.4.3.2.

10.4.3.2 Shelf Reconfiguration

To expand your configuration, perform the following steps:

- Reconfigure your BA350-EA shelf and BA350-SA shelves using the diagram shown in Figure 10–9.

Figure 10–9 Single BA350-EA/Dual BA350-SA Shelf to Single BA350-EA/Quad BA350-SA Shelf Reconfiguration



CXO-3594B-MC

To begin, the BA350-EA shelf is reconfigured by replacing some of the terminators with jumpers, as described below. The jumpers and terminators are located at the back of the unit on the backplane.

- Add the jumpers at SCSI ID slots 2, 3, 4, and 5. This connects SCSI IDs 1, 2, 3, 4, 5, and 0 to channel 5. Note that SCSI IDs 5 and 0 are connected through the backplane and no jumper is required.

Refer to Figure 10–10. The channel and SCSI ID are designated as (x,y) or (channel, SCSI ID). For example, channel 5, SCSI ID 1 is designated as (5,1).

Configurations and Rules

10.4 Recommended Configurations

The BA350-EA shelf now has channel and SCSI IDs (5,1), (5,2), (5,3), (5,4), (5,5), (5,0). The jumpers at SCSI ID slots 2, 3, 4, and 5 have freed up channels 4, 3, 2, and 1 to be routed to the BA350-SA expansion shelves.

The BA350-SA expansion shelves contain eight slots and can hold up to seven drives. Starting at the top of the shelf, the SCSI IDs are 0 to 6. The BA350-SA shelf can be configured as either one SCSI bus with IDs 0 through 6, or as two separate SCSI buses with the following IDs on each bus:

- SCSI IDs 1, 3, 5, and 6
- SCSI IDs 0, 2, and 4

If a terminator is placed in slot 5 (SCSI ID 5), then the SCSI buses will be separate. Otherwise, if you put a jumper in slot 5, then the SCSI IDs are all connected together.

1. For purposes of this configuration, put a terminator in slot 5 as shown in Figure 10-9. As you continue with the reconfiguration, be certain to follow the pattern of jumpers and terminators shown in Figure 10-9.
2. Using Figure 10-10 as a guide, perform the following steps:

Figure 10-10 Single BA350-EA/Dual BA350-SA Shelf Configuration to Single BA350-EA/Quad BA350-SA Drive Reconfiguration

350-E	350-S	350-S	350-E	350-S	350-S	350-S	350-S
	0	(1,0) 0		(1,0) 0	(2,0) 0	ADD 0	ADD 0
(5,1) 1	(3,1) 1	(4,1) 1	(5,1) 1	ADD 1	ADD 1	(3,1) 1	(4,1) 1
(5,2) 2	2	(1,2) 2	(5,2) 2	(1,2) 2	ADD 2	ADD 2	ADD 2
(5,3) 3	(3,3) 3	(4,3) 3	(5,3) 3	ADD 3	ADD 3	(3,3) 3	(4,3) 3
(2,4) 4	4	(1,4) 4	ADD 4	(1,4) 4	(2,4) 4	ADD 4	ADD 4
(2,5) 5	(3,5) 5	(4,5) 5	ADD 5	ADD 5	(2,5) 5	(3,5) 5	(4,5) 5
(2,0) 0	6	6	ADD 0	6	6	6	6

Rank 1: (5,1) (4,1) (3,1) (2,4) (1,0)

Rank 2: (5,2) (4,3) (3,3) (2,5) (1,2)

Rank 3: (5,3) (4,5) (3,5) (2,0) (1,4)

Rank 4: (5,4) (4,4) (3,4) (2,3) (1,5)

Rank 5: (5,5) (4,2) (3,2) (2,2) (1,3)

Rank 6: (5,0) (4,0) (3,0) (2,1) (1,1)

(x,y) = (Channel, SCSI Id)

SHR-XR3018-GRA

- a. Connect the BA350-SA shelves to the BA350-EA shelf.
 - The section of the BA350-EA shelf that contains the HSZ10-AA controller contains a column of six 50-pin, high-density connectors on the right-hand side. These connectors are labeled as channels in Figure 10-8. The top connector is used exclusively for connecting to host adapter SSBs (for example, the DWZZA-VA) in slot 0; this connector is unnumbered in Figure 10-8. The other connectors are labeled as 5, 4, 3, 2, and 1.

- The BA350-SA shelf contains two 50-pin high-density connectors on the backplane at the top of the shelf. If you look at the shelf from the front, the left-hand connector connects SCSI IDs 0, 2, and 4. The right-hand connector connects SCSI IDs 1, 3, and 5. As previously discussed, these SCSI buses are independent of each other unless you put a jumper in slot 5.
 - To connect a channel to the BA350-SA shelf, run a cable from the BA350-EA shelf to the BA350-SA shelf. Use a BN21H-Series cable, which has two 50-pin, high-density, male, straight connectors.
- b. For this configuration, connect channels 1, 2, 3, and 4 from the BA350-EA shelf to the BA350-SA shelves as follows:
- Attach a BN21H cable from the channel in the BA350-EA shelf to the right-hand connector in the BA350-SA shelf. This connects that channel with the SCSI IDs 0, 1, 2, 3, 4, and 5. For example, to connect channel 1 in the BA350-EA shelf, attach the BN21H cable from channel 1 to the right-hand connector of the BA350-SA shelf.

Refer to Figure 10–9.

- c. Insert the original drives in the BA350-EA and BA350-SA shelves in locations with the same channel and SCSI ID as originally assigned to the drive. This allows the operating system to immediately recognize the original LUNs. The original channel and SCSI ID assignments were as follows:
- The BA350-EA shelf contained (5,1), (5,2), (5,3), (2,4), (2,5), (2,0).
 - One of the BA350-SA shelves contained (3,1), (3,3), and (3,5).
 - The other BA350-SA shelf contained (1,0), (4,1), (1,2), (4,3), (1,4), and (4,5).
- In the new configuration, drives (5,1), (5,2), and (5,3) are still in the BA350-EA shelf, while drives (1,0), (1,2), (1,4) are relocated to the first BA350-SA shelf, drives (2,0), (2,4), (2,5) to the second BA350-SA shelf, drives (3,1), (3,3), (3,5) to the third BA350-SA shelf, and drives (4,1), (4,3), (4,5) to the fourth BA350-SA shelf.
- d. To expand to four ranks, add an additional five drives in the slot locations designated by the rank-4 list.
- These drives must be inserted in the following locations: (5,4) in the BA350-EA shelf, and (4,4), (3,4), (2,3), and (1,5) in each of the BA350-SA shelves. Refer to Figure 10–10.
- e. To expand to five ranks, add an additional five drives in the slot locations designated by the rank-5 list.
- These drives must be inserted in the following locations: (5,5) in the BA350-EA shelf, and (4,2), (3,2), (2,2), and (1,3) in each of the BA350-SA shelves. Refer to Figure 10–10.
- f. To expand to six ranks, add an additional five drives in the slot locations designated by the rank-6 list.
- These drives must be inserted in the following locations: (5,0) in the BA350-EA shelf, and (4,0), (3,0), (2,1), and (1,1) in each of the BA350-SA shelves. Refer to Figure 10–10.

Configurations and Rules

10.4 Recommended Configurations

3. Power on the StorageWorks RAID Array 110 Subsystem.

You have completed the shelf reconfiguration. Verify that you have correctly reconfigured your BA350-EA with four BA350-SA shelves as described in Section 10.4.3.3.

10.4.3.3 Postreconfiguration

To verify your configuration, use the StorageWorks RAID Array 110 utilities appropriate for your operating system to perform the following steps:

1. Verify that the original drives are all set to optimal condition.
2. Verify that the new drives are displayed as “spare” drives in the proper array channel/SCSI ID positions.
3. Configure and format the desired RAID level configurations.
4. Reboot the system.

Information on using the StorageWorks RAID Array 110 utilities is contained in Chapter 4.

10.5 Custom Expansions

This section outlines the steps used for expanding a configuration that requires reconfiguration of the array and the rebuilding of user data. Refer to Table 10–1 for details on which expansion paths this includes.

To perform custom expansions, do the following:

Caution

It is very important that you back up your data to other media, since this procedure destroys your data.

1. Back up your data to other media.
2. Using StorageWorks RAID Array 110 utilities, perform the following steps:
 - a. Delete all configured LUNs. This makes all drives spares.
 - b. Reconfigure the drive in array channel 5, SCSI ID 1 to be a 1 drive RAID 0 array.
3. Power down your StorageWorks RAID Array 110 Subsystem.
4. Refer to Table 10–1 for the appropriate section for reconfiguring your shelves.
5. Power on your StorageWorks RAID Array 110 Subsystem.

6. Using the StorageWorks RAID Array 110 utilities, perform the following steps:
 - a. Delete LUN 0 (LUN 0 contains the 1 drive in Channel 5, SCSI ID 1) configured in Step 2.
 - b. Configure and format desired RAID level configurations. To maintain your data, you must conform to the drive mappings detailed in Sections 10.4.1 through 10.4.3. When you elect to use the customized expansion technique detailed here, your data is *not* maintained. Therefore, you do not need to conform to the drive mappings shown in Sections 10.4.1 through 10.4.3 because you are using this customized expansion technique.

Parts Information

This chapter provides information regarding the part numbers and descriptions for the StorageWorks RAID Array 110 Subsystem product including:

- Hardware
- Firmware
- Software

11.1 Hardware

The tables in this chapter list the part numbers and descriptions for all StorageWorks products. If applicable, the number of system building block (SBB) slots that a device occupies in a shelf is also listed. One SBB slot accommodates one 3½-inch device. Three SBB slots are required for one 5¼-inch device. A BA350-SA shelf has eight SBB slots.

Hardware parts include the following:

- Cabinets
- Shelves
- SBBs
- Accessories
- Cables

Table 11-1 Cabinets

Part Number	Description
BA35X-VA	Single-width vertical mounting kit. A double-wide vertical shelf unit is made by mounting two BA35X-VA kits onto a BA350-EA shelf.
H9A02-AD	Data center cabinet—120/208 V ac, 3 Ph, 50/60 Hz. A generic (empty), 800-mm cabinet that includes outer panels, blank doors, power entry units, and cooling systems, but no mounting hardware, cables, or internal power cords.
H9A02-AE	Data center cabinet—240/416 V ac, 3 Ph, 50/60 Hz. A generic (empty), 800-mm cabinet that includes outer panels, blank doors, power entry units, and cooling systems, but no mounting hardware, cables, or internal power cords.

Parts Information

11.1 Hardware

Table 11-2 BA350 Shelves

Part Number	Description
BA350-EA	A special purpose, double-height shelf used for a maximum of two controllers, six 3½-inch SBBs or two 5¼-inch SBBs, four power supplies, and six SCSI-2 buses.
BA350-LA	A storage shelf kit that includes a BA350-SA storage shelf, an S-PIU rack mounting kit, and a dc power supply. This shelf is for use in DEC 10000 AXP and DEC 7000 AXP systems and cannot be used in any other cabinet. This shelf can be configured for either one or two SCSI-2 buses. Its maximum capacity is seven 3½-inch SBBs, or two 5¼-inch SBBs and one 3½-inch SBB.
BA350-SA	A basic, single-height modular storage shelf that can be configured for either one or two SCSI-2 buses. Its maximum capacity is seven 3½-inch SBBs, or two 5¼-inch SBBs and one 3½-inch SBB.

Table 11-3 SBBs

Part Number	SBB Slots	Description
Power Supplies		
BA350-HA	1	131-watt universal ac input power supply; outputs +5 V and +12 V
BA350-HB	1	131-watt universal 48 V dc input power supply; outputs +5 V and +12 V
Controllers		
HSZ10-AA	NA	RAID-ready controller that fits into the BA350-EA shelf
Disk Drives		
RZ25-VA	1	3½-inch, 0.426-GB, fixed disk drive in a modular carrier
RZ26-VA	1	3½-inch, 1.05-GB, fixed disk drive in a modular carrier
RZ73-VA	3	5¼-inch, 2-GB, fixed disk drive in a modular carrier
Tape Drives		
TLZ06-VA	1	3½-inch, 4-mm, 4-GB DAT tape drive mounted in a modular carrier

Table 11-4 Accessories

Part Number	Description
Filler Panels	
BA35X-PA	A kit of six panels used with the BA350-LA shelf module to improve airflow. These panels are not required with any other shelf.
Replacement Fans	
BA35X-MA	Single replacement fan carrier assembly for BA350 shelves
Power Cords for BA35X-VA Vertical Mounting Kit	
BN18L-2E	Israel 220 V ac power cord (2.5 m)
BN19A-2E	United Kingdom/Ireland 240 V ac power cord (2.5 m)
BN19E-2E	Switzerland 220 V ac power cord (2.5 m)
BN19H-2E	Australia/New Zealand 220 V ac power cord (2.5 m)
BN19K-2E	Denmark 220 V ac power cord (2.5 m)
BN19M-2E	Italy 220 V ac power cord (2.5 m)
BN19S-2E	India/South America 220 V ac power cord (2.5 m)
BN19W-2E	Central Europe 220 V ac power cord (2.5 m)
BN27S-03	United States/Japan 120 V ac power cord, (3 m)
BA350 Backplane SCSI Components	
BA35X-MB	Spare SCSI terminator board for BA350 shelves
BA35X-MC	Spare SCSI jumper board for BA350 shelves
HSZ10-AA Controller to KZESA Adapter Accessories	
H879-AA	68-pin differential terminator for use with a Y cable.
H885-AA	A trilink connector block that has two 68-pin, female connectors on one side and a 68-pin, male connector on the other. The trilink connector block enables an HS-series controller to be removed from the SCSI bus without degrading the bus.

Parts Information

11.1 Hardware

Table 11-5 Cables

Part Number	Description
BN21H-Series SCSI A Cables	
BN21H-0C	0.3 m (0.98 ft); SCSI-2, high-density, 50-conductor cable; male standard straight to male standard straight connectors
BN21H-0E	0.5 m (1.6 ft)
BN21H-01	1.0 m (3.3 ft)
BN21H-0E	1.5 m (4.8 ft)
BN21H-02	2.0 m (6.5 ft)
BN21H-03	3.0 m (9.8 ft)
BN21H-05	5.0 m (16.4 ft)
BN21H-01	10.0 m (32.8 ft)
BN21H-15	15.0 m (49.2 ft)
BN21H-20	20.0 m (65.6 ft)
BN21J-Series SCSI A Cables	
BN21J-0E	0.5 m (1.6 ft); SCSI-2, high-density, 50-conductor cable; male standard straight connector on one end; male special right-angle connector (using jackscrews) on the other
BN21J-01	1.0 m (3.3 ft)
BN21J-02	2.0 m (6.5 ft)
BN21J-03	3.0 m (9.8 ft)
BN21J-05	5.0 m (16.4 ft)
BN21J-10	10.0 m (32.8 ft)
BN21J-15	15.0 m (49.2 ft)
BN21J-20	20.0 m (65.6 ft)
BN21K-Series SCSI P Cables	
BN21K-02	2.0 m (6.5 ft); SCSI-3, standard, 68-conductor P cable; male standard straight connector on one end; male standard right-angle connector on the other; typically used to connect an HSZ10-AA controller to a KZTSA host adapter
BN21K-03	3.0 m (9.8 ft)
BN21K-05	5.0 m (16.4 ft)
BN21K-10	10.0 m (32.8 ft)
BN21K-15	15.0 m (49.2 ft)
BN21K-20	20.0 m (65.6 ft)

(continued on next page)

Table 11-5 (Cont.) Cables

Part Number	Description
BN21L-Series SCSI P Cables	
BN21L-0B	0.15 m (.49 ft); SCSI-3, standard, 68-conductor P cable; male standard right-angle connector on both ends; typically used to connect one HSZ10-AA controller to another HSZ10-AA controller
BN21L-02	2.0 m (6.5 ft)
BN21L-03	3.0 m (9.8 ft)
BN21L-05	5.0 m (16.4 ft)
BN21L-10	10.0 m (32.8 ft)
BN21L-15	15.0 m (49.2 ft)
BN21L-20	20.0 m (65.6 ft)
BN21M-Series SCSI Transition Cables	
BN21M-02	2.0 m (6.5 ft); SCSI-2 adapter, 50-conductor A cable; a 50-pin, low-density, male, straight connector on one end and a SCSI-3, 68-pin, special male, right-angle connector on the other; only 50 pins are used; typically used to connect an HSZ10-AA controller to a KZESA bus adapter
BN21M-03	3.0 m (9.8 ft)
BN21M-05	5.0 m (16.4 ft)
BN21M-10	10.0 m (32.8 ft)
BN21M-15	15.0 m (49.2 ft)
BN21M-20	20.0 m (65.6 ft)
BN21N-Series SCSI Transition Cables	
BN21N-02	2.0 m (6.5 ft); SCSI-2 adapter, 50-conductor A cable; a 50-pin, high-density, male, straight connector on one end and a SCSI-3, 68-pin, special male, right-angle connector on the other; only 50 pins are used; typically used to connect a DWZZA-AA adapter to a PMAZB host adapter
BN21N-03	3.0 m (9.8 ft)
BN21N-05	5.0 m (16.4 ft)
BN21N-10	10.0 m (32.8 ft)
BN21N-15	15.0 m (49.2 ft)
BN21N-20	20.0 m (65.6 ft)
BN21P SCSI Transition Cables	
BN21P-02	2.0 m (6.5 ft); SCSI-2 adapter, 38-conductor A cable; a SCSI-2, 50-pin, male, straight, microD connector on one end and a SCSI-3, 68-pin, special female, straight connector on the other; only 38 pins are used; typically used to connect a VAXstation 3100 to a BA350 shelf

(continued on next page)

Parts Information

11.1 Hardware

Table 11-5 (Cont.) Cables

Part Number	Description
BN21P SCSI Y Cables	
BN21P-0B	.152 m (0.5 ft). SCSI-3, 68-pin connector to 50-pin connector on a KZESA adapter. Must be used with the H879-AA 68-pin terminator.
BC09D-Series SCSI Transition Cables	
BC09D-03	1.09 m (3.0 ft). Low-density 50-conductor to high-density 50-conductor
BC09D-06	1.8 m (6.0 ft). Low-density 50-conductor to high-density 50-conductor
BC09D-12	3.67 m (12.0 ft). Low-density 50-conductor to high-density 50-conductor

11.2 Firmware

The HSZ10-AA controller firmware is the portion of the code that resides with the controller's EPROM. The HSZ10-AA controller is described in Section 1.2.1.3 and Section 1.2.1.5. The physical layout of the HSZ10-AA controller is shown in Figure 1-4.

11.3 Software

The StorageWorks RAID Array 110 utilities are a collection of utilities designed to provide RAID configuration and management functionality in the OpenVMS VAX, MS-DOS, NetWare, and SCO UNIX operating system environments. This section describes the necessary hardware and software components required for each operating system. It also describes the various StorageWorks RAID Array 110 utility kits and their components.

Important Note

The DEC RAID OpenVMS VAX Utility is renamed to the StorageWorks RAID Array 110 Utility for OpenVMS VAX. This guide refers to DEC RAID OpenVMS VAX Utility in code examples only. In all other cases, the new product name is used.

11.3.1 Hardware Requirements

The StorageWorks RAID Array 110 utilities are designed to run with the following hardware components:

- StorageWorks RAID Array 110 Subsystem with the HSZ10-AA controller
- KZESA host adapter to SCSI differential host adapter
- DEC DWZZA-VA single-ended to differential SCSI adapter (for OpenVMS VAX only)
- Supported drives: RZ25 and RZ26

Refer to the *StorageWorks RAID Array 110 Subsystem User's Guide* for further information on the StorageWorks RAID Array 110 Subsystem.

11.3.2 Software Requirements

The StorageWorks RAID Array 110 utilities offer RAID management services in the OpenVMS VAX, MS-DOS, NetWare, and SCO UNIX environments. The following sections discuss StorageWorks RAID Array 110 utility support for each operating system. The software kits are as follows:

- QA-0FFAA-H5—StorageWorks RAID Array 110 Utility for OpenVMS VAX (TK50)
- QA-0FFAA-H8—StorageWorks RAID Array 110 Utility for OpenVMS VAX (CDROM)
- QA-0FCAA-BC—StorageWorks RAID Array 110 Utility for MS-DOS
- QA-0FDAA-BC—StorageWorks RAID Array 110 Utility for Novell NetWare
- QA-0FEAA-BC—StorageWorks RAID Array 110 Utility for SCO UNIX

Note

The QA part numbers contain a zero, not the letter O.

11.3.2.1 OpenVMS VAX V5.5-2

The StorageWorks RAID Array 110 Utility for OpenVMS VAX provides RAID configuration and management services for the OpenVMS VAX V5.5-2 environment. The software kit is distributed on a TK50 cartridge or CDROM and contains utilities to support the HSZ10-AA controller.

11.3.2.2 MS-DOS

The StorageWorks RAID Array 110 Utility for MS-DOS is an MS-DOS-based utility used primarily to configure the RAID array in an MS-DOS and/or SCO UNIX environment. The utility is provided on an MS-DOS Version 5.0 bootable floppy.

11.3.2.3 NetWare

In addition to the StorageWorks RAID Array 110 Utility for MS-DOS, the StorageWorks RAID Array 110 Utility for Novell NetWare (or NetWare RAID Manager™) provides the same level of functionality in the NetWare Version 3.11 operating system environment. Additional features include:

- Array Parity Check/Repair
- Array Monitor Daemon to notify users of drive status changes

11.3.2.4 SCO UNIX

Configuration in SCO UNIX Version 3.2.4 is accomplished through the StorageWorks RAID Array 110 Utility for SCO UNIX. Additional functionality is provided through a set of SCO UNIX-based utilities including:

- Array parity check/repair utility
- RAID Monitor to notify users of changes in LUN and drive status
- RAID status utility to display current LUN and drive status

Parts Information

11.3 Software

11.3.3 StorageWorks RAID Array 110 Utility Kits

Software kits are available for the StorageWorks RAID Array 110 utilities for each supported operating system. The following sections describe the kits and the files contained within each kit.

11.3.3.1 StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit

The StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit contains the files listed in Table 11–6, which can be extracted by using the VMSINSTAL procedure.

Table 11–6 StorageWorks RAID Array 110 Utility for OpenVMS VAX Kit

Saveset A Files	Description
KITINSTAL.COM	Installation script
HSZ10010.RELEASE_NOTES	Release notes
Saveset B Files	Description
92000304.AFW	HSZ10 Downloadable software
CCF_A.COM	Sample configuration file
CCF_B.COM	Sample configuration file
CCMFLIST	Configuration file list
EE_HSZ10.DEF	EEPROM default settings
EE_HSZ10.VMS	EEPROM OpenVMS VAX settings
HSZ10\$ACE.EXE	Array Configuration Editor (ACE) utility
HSZ10\$ACF.EXE	Array Initialization Facility†
HSZ10\$APC.EXE	Array Parity Check utility
HSZ10\$APR.EXE	Array Parity Repair utility
HSZ10\$ASM.EXE	Array Status Monitor utility
HSZ10\$HSZ_EEPROM.EXE	Utility to set EEPROM values
HSZ10\$STARTUP.COM	Startup file for HSZ10-AA controller
HSZ10_UTIL.CLD	Command definitions
HSZ10_UTIL.HLP	Help file
UTILEXE.DAT	Executable directory locations
UTILSRC.DAT	Source file directory locations
Help files	Help files

†This file is used only at the time of initialization.

11.3.3.2 StorageWorks RAID Array 110 Utility for MS–DOS Kit

The StorageWorks RAID Array 110 Utility for MS–DOS Kit contains a bootable MS–DOS Version 5.0 diskette with the files listed in Table 11–7.

Table 11-7 StorageWorks RAID Array 110 Utility for MS-DOS Kit

Files	Description
COMMAND.COM	MS-DOS Version 5.0
DPTDDL.SYS	KZESA MS-DOS driver
CONFIG.SYS	Configuration file
RAIDMGR.CWA	RAID Manager overlay C-Worthy file
RAIDMGR.EXE	Stand-Alone RAID Manager
RAIDMGR.MSG	RAID Manager message file
RAIDMGR.HLP	RAID Manager help file
RAIDMGR.CFG	RAID Manager configuration file
ACE.EXE	Array Configuration Editor (ACE)
ACF.EXE	Array Configuration Utility (ACF)
FAD.EXE	Array Diagnostic Utility (FAD)
README.BAT	RAID Manager online introduction
READ.ME	RAID Manager online introduction text
PAGE.EXE	Pagination for online introduction
FDISK.EXE	Fdisk utility
ACEHELP <DIR>	Directory of help files
DOS <DIR>	Directory of DOS utilities

11.3.3.3 StorageWorks RAID Array 110 Utility for Novell NetWare Kit

The StorageWorks RAID Array 110 Utility for Novell NetWare Kit contains the StorageWorks RAID Array 110 Utility for MS-DOS Kit as well as a diskette with the files listed in Table 11-8.

Table 11-8 StorageWorks RAID Array 110 Utility for Novell NetWare Kit

Files	Description
RMINST.NLM	Installation file
RAIDMGR.NLM	NetWare RAID Manager
RAIDMGR.CFG	Configuration file
RAIDMGR.DAT	Menu and form definitions for RAID Manager
RAIDMGR.MSG	Message file
RAIDMGR.HLP	Help file
RAIDIF1.NLM	NetWare file
ACE.NLM	Array Configuration Editor (ACE)
APC.NLM	Array Parity Check/Repair Utility (APC)
AMD.NLM	Array Monitor Daemon Utility (AMD)
RELEASE.ID	Release version number
ACEHELP <DIR>	Directory of help files

Parts Information

11.3 Software

11.3.3.4 StorageWorks RAID Array 110 Utility for SCO UNIX Kit

The StorageWorks RAID Array 110 Utility for SCO UNIX Kit contains the StorageWorks RAID Array 110 Utility for MS-DOS Kit as well as a diskette with the files listed in Table 11-9.

Table 11-9 StorageWorks RAID Array 110 Utility for SCO UNIX Kit

Files	Description
makedpth	Create device files used by utility
help.txt	Help text file
monitor	RAID Monitor executable
monitorsh	Shell script to run RAID Monitor
parity	Parity Check/Repair Utility executable
paritysh	Shell script to run Parity Check/Repair Utility
raidstat	RAID Status Utility executable
raidstatsh	Shell script to run the RAID Status Utility
read.me	Read me file with general information about the kit

Preventive and Proactive Maintenance

This chapter contains information on the following:

- Upgrading software
- Checking/repairing array parity for the StorageWorks RAID Array 110 Utility for Novell NetWare
- Array Monitor Daemon for the StorageWorks RAID Array 110 Utility for Novell NetWare
- Array parity check/repair for the StorageWorks RAID Array 110 Utility for SCO UNIX
- RAID Monitor Utility for the StorageWorks RAID Array 110 Utility for SCO UNIX
- RAID Status Utility for the StorageWorks RAID Array 110 Utility for SCO UNIX

Note

The Array Parity Check (APC), the Array Parity Repair (APR), and the Array Status Monitor (ASM) utilities for the OpenVMS VAX operating system are described in Chapter 4.

12.1 Upgrading Software

The disk array software can be upgraded using a StorageWorks RAID Array 110 utility, which provides downloading capability. Upgrading the software does not affect the user data.

Refer to the *DEC RAID Utilities User's Guide* for instruction regarding upgrading the software.

12.2 Checking/Repairing Array Parity for Novell NetWare

This section provides information on the following:

- What is Parity Check/Repair?
- How automatic Parity Check/Repair works
- How to run manual Parity Check/Repair
- Changing Parity Check scheduling

12.2.1 What is Parity Check/Repair?

Note

Parity Check/Repair applies only to RAID 1 and 5. RAID 0 does not have array parity, and therefore cannot be checked and repaired. RAID 1 does not really have "parity" either, but parity check compares data on the mirrored drives. In addition, you cannot run Parity Check/Repair on a degraded RAID 1 or 5 LUN (a LUN with a status of *Degraded*).

Parity Check/Repair performs the following functions:

- Scans the LUN and checks the array parity for each block in the LUN. On a RAID 1 LUN, parity check compares the data on each mirrored pair, block by block.
- Repairs any array parity errors found during the parity check. On a RAID 1 LUN, the array controller changes the data on the mirror disk to make it match the data on the data disk. On a RAID 5 LUN, the controller changes the parity segment so that it is consistent with the data segments.

Note that if the array parity errors resulted from corrupted data, the data is not repaired, only the array parity. Also, note that you may still lose some data as a result of the power failure or abnormal shutdown, especially if you do not have an uninterruptible power supply (UPS).

12.2.1.1 When to Run Parity Check/Repair

Run Parity Check/Repair at the following times:

- Automatically at a time determined during the RAID Manager installation (see Section 12.2.2). Automatic Parity Check/Repair helps guarantee the data integrity of the LUN so that you can reconstruct the data on the array if a drive fails.
- Manually after an abnormal server or array shutdown (see Section 12.2.3). As the result of such a shutdown, required array parity may not have been updated, resulting in potential data corruption. You must run Parity Check/Repair on the LUN after you run the NetWare VRepair command to repair any volumes on the unit.

12.2.2 How Automatic Parity Check/Repair Works

This section describes how the automatic Parity Check/Repair works. See Section 12.2.3 for information on manual operation.

During the RAID Manager installation, you specify when you want the Parity Check/Repair utilities to run automatically. Automatic Parity Check/Repair ensures the data integrity of your array, which allows you to accurately reconstruct data on the array after a device fails.

12.2.2.1 Automatic Parity Check/Repair Operation

- During the RAID Manager installation, you specify a time to run automatic Parity Check/Repair. You can also specify a parity check without repairing any errors found.
- At the specified time of day, the system starts the parity check operation. A new screen is created on the system console for the parity check operation.

Preventive and Proactive Maintenance

12.2 Checking/Repairing Array Parity for Novell NetWare

- Parity check writes messages to this screen. If any errors are found, they are also written to a parity error log file.
- The log file is named device-ID.chk, where device-ID is the device ID of the LUN checked (for example, 8B001.chk is the name of the parity error log file for device 8B001). The file is written in the SYS/raidmgr directory. A separate log file is created for each LUN.
- If parity errors are found, and you specified automatic parity repair, the system repairs the parity errors.

If you did not specify automatic parity repair, you have to check the parity error log files to see if any errors are recorded, and then run Parity Check/Repair manually. There is no reason not to specify automatic parity repair.
- Automatic Parity Check/Repair is performed on each LUN in turn.
- After the automatic Parity Check/Repair process, the parity error log file contains one of the following:

If no errors were found, the file contains only the date and time automatic Parity Check/Repair was run.

If errors were found, the file contains a list of the bad blocks found and repaired (if parity repair was run).
- If desired, you can change the scheduled automatic Parity Check/Repair time (see Section 12.2.4).
- You can check the current settings for automatic Parity Check/Repair by selecting *RAID Manager Configuration* from the *Disk Array Services* menu, then selecting *General Configuration* to view RAID Manager configuration parameters.

Preventive and Proactive Maintenance

12.2 Checking/Repairing Array Parity for Novell NetWare

12.2.3 How to Run Manual Parity Check/Repair

This section describes how to run Parity Check/Repair utilities manually. See Section 12.2.2 for information on automatic operation.

You run Parity Check/Repair manually after an abnormal shutdown (power failure, system crash, and so on). Note that you may still lose some data as a result of the power failure or abnormal shutdown, especially if you do not have an uninterruptible power supply (UPS).

You must run Parity Check/Repair after you run VRepair on the affected volumes.

12.2.3.1 Manual Parity Check/Repair Operation

- Manual Parity Check/Repair is run from the RAID Manager. You should run manual parity check after you have run VRepair on any affected volumes.
- You can run parity check without repairing parity (although normally there is no reason not to repair parity).
- An entry is made in the activity log showing the parity check operation.
- Parity check creates a file to store parity errors found and repaired. This file is named device-ID.chk, where device-ID is the device ID of the LUN checked (for example, 8B001.chk is the name of the parity error log file for device 8B001). The file is written in the SYS/raidmgr directory.
- You can temporarily change the name of the parity error log file using the *Change Parity Report File Name* option in the *Disk Array Parity* menu. This change applies only to this LUN, and only for the next run of parity check. If you leave the RAID Manager without running parity check, the change no longer applies.
- After you run manual Parity Check/Repair, the parity error log report is automatically displayed, without your having to enter the file name of the file.
- You must run Parity Check/Repair on each LUN separately.

12.2.3.2 Running the Check

Use the following procedure to check and repair parity manually:

1. If you are running Parity Check/Repair after an abnormal system shutdown, first run VRepair on all affected volumes before checking parity.
2. Start the RAID Manager by entering the following command from the console:

```
load raidmgr
```
3. Use the *Device Selection* option in the *Disk Array Services* menu to select the LUN you want to check.
4. After you have selected the LUN you want to check, select *Array Parity* from the *Disk Array Services* menu.
5. From the *Disk Array Parity* menu, you can select the following options:
 - Check and Repair Parity—use this option to check and repair parity on the selected LUN. The parity error log file contains a list of the errors found and repaired.
 - Check Parity (no repair)—use this option to check parity on the selected LUN without repairing any errors found. The parity error log file contains a list of the errors found.

Preventive and Proactive Maintenance

12.2 Checking/Repairing Array Parity for Novell NetWare

- **Change Parity Report File Name**—use this option to change the name of the parity error log file for the current LUN before you run Parity Check/Repair. This change applies only to this LUN, and only for the next run of parity check.
 - **View Parity Report**—use this option to view the parity report file for any LUN.
 - **Change Automatic Parity Time**—use this option to change the scheduled time for automatic parity check. You cannot change the automatic parity repair parameter here.
6. If you want to write the parity error report to a file other than the default file, select the *Change Parity Report File Name* option and enter the new name. Any new name applies only to this LUN and only for this run of Parity Check/Repair.
 7. Select either *Check and Repair Parity* or *Check Parity (no repair)* to start parity check.
 8. A prompt is displayed to confirm that you want to check parity. Press the Y key and the Enter key to start the selected operation.
 9. A popbox is displayed as parity check is performed showing the percentage of the check completed. When the operation is completed, press the Enter key.
 10. The parity error report for the LUN you just checked is displayed.
 - If you selected Parity Check/Repair, this is a list of the parity errors found and repaired. If no errors were found the report indicates this.
 - If you selected parity check without repair, this is a list of the parity errors found. To repair these errors, you have to run parity check with repair. If no errors were found the report indicates this.
 11. After viewing the report, press the Esc key to return to the *Disk Array Parity* menu.
 12. If you want to check more LUNs, press the Esc key to return to the *Disk Array Services* menu, then use the *Device Selection* option to select another LUN. Then go back to Step 4 and repeat this procedure.
 13. Press the Esc key until you return to the RAID Manager exit prompt, then exit the RAID Manager.

12.2.4 Changing Parity Check Scheduling

Use the following procedure to change the scheduled time for automatic Parity Check/Repair.

Note

You can also change the scheduled time (but not the parity repair option) by selecting *Change Automatic Parity Time* from the *Disk Array Parity* menu.

1. Start the RAID Manager by entering the following command from the console:

```
load raidmgr
```

2. From the *Disk Array Services* menu, select *RAID Manager Configuration*.
3. From the *RAID Manager Configuration* menu, select *General Configuration*.
4. Use the down-arrow key to highlight the *Parity Check Start Time* parameter on the General Configuration display.
5. Enter the new time for automatic parity check. Enter the time as a 24-hour value (that is, 1:00 PM is entered as 13:00).

Important

While the Parity Check/Repair utility is running, workstations connected to the server may slow considerably until the check function is completed. During the utility's daily run, try and choose a time when this slowdown does not affect your operations.

6. After you enter the new time, press the Enter key. The cursor automatically moves down one line to highlight the *Automatic Parity Repair* option.
7. If you want to change the current value, press any key except the space bar or the Enter key to toggle to the other value (*yes* or *no*). There is normally no reason not to specify automatic parity repair.
8. After you have specified the parity repair parameter, press the Enter key to save the new value. The cursor automatically highlights the next parameter.
9. Press the Enter key to exit the *General Configuration* menu.
10. You are then prompted to determine whether you want to leave the menu with the current changes, leave without the changes, or return to the *General Configuration* menu to make more changes.
11. After you select *Yes*, you return to the *RAID Manager Configuration* menu.
12. Press the Esc key until you get the RAID Manager exit prompt, then exit the RAID Manager.

12.3 The Array Monitor Daemon for Novell NetWare

This section provides information on the following:

- Array Monitor Daemon (AMD)
- Error message types
- AMD error message format
- AMD error messages
- RAID Manager messages

12.3.1 Array Monitor Daemon

The AMD scans for nonoptimal status on your array devices and notifies you when one is found. While these status changes are also written to the system error log and on the console screen, the AMD notifies you immediately of the changes in status, and displays a message on the console terminal, regardless of which screen you happen to be in.

12.3.1.1 How the AMD Works

The AMD operates as follows:

1. At a specified interval (determined by the delay parameter value), the daemon polls the array devices (both LUNs and individual drives) for their status.
2. If a status other than *Ok - Optimal* is returned (for any array component), the daemon *takes over* the console screen and displays the other status on the console screen for a specified period (determined by the persistence parameter value).
3. If at the end of the persistence interval you have not acknowledged the message, the daemon releases the screen. It displays the status again after the specified delay interval.
4. To acknowledge the message, press any key except the Enter key. The message then remains on the screen until you press the Enter key. After you press the Enter key, the delay interval begins again.

The AMD continues to display the same status until you correct the condition that caused the status.

Important Note

The AMD does not poll array devices if RAID Manager, the NetWare Install utility, or the NetWare VRepair utility is currently loaded.

12.3.1.2 What to do When an Error is Displayed

When the AMD displays an error status on the console, you should take the following steps:

1. When the message first appears, press any key except the Enter key to hold the message on the screen (pressing the Enter key clears the screen—if that happens, you have to wait for the message to be displayed again).
2. Read the message, taking note of the array device named in the message (see the AMD message format Section 12.3.3).

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

3. Change to the console screen (hold down the Ctrl Key and press the Esc key) to see if any messages are displayed there (they may have scrolled off the screen).
4. Start the RAID Manager and check the LUN status and drive status of the devices named in the error message.
5. Determine the nature of the error. See Section 4.3.8 for a list of LUN status and drive status.
6. Read the system error log to see if there are any associated messages in the log that might further explain the error and its consequences. Frequently, there are also standard NetWare error messages associated with the error.
7. Take whatever action is required to restore your array (refer to Section 8.7 of this manual).

12.3.1.3 Error Log Example

This section gives an example of AMD operations on a RAID 5 LUN. The example represents one set of error messages you might see if a single drive fails on a RAID 5 LUN. It is not necessarily the only sequence (the circumstances which caused the drive to fail may result in different error messages being displayed).

1. You are using the LUN when the following message appears on the console terminal screen (which is not currently set to the system console) at 1:40 PM:

```
Disk Array Monitor  
Version xx.xx.xx
```

```
Device #7 (8B001) Degraded: Drive Failure  
Array Drive Error: Channel 1, ID 0: Drive Failure
```

Press the Enter key to clear screen, any key to hold screen.

2. You press any key except the Enter key to keep the message on the screen so that you can read it.
3. From this message, you can see that a drive has failed in a RAID 1 or RAID 5 LUN, leaving the LUN in the degraded mode. The LUN involved is Device #7.
4. After you have read the message, press the Enter key to clear the error message from the screen.
5. Hold down the Ctrl Key and press the Esc key to switch to the console screen and see the following messages:

```
7/29/92 1:39 PM: 0.0.0 Array Monitor Daemon: Array Device Error:  
Device #7 (8B001) Status: Degraded: Drive Failure
```

```
7/29/92 1:39 PM: 0.0.0 Array Monitor Daemon: Array Component Disk Error:  
Channel 1 ID 0, Status: Failed: Drive Failure
```

6. Next, start the RAID Manager and use the *Device Selection* option to select Device #7. You can see from the display that Device #7 is a RAID 5 LUN and that one of the drives in the array has failed.
7. Check your system error log until you find the corresponding messages (from the *Disk Array Services* menu, select *System Messages*, then *System Error Log*). In this case, you see the following messages:

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

7/29/92 1:39 PM Severity = 3
0.0.0 Array Monitor Daemon: Array Device Error:
Device #7 (8B001) Status: Degraded: Drive Failure

7/29/92 1:39 PM Severity = 3
0.0.0 Array Monitor Daemon: Array Component Disk Error:
Channel 1 ID 0, Status: Failed: Drive Failure

8. Check any messages preceding and following to see if they clarify the reasons for the drive failure. In this case, there are no further messages.

Note that you could see standard NetWare error messages relating to the array error messages. For example, for an error more serious than the one in this example you might see:

9/27/91 10:57:38 am Severity = 4.
1.1.10 Device #1 (8B001) *** NCR ADP-92/01 0100 deactivated due to drive failure.
9/27/91 10:57:38 am Severity = 4.
1.1.86 Volume ARRAY dismounted due to drive deactivation

These messages from NetWare indicate that the array device has been deactivated as the result of a drive failure, and that the volume ARRAY (contained on the deactivated device), has been dismounted. These messages would follow array-specific messages indicating drive or LUN failure.

9. Search through the rest of the error log to make sure there are no errors affecting other LUNs or other errors affecting this unit.
10. Given the error indicated, you would then go to Section 8.7 for the steps required to restore a RAID 5 array. (In this case, you would replace the failed drive and reconstruct the data on it.)
11. Clear the error log (copy it to a file first so that you can refer to it later if necessary).

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

12.3.2 Error Message Types

The RAID Manager and the disk array software generate three basic types of error messages:

- Array-specific error messages, which are displayed only for the disk array. These messages are displayed by AMD on the current screen, the console screen, and in the system error log (see Section 12.3.4).
- Standard NetWare error messages. These messages are displayed on the console and written to the system error log. These messages are also used for devices other than the disk array. See your NetWare System Messages documentation for an explanation of these messages. These messages are not displayed by the AMD.
- Error messages that occur while you are running the RAID Manager. These messages are displayed on the console as soon as they occur. They are not written to the system error log (see Section 12.3.4.4).

12.3.3 AMD Error Message Format

There are three kinds of AMD messages:

- AMD messages that appear in the pop-up screen.
- AMD messages displayed on the system console.
- AMD messages written in the system error log.

12.3.3.1 AMD Pop-Up Message Format

AMD pop-up messages are displayed periodically on the console terminal, regardless of which screen is currently displayed, until you acknowledge the message. They have the following format:

```
Disk Array Monitor  
Version xx.xx.xx
```

```
Device #n (devID) lun-status  
Array Drive Error: Channel x, ID y: drive-status
```

Where:

- **n** is the device number of the LUN affected by the status change.
- **devID** is the device ID of the LUN affected by the status change. Device IDs have the following format:

The first two digits indicate the type of host adapter. These digits are assigned by NetWare. For example, 8B is the KZESA differential host adapter.

The third digit is the board number assigned to the host adapter by the NetWare system.

The fourth digit is the SCSI ID of the array controller.

The fifth, or last, digit is the LUN number.

- **lun_status** is the current status of the affected LUN.
- **x** is the channel number of the drive affected by the status change.
- **y** is the SCSI ID of the drive affected by the status change.
- **drive_status** is the current status of the affected drive.

12.3.3.2 AMD Console Message Format

AMD console messages are displayed on the system console once when the status change is detected. AMD console messages have the following format (the first message given is for a LUN; the second message is for an array drive):

```
mm/dd/yy hh:mm : 0.0.0 Array Monitor Daemon: Array Device Error:  
Device #n (devID) Status: lun-status
```

```
mm/dd/yy hh:mm : 0.0.0 Array Monitor Daemon: Array Component Disk Error:  
Channel x ID y, Status: drive-status
```

Where:

- mm/dd/yy hh:mm is the time the status change was detected.
- n is the device number of the LUN affected by the status change.
- devID is the device ID of the LUN affected by the status change. Device IDs have the following format:
The first two digits indicate the type of host adapter. These digits are assigned by NetWare. For example, 8B is the KZESA differential host adapter.
The third digit is board number assigned to the host adapter by the NetWare system.
The fourth digit is the SCSI ID of the array controller.
The fifth, or last, digit is the logical unit number.
- lun_status is the current status of the affected LUN.
- x is the channel number of the drive affected by the status change.
- y is the SCSI ID of the drive affected by the status change.
- drive_status is the current status of the affected drive.

12.3.3.3 AMD Error Log Message Format

AMD error log messages are written into the system error log when the status change is first detected. They have the following format (note that this is a standard NetWare error log format):

```
dd/mm/yy hh:mm Severity = z  
0.0.0 Array Monitor Daemon: Array Device Error:  
Device #n (devID) Status: lun-status
```

```
dd/mm/yy hh:mm Severity = z  
0.0.0 Array Monitor Daemon: Array Component Disk Error:  
Channel x ID y, Status: drive-status
```

Where:

- mm/dd/yy hh:mm is the time the status change was detected.
- z is the severity of the error:
 - 0 = Informational
 - 1 = Warning
 - 3 = Critical
 - 4 = Fatal
- n is the device number of the LUN affected by the status change.

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

- **devID** is the device ID of the LUN affected by the status change. Device IDs have the following format:
The first two digits indicate the type of host adapter. These digits are assigned by NetWare. For example, 8B is the KZESA differential host adapter.
The third digit is the board number assigned to the host adapter by the NetWare system.
The fourth digit is the SCSI ID of the array controller.
The fifth, or last, digit is the logical unit number.
- **lun_status** is the current status of the affected LUN.
- **x** is the channel number of the drive affected by the status change.
- **y** is the SCSI ID of the drive affected by the status change.
- **drive_status** is the current status of the affected drive.

12.3.4 AMD Error Messages

The following messages are AMD messages, and appear on your system console if AMD detects a status change in the array.

12.3.4.1 Array Component Errors

These errors have the following format:

```
Array Monitor Daemon: Array Component Error
<drive name > Status: Array Formatting
```

where <drive name > is the device ID of the drive.

Note

The errors in the following tables are all Array Component Errors, so the words *Array Monitor Daemon: Array Component Error* are not displayed in the following tables.

Table 12-1 Array Component Errors

Message	Probable Cause	Action To Take
<drive name> Status: Array Formatting	The array controller is formatting a drive.	No action.
<drive name> Status: Bad Channel	You have installed a drive that has a different channel than the drive that was originally in its place.	Replace the drive with a drive having the correct channel.
<drive name> Status: Drive Not Ready	The drive failed the Test Unit Ready (that occurs during a LUN format) after the drive was hot swapped, added, or after the subsystem was started.	Replace the drive, and make sure that the data is restored to the LUN containing it.

(continued on next page)

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

Table 12-1 (Cont.) Array Component Errors

Message	Probable Cause	Action To Take
<drive name> Status: Failed	The drive failed during a write to the LUN, or failed the initialization process.	Replace the drive, and make sure that the data is restored to the LUN containing it.
<drive name> Status: Failed by User	The drive was failed through ACE.	Replace the drive, and make sure that the data is restored to the LUN containing it.
<drive name> Status: Format Failed	The SCSI format issued to the drive failed. This format was issued through ACE as a LUN format or drive format before reconstruction.	Try to format the LUN again. If it fails again, replace the drive and make sure that the data is restored to the LUN containing it.
<drive name> Status: Illogical SCSI ID	You have installed a drive that has a different ID than the drive that was originally in its place.	Replace the drive with a drive of the correct ID.
<drive name> Status: Newly Replaced	The drive has just been replaced.	No action.
<drive name> Status: OK - Warning Condition	The drive has experienced errors while read operations were occurring.	Back up the drive (if possible) and restore the LUN.
<drive name> Status: OK - Warning Condition	The drive is part of a RAID 0 LUN and has experienced errors while reading or writing.	Back up the drive (if possible) and restore the LUN.
<drive name> Status: Reconstructing	The drive data is being reconstructed by the array controller.	No action.
<drive name> Status: Start of Day Failure	At system boot-up, the drive failed because it did not accept writes from the array controller.	Back up the drive (if possible) and restore the LUN.
<drive name> Status: Write Failed	The drive failed when the array controller was initializing it for array parity.	Back up the drive (if possible) and restore the LUN.
<drive name> Status: Wrong Capacity	The drive is incompatible with the other drives in the LUN because its capacity is different.	Replace the drive with a certified drive.
<drive name> Status: Wrong Ctrl Serial No.	The drive that was just installed was originally connected to a different array controller and was formatted with special information about that array controller, AND it replaced a drive that was optimal, spare, or warning.	Fail the drive using ACE, and restore the LUN.
<drive name> Status: Wrong Parameters	The drive is incompatible with the other drives in the LUN because it is a different size.	Replace the drive with a drive of the same size as the other drives.

(continued on next page)

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

Table 12–1 (Cont.) Array Component Errors

Message	Probable Cause	Action To Take
<drive name> Status: Wrong Replacement Disk	An optimal or warning drive in a degraded LUN was replaced, or an optimal or spare drive or a drive that is being formatted was replaced and then the subsystem was shutdown.	Replace the drive that was mistakenly removed from the subsystem, turn on the subsystem, and remove the drive that you had intended to remove. CAUTION: DO NOT send I/O to the LUN until you have completed all the steps above.
<drive name> Status: Wrong Sector Size	The drive is incompatible with the other drives in the LUN because its sector size is different.	Replace the drive with a certified drive.

12.3.4.2 Array Device Errors

Table 12–2 lists the array device error messages, their probable cause, and action(s) to take.

Array Device errors have the following format:

Array Monitor Daemon: Array Device Error
<LUN name> Status: Bad Channel

Where <LUN name> is the device name of the LUN.

Note

The errors in the following tables are all Array Device Errors, so the words *Array Monitor Daemon: Array Device Error* are not displayed in the following tables.

Table 12–2 Array Device Errors

Message	Probable Cause	Action To Take
<LUN name> Status: Bad Channel	The LUN is dead because a drive that has a different channel than the drive that was originally in its place was installed.	Replace the drive with the wrong channel with a drive of the correct channel.
<LUN name> Status: Bad Identifier	The LUN is dead because a drive has a different ID than the drive that was originally in its place was installed.	Replace the drive with the wrong channel with a drive of the correct ID.
<LUN name> Status: Bad Parameters	The LUN is dead because a drive that has a different capacity from the other drives in the LUN has been installed.	Replace the drive with the wrong capacity with a certified drive.

(continued on next page)

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

Table 12-2 (Cont.) Array Device Errors

Message	Probable Cause	Action To Take
<LUN name> Status: Component Failure	Two drives were removed from a RAID 5 LUN, and then the multirank subsystem was shutdown.	Shut down the subsystem, replace the drives in their original positions, and restart the subsystem. CAUTION: DO NOT send I/O to the LUN until you have completed the steps above.
<LUN name> Status: Degraded - Drive Failure	A drive has failed in the LUN.	Replace the drive as soon as possible. The LUN remains degraded until the drive is replaced.
<LUN name> Status: Degraded - Drive Formatting	The LUN is degraded and the drive that was just replaced is being formatted by the array controller.	No action.
<LUN name> Status: Drive Failures	Multiple drives are failed in the LUN.	Replace the drive. Then format the LUN and copy the data from the backup media to the LUN.
<LUN name> Status: Formatting	The LUN is being formatted by the array controller.	No action.
<LUN name> Status: OK - Reconstructing	The array controller is reconstructing the data on the newly-replaced drive.	No action.
<LUN name> Status: Waiting for Format	The LUN is dead and must be formatted.	Format the LUN using ACE.
<LUN name> Status: Wrong Drive.	The LUN is dead because a drive that was either optimal or in warning was replaced.	First Option: If no writes are sent to the LUN, do the following: <ol style="list-style-type: none">a. Remove the new drive.b. Replace the original drive.c. Either, install the new drive in the correct location if the wrong drive was replaced, or fail the warning drive using ACE and replace it. Second Option: If writes are sent to the LUN, reformat the LUN and copy the data from the backup media to the LUN.

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

12.3.4.3 Other Status Change Errors

Table 12–3 lists the error messages, their probable cause, and action(s) to take.

Table 12–3 Other Status Change Errors

Message	Probable Cause	Action To Take
Automatic Disk Array Parity Check Complete	Displayed once daily when the automatic parity check utility completes.	No action.
Automatic Disk Array Parity Check Started	Displayed once daily when the automatic parity check utility starts.	No action.
Array Monitor Daemon: <LUN 0 name> Obsolete Downloadable Controller Firmware [Version xx.yy.zz]	An old version of firmware is on the array drives.	Download new firmware.
Array Monitor Daemon: Power Supply #x Failure	A power supply failed.	Replace the failed power supply.
Array Monitor Daemon: Subsystem Unique Device #x failure	A subsystem component failed (for example, a fan).	Replace the failed component.
Disk Array Parity Check: Device #n (DevID)	The automatic parity check started for this device.	No action.

12.3.4.4 RAID Manager Messages

Table 12–4 lists the error message explanations given in the RAID Manager messages. These explanations may appear whenever you are using a RAID Manager function or option.

Table 12–4 RAID Manager Messages

Explanation	Probable Cause	Action To Take
ABORT: Insufficient Memory	Software problem.	Call your Customer Support Representative.
ABORT: Invalid option <option> in the command line.	The incorrect option was typed.	Type the command again using the correct option.
ABORT: No message text for confirmation header <#>, <message>	Software problem.	Call your Customer Support Representative.
ABORT: Unable to create confirmation menu	Software problem.	Call your Customer Support Representative.
ABORT: Unable to locate the RAID Manager Data File	The file is deleted, or it was never installed properly.	Remove the RAID Manager and reinstall it.

(continued on next page)

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

Table 12-4 (Cont.) RAID Manager Messages

Explanation	Probable Cause	Action To Take
ABORT: Unable to locate the RAID Manager Help File	The file is deleted, or it was never installed properly.	Remove the RAID Manager and reinstall it.
ABORT: Unable to locate the RAID Manager Message File	The file is deleted, or it was never installed properly.	Remove the RAID Manager and reinstall it.
Error encountered during disk array parity checking.	The array controller was unable to complete the array parity operation due to an array controller software error.	Call your Customer Support Representative.
Error encountered during disk array parity repair. The array device may be require other servicing.	The array controller was unable to complete the array parity operation due to an array controller software error.	Call your Customer Support Representative.
File <file> on path <path> cannot be created or opened for writing.	There is a problem with the file, or there is a RAID Manager software problem.	Try the operation again. If it fails again, call your Customer Support Representative.
File <file> on path <path> cannot be opened for reading.	There is a problem with the file, or there is a RAID Manager software problem.	Try the operation again. If it fails again, call your Customer Support Representative.
Firmware data file <file> is not present in path <path>.	The RAID Manager cannot find the file in the path you selected.	Try the operation again using the correct path name. If the problem occurs again, call your Customer Support Representative.
Incomplete download file or invalid download record format.	There is a problem with the diskette or with the file.	Try to download the firmware again. If the problem occurs again, call your Customer Support Representative.
The RAID Manager Help file <file> is corrupt.	Software problem.	Remove the RAID Manager utilities and reinstall them. If the problem occurs again, call your Customer Support Representative.
Unable to get exclusive access to the array device. A volume on the device may be mounted, or another process may have the device locked.	A volume may be mounted, or there is another process that is presently accessing the LUN (for example, the Install utility, or another NetWare NLM).	<p>First Option: Unmount the volume if it is mounted, then try the operation again.</p> <p>Second Option: Try to terminate the other process accessing the LUN, then try the operation again.</p>

(continued on next page)

Preventive and Proactive Maintenance

12.3 The Array Monitor Daemon for Novell NetWare

Table 12-4 (Cont.) RAID Manager Messages

Explanation	Probable Cause	Action To Take
Unable to complete the operation because an error occurred during a SCSI command. For detailed information, examine the RAID Manager Activity Log.	There is a software problem or a host adapter problem.	Read the RAID Manager Activity Log to determine the exact problem that occurred.
Unable to copy download file from disk. File may not be present, or SYS volume is full.	SYS volume is out of space, or there is a problem with the DOS copy.	First Option: Check the SYS volume to determine if it is full, and make some space on the disk if it is out of space. Second Option: Fix the DOS copy problem.
Unable to copy file form flexible disk. DOS must be present to perform the copy.	DOS is not available and must be for the firmware download.	<ol style="list-style-type: none">Reboot your system.At the DOS prompt, enter <i>server -na</i>, then press the Enter key.Download the firmware again.

12.4 Array Parity Check/Repair for SCO UNIX

The RAID Parity Check/Repair utility runs a parity check and/or repair on a given LUN. Parity Check/Repair applies only to RAID 1 and RAID 5. RAID 0 does not have array parity. Parity check on a RAID 1 configuration simply compares the data on the mirrored drives. Parity Check/Repair cannot be run on a degraded RAID 1 or RAID 5 LUN. See the discussion at the beginning of Section 12.2 for further definition of parity check and repair.

12.4.1 Getting Started

The RAID Parity Check/Repair can be invoked in one of two ways: through the shell script or by invoking the executable directly. The shell script method prompts you for options and invokes the monitor executable. The executable requires switch options to be specified, and is provided in the event that you would prefer to include it in customized shell scripts.

12.4.1.1 Using the Shell Script

The RAID Parity Check/Repair utility is invoked by typing *paritysh* at the system prompt. The following output appears:

```
Parity Check/Repair Utility Shell Script  (C) Digital Equipment Corporation 1992  
Version XX.X
```

```
Enter values for the following parameters or hit <CR>  
for default value:
```

```
Enter SCSI ID [0]:
```

The format of the input prompt line shows the parameter to enter, followed by square brackets containing the default value of that parameter. Press the Return key or the Enter key to select the default value, or enter some other value for that parameter.

A full help menu that describes the utility options is available by typing *paritysh -h* at the system prompt. If help information is needed for a specific parameter, enter *h* when that parameter is prompted for, and a description is displayed.

Once all the parameters have been specified, the shell script invokes the parity executable, redirecting the output to a specified log file, and piping the execution statement to the SCO UNIX *at* utility. This schedules the Parity Check/Repair to occur at the specified time.

The Parity Check/Repair runs as a background process. The first phase reads each block on the LUN and verifies that the parity is *good*. It creates a unique file in */tmp*, which lists the blocks that had *bad* parity. If a repair is to be done, this file is opened, and the parity of the blocks in that file is repaired.

12.4.1.2 Invoking the Executable

The RAID Monitor can also be started by invoking the executable directly. Enter *parity [options]* to start execution. The *-h* option prints out a help menu with a description of the various switches.

Preventive and Proactive Maintenance

12.4 Array Parity Check/Repair for SCO UNIX

12.4.2 Options

The options available to both the shell script and executable are outlined in Table 12-5.

Table 12-5 Parity Check/Repair Utility Options

Options	Symbol	Description
SCSI ID	-iID	The SCSI ID of the StorageWorks RAID Array 110 Subsystem. The value must be from 0 to 7, with a default value of 0.
SCSI LUN	-lLUN	The SCSI LUN. This value is not used at this time and should be set to the default value of 0.
Device File	-dFILENAME	<p>The device file name for access to the KZESA Host Adapter.</p> <ul style="list-style-type: none">For the shell script, the default is 0 for /dev/dptH0, which is the device file for the KZESA host adapter as a primary controller. Enter 1 for /dev/dptH1, which is the secondary controller.Invoking the executable, this option does not need to be specified for the default value of 0. If the secondary controller is required, either -d/dev/dptH1 or -1 can be specified.
Check only	-c	Select parity check option only (no repair).
Time	n/a	Time of day to run the Parity/Check utility. Used by the paritysh only. If no time is specified, then the utility runs within the next few minutes. Invoking the executable causes the Parity Check/Repair to run in the foreground immediately, unless piped to <i>at</i> or some other scheduling utility.
Log File Name	n/a	Name of the log file to write errors. Used by the paritysh only. The user must manually redirect the output to a file if the executable is invoked.

12.4.3 Output

The output goes to the screen or to a log file depending on how the utility is run. The output is shown below:

```
Parity Check/Repair Utility  (c) Digital Equipment Corporation 1992
Version X.XX
```

```
Parity Check begun timestamp
```

```
Comparing block 0 H to FFFF H
```

```
.
.
.
```

```
Parity Check completed timestamp
```

```
Number of blocks with bad parity: 0
```

```
Parity Repair begun timestamp
```

```
Parity Repair completed timestamp
```

```
Number of repaired blocks: 0
```

The output lists the range of blocks whose parity is checked. If at any time a block with bad parity is found, the following message appears:

```
*** Bad parity found at block a ****
```

The total number of blocks with bad parity is displayed after the check is complete. The repair portion of the display shows any blocks repaired with the following message:

```
Repairing parity at block a
```

A final count of all blocks repaired is given at the end of the display.

12.4.4 Error Messages

The error messages returned by the Parity Check/Repair utility use the following format:

- For general error messages:

```
PARITY: message
```

Generically, if an error occurs during a command to the controller that results in either a KZESA Host Adapter error, or a SCSI status other than *good*, that information is presented to the user in the following format:

- SCSI Check Condition Status (2 H)

```
SCSI Check Condition
SCSI Sense Key: 6 H  ASC: 29 H  ASCQ: 0 H
```

In this example, the SCSI status is a Check Condition (2 H), the Sense Key is 6 H, the Additional Sense Code (ASC) is 29 H, and the Additional Sense Code Qualifier (ASCQ) is 0 H indicating a power on, reset or bus device reset occurred. Appendix A lists all the ASCs, ASCQs and Sense Keys valid for the array controller. Note that these values are given in hexadecimal notation.

Preventive and Proactive Maintenance

12.4 Array Parity Check/Repair for SCO UNIX

- Host Adapter error or SCSI Status (other than *Check Condition*)

SCSI Status: 0 H
Host Adapter Status: 7 H

In this example, the SCSI status is 0 H and the Host Adapter Status is 7 H indicating a bus parity error. Appendix A lists all the SCSI status codes for the array controller as well as the host adapter error codes.

The following list describes the general error messages returned from the utility:

Error Opening Controller Device File x—Results from problems opening /dev/dptH0 or /dev/dptH1. Verify that those device files exist.

Error issuing read capacity to device—Array controller was unable to respond to the request sense command.

Error issuing verify command to device—Array controller was unable to respond to the verify command.

Error repairing parity on device—An error occurred while repairing the parity blocks.

Illegal switch - x—An illegal option was specified.

Error Opening Bad Parity Blocks File—An error occurred while attempting to open the bad parity blocks file in /tmp.

Error Reading Bad Parity Blocks File—An error occurred during the repair phase trying to open the bad parity blocks file in /tmp.

12.5 RAID Monitor Utility for SCO UNIX

The RAID Monitor performs the following functions:

The RAID Monitor notifies the user of any change of status in the LUN or devices in the array. Messages can appear on /dev/console and/or in a file depending on option selection.

12.5.1 Getting Started

The RAID Monitor can be invoked in one of two ways: through the shell script or by invoking the executable directly. The shell script method prompts the user for options and invokes the monitor executable. The executable requires switch options to be specified, and is provided in the event that the user would prefer to include it in customized shell scripts.

12.5.1.1 Using the Shell Script

The RAID Monitor can be invoked by typing *monitorsh* at the system prompt. The following output appears:

```
RAID Monitor Utility Shell Script      (C) Digital Equipment Corporation 1992  
Version XX.X
```

```
Enter values for the following parameters or hit <CR>  
for default value:
```

```
Enter SCSI ID [0]:
```

The format of the input prompt line shows the parameter to enter, followed by square brackets containing the default value of that parameter. Press the Return key or the Enter key to select the default value, or enter some other value for that parameter.

A full help menu that describes the utility options is available by typing *monitorsh -h* at the system prompt. If help information is needed for a specific parameter, enter *h* when that parameter is prompted for, and a description is displayed.

Once all the parameters have been specified, the shell script invokes the monitor executable, redirecting information output to /dev/console, and error information to a log file.

The monitor runs as a background process in the system. When the monitor is started, a status of the LUNs and drives is displayed. This information is similar to the information displayed when the RAID status utility is displayed on the screen.

12.5.1.2 Invoking the Executable

The RAID Monitor can also be started by invoking the executable directly. Enter *monitor [options]* to start the execution. The *-h* option prints out a help menu with a description of the various switches. It is the user's responsibility to redirect standard output and standard error if necessary.

Preventive and Proactive Maintenance

12.5 RAID Monitor Utility for SCO UNIX

12.5.2 Options

The options available to both the shell script and executable are outlined in Table 12-6.

Table 12-6 RAID Monitor Options

Options	Symbol	Description
SCSI ID	-iID	The SCSI ID of the StorageWorks RAID Array 110 Subsystem. The value must be from 0 to 7, with a default value of 0.
SCSI LUN	-lLUN	The SCSI LUN. This value is not used at this time and should be set to the default value of 0.
Device File	-dFILENAME	The device file name for access to the KZESA Host Adapter. <ul style="list-style-type: none">• For the shell script, the default is 0 for /dev/dptH0, which is the device file for the KZESA host adapter as a primary controller. Enter 1 for /dev/dptH1, which is the secondary controller.• If the executable is invoked, this option does not need to be specified for the default value of 0. If the secondary controller is required, either -d/dev/dptH1 or -1 can be specified.
Time	-tTIME	Interval of time to update status. This time is specified in minutes.
Log File Name	n/a	Name of the log file where errors are written. Used by the monitorsh only. The user must manually redirect the output to a file if the executable is invoked.

12.5.3 Error Messages

The error messages returned by the RAID Monitor use the following format:

- For general error messages:

MONITOR: message

- For a change in LUN status:

```
MONITOR: Lun 0 Status Changed Thu Oct 15 13:11:19 EDT 1992
MONITOR: Drives: (1,2) (3,4)
MONITOR: Previous status: Optimal
MONITOR: New status      : Degraded
```

The LUN is identified in the first message with a time-stamp of when the change was displayed. The next line shows the drives associated with that LUN referenced by channel and SCSI ID of the drive. In the above example, 2 drives are associated with LUN 0, a drive in channel 1 and at SCSI ID 2, and at channel 3, and SCSI ID 4. Finally a description of the previous status and new status is displayed.

- For a change in drive status:

```
MONITOR: Drive Channel 1 ID 2 Status Changed Thu Oct 15 13:11:19 EDT 1992
MONITOR: Previous status: Optimal
MONITOR: New status      : Degraded: Drive failure
```

The drive is identified by channel and SCSI ID with a time-stamp of when the change was displayed. This is followed by a description of the previous status, and now the new status.

Generically, if an error occurs during a command to the controller that results in either a KZESA Host Adapter error, or a SCSI status other than *good*, that information is presented in the following format:

- SCSI Check Condition Status (2 H)

```
SCSI Check Condition
SCSI Sense Key: 6 H  ASC: 29 H  ASCQ: 0 H
```

In this example, the SCSI status is a Check Condition (2 H), the Sense Key is 6 H, the Additional Sense Code (ASC) is 29 H, and the Additional Sense Code Qualifier (ASCQ) is 0 H indicating a power on, reset, or bus device reset occurred. Section A.2 lists all the ASCs, ASCQs and Sense Keys valid for the array controller. Note that these values are given in hexadecimal notation.

- Host Adapter error or SCSI Status (other than *Check Condition*)

```
SCSI Status: 0 H
Host Adapter Status: 7 H
```

In this example, the SCSI status is 0 H and the host adapter status is 7 H indicating a bus parity error. Appendix A lists all the SCSI status codes for the array controller. Appendix B lists all the KZESA host adapter error codes.

The following list describes the general error messages returned from the utility:

Error Opening Controller Device File x Monitor process terminating—Results from problems opening /dev/dptH0 or /dev/dptH1. Verify that those device files exist.

Error issuing inquiry command to device Monitor process terminating—Array controller was unable to respond to the inquiry command. Verify that the array controller is booted and the drives are spun up.

Lun x Status Changed Not Configured—LUN originally was configured, but no longer responds as configured.

Lun x Status Changed Configured—LUN originally was not configured, but now responds as if configured.

Error issuing mode sense command to logical array page Monitor process terminating—Unable to successfully issue a mode sense command to the array controller, and is terminating the monitor process.

Error issuing mode sense command to physical array page Monitor process terminating—Unable to successfully issue a mode sense command to the array controller, and is terminating the monitor process.

Illegal switch - x Monitor process terminating—An illegal option was specified.

12.6 RAID Status Utility for SCO UNIX

The RAID status utility displays the current LUN status and device status.

12.6.1 Getting Started

The RAID status utility can be invoked in one of two ways: through the shell script or by invoking the executable directly. The shell script method prompts the user for options and invokes the monitor executable. The executable requires switch options to be specified, and is provided in the event that the user prefers to include it in customized shell scripts.

12.6.1.1 Using the Shell Script

The RAID status utility can be invoked by typing *raidstatsh* at the system prompt. The following banner appears:

```
RAID Status Utility Shell Script      (C) Digital Equipment Corporation 1992
Version XX.X
```

```
Enter values for the following parameters or hit <CR> for
default value:
```

```
Enter SCSI ID [0]:
```

The format of the messages include the parameter to enter, followed by square brackets containing the default value of that parameter. Press the Return key or the Enter key to select the default value, or enter some other value for that parameter.

A full help menu that describes the utility options is available by typing *raidstatsh -h* at the system prompt. If help information for a specific parameter is needed, enter *h* when that parameter is prompted for, and a description is displayed.

Once all the parameters have been specified, the shell script invokes the *raidstat* executable, displaying all status information to the screen.

12.6.1.2 Invoking the Executable

The RAID status utility can also be started by invoking the executable directly. Enter *raidstat [options]* to start execution. The *-h* option prints out a help menu with a description of the various switches.

12.6.2 Options

The options available to both the shell script and executable are outlined in Table 12-7.

Table 12-7 RAID Status Utility Options

Options	Symbol	Description
SCSI ID	-iID	The SCSI ID of the StorageWorks RAID Array 110 Subsystem. The value must be from 0 to 7, with a default value of 0.
SCSI LUN	-lLUN	The SCSI LUN. This value is not used at this time and should be set to the default value of 0.
Device File	-dFILENAME	The device file name for access to the KZESA Host Adapter. <ul style="list-style-type: none"> For the shell script, the default is 0 for /dev/dptH0, which is the device file for the KZESA host adapter as a primary controller. Enter 1 for /dev/dptH1, which is the secondary controller. Invoking the executable, this option does not need to be specified for the default value of 0. If the secondary controller is required, either -d/dev/dptH1 or -1 can be specified.

12.6.3 Output

The status information displayed by the RAID status utility is shown below:

RAID Status Utility Version X.XX (C) Digital Equipment Corporation 1992
LUN STATUS

Lun	RAID	Status	Drives (Chan, ID)	Size	Type
0	RAID 5	Optimal	(2,0) (3,0) (4,0)	3318041	Reg
1	RAID 1	Optimal	(1,0) (5,0)	1658880	Sub

DRIVE STATUS

ID	Channel	1	2	3	4	5
0	0	0	0	0	0	0
1	1	1	1	1	1	1
2	1	1	1	1	1	1
3	1	1	1	1	1	1
4	1	1	1	1	1	1
5	1	1	1	1	1	1
6	1	1	1	1	1	1
7	7	7	7	7	7	7

0-opt 1-no-exist 2-spare 3-fail 4-repl 5-warn 6-mismatch 7-ctrl 8-fmt 9-wrong

The LUN status is displayed first. Each LUN is listed, followed by its RAID number, and status. The drives associated with the LUN are listed next as (x,y) where x is the channel number, and y is the SCSI ID of the drive. The size is displayed as the number of blocks in this LUN. The LUN type is either regular or sub-LUN.

Preventive and Proactive Maintenance

12.6 RAID Status Utility for SCO UNIX

The drive status shows the status of each drive referenced in the matrix by channel and SCSI ID. The definition of the statuses is listed in the last line of the display (see Section 4.4.3.2 for an explanation of the drive statuses).

12.6.4 Error Messages

The error messages returned by the RAID status utility use the following format:

- For general error messages:

RAIDSTAT: message

Generically, if an error occurs during a command to the controller that results in either a KZESA Host Adapter error, or a SCSI status other than *good*, that information is presented in the following format:

- SCSI Check Condition Status (2 H)

```
SCSI Check Condition
SCSI Sense Key: 6 H  ASC: 29 H  ASCQ: 0 H
```

In this example, the SCSI status is a Check Condition (2 H), the Sense Key is 6 H, the Additional Sense Code (ASC) is 29 H, and the Additional Sense Code Qualifier (ASCQ) is 0 H indicating a power on, reset or bus device reset occurred. Appendix A lists all the ASCs, ASCQs and Sense Keys valid for the array controller. Note that these values are given in hexadecimal notation.

- Host Adapter error or SCSI Status (other than *Check Condition*)

```
SCSI Status: 0 H
Host Adapter Status: 7 H
```

In this example, the SCSI status is 0 H and the Host Adapter Status is 7 H indicating a bus parity error. Appendix A lists all the SCSI status codes for the array controller as well as the host adapter error codes.

The following list describes the general error messages returned from the utility:

Error Opening Controller Device File x—Results from problems opening /dev/dptH0 or /dev/dptH1. Verify that those device files exist.

Error issuing inquiry command to device—Array controller was unable to respond to the inquiry command. Verify that the array controller is booted and the drives are spun up.

Lun x Status Changed Not Configured—LUN originally was configured, but no longer responds as configured.

Lun x Status Changed Configured—LUN originally was not configured, but now responds as if configured.

Error issuing mode sense command to logical array page—Unable to successfully issue a mode sense command to the array controller.

Error issuing mode sense command to physical array page—Unable to successfully issue a mode sense command to the array controller.

Illegal switch - x—An illegal option was specified.

SCSI-2 Error Codes

This appendix contains the following information:

- SCSI-2 status codes
- SCSI sense keys and additional sense codes

A.1 SCSI-2 Status Codes

This section lists the supported SCSI-2 status codes for the array controller. Refer to Table A-1.

Table A-1 SCSI-2 Status Codes for the Array Controller

Status Description	Code
Good	00 H
Check Condition	02 H
Busy	08 H
Reservation Conflict	18 H
Queue Full	28 H

A.2 SCSI Sense Keys and Additional Sense Codes

This section lists the SCSI sense keys, additional sense code (ASC), and additional sense code qualifier (ASCQ) values that the controller returns in the sense data. SCSI-2-defined codes are used when possible. Array-specific error codes are used when necessary and are assigned SCSI-2 vendor-unique codes 80-FFH.

Some errors that could be mapped into SCSI-2-defined ASC bytes are assigned vendor-unique codes to provide more error information than the SCSI-2 defined ASCQ bytes. In these cases, the SCSI-2-defined byte is returned in the ASC field and the more detailed error information byte is returned in the ASCQ field as a vendor-unique qualifier of the standard ASC. These ASCQ bytes have values of 80-FFH.

Table A-2 defines the SCSI sense keys.

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-2 SCSI Sense Keys

Sense Key	Description
0x00	No Sense
0x01	Recovered Error
0x02	Not Ready
0x03	Medium Error
0x04	Hardware Error
0x05	Illegal Request
0x06	Unit Attention
0x07	Data Protect (returned only if the drive returns this sense key)
0x08	Blank Check (not used)
0x09	Vendor-Specific (not used)
0x0A	Copy Aborted (not used)
0x0B	Aborted Command
0x0C	Equal (not used)
0x0D	Volume Overflow (not used)
0x0E	Miscompare (not used)
0x0F	Reserved (not used)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 defines the additional sense codes (ASC) and additional sense code qualifiers (ASCQ) affiliated with a given sense key.

Table A-3 SCSI Error Codes

ASC	ASCQ	Sense Key	Description
00	00	0	No Additional Sense Information
00	06	0	I/O Process Terminated
01	00	1,4	No Index/Sector Signal
02	00	1,4	No Seek Complete
03	00	1,4	Peripheral Device Write Fault
04	00	1,2	Logical Unit Not Ready, Cause Not Reportable
04	01	1,2	Logical Unit Is In Process Of Becoming Ready
04	02	1,2	Logical Unit Not Ready, Initializing Command Required
			Start Unit command required
04	03	1,2	Logical Unit Not Ready, Manual Intervention Required
04	04	1,2	Logical Unit Not Ready, Format In Progress
04	80	2	Logical Unit Not Ready, ROM Installed Does Not Support Redundant Controller Configuration
05	00	1,2	Logical Unit Does Not Respond To Selection
06	00	1,3,4	No Reference Position Found
07	00	1,4	Multiple Peripheral Devices Selected
08	00	1,4	Logical Unit Communication Failure
08	01	1,4	Logical Unit Communication Timeout
08	02	1,4	Logical Unit Communication Parity Error
09	00	1,4	Track Following Error
0A	00	1,6	Error Log Overflow
0C	01	1	Write Error Recovered With Auto Reallocation
0C	02	1,3,4	Write Error—Auto Reallocation Failed
10	00	1,3,4	ID CRC or ECC Error
11	00	1,3,4	Unrecovered Read Error
			The description of this error as being unrecovered may be misleading because a drive-reported Unrecovered Read Error can be recovered by the array using parity. The sense key for this case is Recovered and the Recovery Actions field in the sense indicates that parity was used.
11	01	1,3,4	Read Retries Exhausted
11	02	1,3,4	Error Too Long To Correct
11	04	1,3,4	Unrecovered Read Error—Auto Reallocate Failed
11	0A	1,4	Miscorrected Error

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
11	0B	1,3	Unrecovered Read Error—Recommend Reassignment
11	0C	1,3	Unrecovered Read Error—Recommend Rewrite The Data
12	00	1,3	Address Mark Not Found For ID Field
13	00	1,3	Address Mark Not Found For Data Field
14	00	1,3	Recorded Entity Not Found
14	01	1,3	Record Not Found
15	00	1,3,4	Random Positioning Error
15	01	1,4	Mechanical Positioning Error
15	02	1,3,4	Positioning Error Detected By Read Of Medium
16	00	1,3,4	Data Synchronization Mark Error
17	00	1	Recovered Data With No Error Correction Applied
17	01	1	Recovered Data With Retries
17	02	1	Recovered Data With Positive Head Offset
17	03	1	Recovered Data With Negative Head Offset
17	05	1	Recovered Data Using Previous Sector ID
17	06	1	Recovered Data Without ECC—Data Auto-Reallocated
17	07	1	Recovered Data Without ECC—Recommend Reassignment
18	00	1	Recovered Data With Error Correction Applied
18	01	1	Recovered Data With Error Correction And Retries Applied
18	02	1	Recovered Data—Data Auto-Reallocated
18	05	1	Recovered Data—Recommend Reassignment
19	00	1,3	Defect List Error
19	01	1,3	Defect List Not Available
19	02	1,3	Defect List Error In Primary List
19	03	1,3	Defect List Error In Grown List
1A	00	1,5	Parameter List Length Error
1B	00 ¹	1,4	Synchronous Data Transfer Error
1C	00	0,1,3	Defect List Not Found
1C	01	0,1,3	Primary Defect List Not Found
1C	02	0,1,3	Grown Defect List Not Found
1D	00	1,E	Miscompare During Verify Operation
1E	00	1	Recovered ID With ECC Correction
20	00	1,5	Invalid Command Operation Code
21	00	1,5	Logical Block Address Out Of Range
22	00	1,5	Illegal Function

¹The FRU code in byte 14 of the sense data indicates where the error was detected.

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
24	00	1,5	Invalid Field In CDB
25	00	1,5	Logical Unit Not Supported
26	00	5	Invalid Field In Parameter List
26	01	5	Parameter Not Supported
26	02	5	Parameter Value Invalid
26	03	5	Threshold Parameters Not Supported
27	00	1,7	Write Protected
28	00	6	Not Ready To Ready Transition (Medium May Have Changed)
29	00 ¹	6	Power On, Reset, Or Bus Device Reset Occurred
2A	00	6	Parameters Changed
2A	01	6	Mode Parameters Changed
2A	02	6	Log Parameters Changed
2C	00 ¹	1,4	Command Sequence Error
2F	00	6	Commands Cleared By Another Initiator
30	00	1,3	Incompatible Medium Installed
30	01	1,3	Cannot Read Medium - Unknown Format
30	02	1,3	Cannot Read Medium - Incompatible Format
31	00	1,3	Medium Format Corrupted
31	01	3,4	Format Command Failed
31	80	5	Format Command Cannot Execute Because Of Multiple LUN Assignments Per Physical Drive
32	00	3,4	No Defect Spare Location Available
32	01	3,4	Defect List Update Failure
37	00	1	Rounded Parameter
39	00	5	Saving Parameters Not Supported
3A	00	1,2	Medium Not Present
3D	00 ¹	1,5	Invalid Bits In Identify Message
3E	00	1,2	Logical Unit Has Not Self-Configured Yet
3F	00	6	Target Operating Conditions Have Changed
3F	01	6	Microcode Has Been Changed Returned when new drive microcode has been downloaded.
3F	02	6	Changed Operating Definition
3F	03	6	Inquiry Data Has Changed

¹The FRU code in byte 14 of the sense data indicates where the error was detected.

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
3F	80 ³	(6) ²	<p>Drive Failed Because Of A Failed Write Operation</p> <p>This is the result of an error that prevents access to user data on this drive or that may result in erroneous data being read/written (for example, an Unrecovered Write Error).</p>
3F	81 ³	(6) ²	<p>Drive Failed—Automatic Reallocation Failed</p> <p>Either reassign block or recovery of data on the reassigned sector failed.</p>
3F	82 ³	(6) ²	<p>Drive Failed—Reconstruction Failed</p> <p>(Error On Drive Being Reconstructed)</p> <p>The reconstruction operation cannot complete because of an unrecoverable Write error or a failed drive format (before reconstruction starts) on the drive being reconstructed. It must be replaced (again) before the reconstruction can be retried.</p>
3F	83 ³	(6) ²	<p>Drive Failed—Reconstruction Failed</p> <p>(Error On Drive Required For Reconstruction)</p> <p>The reconstruction operation cannot complete because of an unrecoverable Read error on one of the drives needed for reconstruction. You can still run the array in degraded mode, but if user data is stored at the address in error, you can not read it. You should attempt a backup of the array to recover the remaining data. To bring the LUN to an optimal state, replace the failed drive, format the array, and restore the data from a backup copy (the one just made or a previous successful backup copy).</p>
3F	84 ³	(6) ²	Drive Failed Due To A Hardware Component Diagnostics Failure
3F	85 ³	(6) ²	Drive Failed Because It Failed A Test Unit Ready Command (during start-of-day) Or Read Capacity Command (during start-of-day or during a format or reconstruction operation)
3F	86 ³	(6) ²	Drive Failed Because It Failed A Format Unit Command
3F	87 ³	(6) ²	Drive Failed By A Host Mode Select Command
3F	88 ³	(6) ²	Drive Failed Because Of Deferred Error Reported By The Drive
3F	89 ³	(6) ²	Drive Failed By Start-of-Day Application Code Because Of A Drive Replacement Error
3F	90	(6) ²	<p>Unrecovered Read/Write Error</p> <p>Unrecovered Read errors always generate this condition.</p>

²The Unit Attention sense key in parentheses (6) indicates that this code depends on the Sense Key for Drive Failures field in the EEPROM data.

³These conditions (3F 8x) require that the drive be replaced as soon as possible. User data is still available by using parity, but a second failure at this time (when a drive has been downed) results in the loss of user data for the entire LUN (RAID 5) or results in an unrecoverable error (other RAID levels).

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
			Unrecovered Write errors cause this condition if the state of the LUN or the RAID level guarantees that data at the other addresses will not be affected by the failed operation and continues access to this drive.
			A second drive failure or warning condition on another drive can prevent reconstruction of this drive.
3F	91	(6) ²	Drive Reported Deferred Error Caused Drive To Be Placed In Warning
3F	AX	(6) ²	Single Drive Array Assurance Error * X = 0—Drive ECC Test Failed * X = Other—Not Currently Implemented
3F	B0 ⁴	(6) ²	Excessive Media Error Rate
3F	B1 ⁴	(6) ²	Excessive Seek Error Rate
3F	B2 ⁴	(6) ²	Excessive Grown Defects
3F	C0 ⁵	(6) ²	No Response From One Or More Drives
3F	C1 ⁵	(6) ²	Communication Errors
3F	C2 ⁵	(6) ²	Firmware Indicates No Drive Is Present Although Information Stored On Disk Indicates Drive Should Be Present
3F	C7	(6) ²	Subsystem Component Failure FRU code and qualifiers identify the failed component.
3F	C8	(6) ²	AC Power Is Lost, DC Power Is Being Supplied By A Hold-Over Battery
3F	C9	(6) ²	AC Power Is Lost, DC Power Supplied For A Maximum Of Two Minutes The host should start its shutdown procedure.
3F	CA	(6) ²	AC Power Is Lost, DC Power Is Exhausted The controller completes currently executing drive commands to maintain data integrity.
3F	CB	(6) ²	AC Power Was Lost, But Is Now Restored
3F	D0 ⁶	(6) ²	>75% Of Transfer Delays Caused By One Drive
3F	D1 ⁶	(6) ²	>75% Of Non-Aligned Reselections Caused By One Drive
3F	D2 ⁶	(6) ²	Synchronous Transfer Value Differences Between Drives
3F	D3 ⁶	(6) ²	Software Measured Performance Degradation
3F	D4 ⁶	(6) ²	Mode Parameter Differences Between Drives May Cause Performance Degradation

²The Unit Attention sense key in parentheses (6) indicates that this code depends on the Sense Key for Drive Failures field in the EEPROM data.

⁴These errors are warnings to allow early drive replacement (drive still being used until a failed Write of user data—3F 80).

⁵These errors are warnings to check installation (drive still being used until a failed Write operation—3F 80).

⁶These errors are indications of potential performance degradation (drive still being used).

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
3F	E0 ⁷	(6) ²	LUN Downed
3F	E1 ⁷	4	Multiple Drives Have Been Downed
3F	E2 ⁷	4	Mode Parameters For Drives In LUN Do Not Match
3F	E3 ⁷	4	Drive Channel Verification Failed
3F	E4 ⁷	4	SCSI ID Verification Failed
3F	E5 ⁷	4	Wrong Drive Was Replaced
3F	E6 ⁷	4	Component Failure Affecting Multiple Channels
3F	F0 ⁸	(6) ²	EEPROM Error
3F	F1 ⁸	(6) ²	EEPROM Hard Checksum Error
3F	F2 ⁸	(6) ²	Maximum EEPROM Write Count Exceeded
3F	F8	(6) ²	Application Software Copy To More Than One Drive Failed. Another download is required if the drive with the current application code failed.
40	00	4	RAM Failure
40	NN	4	Diagnostic Failure On Component NN (80H-FFH) In this release of the software, this ASC and ASCQ are only returned if reported by a drive. FRU codes distinguish between failures on components of the array controller and components of the drive controllers.
41	00 ¹	1,4	Data Path Failure, Carryover From CCS (same as 40 NN)
42	00	4	Power-On Or Self-Test Failure, Carryover From CCS (same as 40 NN)
43	00 ¹	1,4	Message Error
44	00	1,4	Internal Target Failure
45	00 ¹	1,4	Select/Reselect Failure
46	00 ¹	4	Unsuccessful Soft Reset
47	00 ¹	1,B	SCSI Parity Error
48	00	1,4	Initiator Detected Error Message Received
49	00 ¹	1,4	Invalid Message Error
4A	00 ¹	1,4	Command Phase Error
4B	00 ¹	4	Data Phase Error
4B	80	4	Data Overrun/Underrun

¹The FRU code in byte 14 of the sense data indicates where the error was detected.

²The Unit Attention sense key in parentheses (6) indicates that this code depends on the Sense Key for Drive Failures field in the EEPROM data.

⁷Subsequent media access commands to this LUN do not execute and return an ASC of Command Cannot Execute Because The LUN Has Been Downed with an ASCQ that indicates why the LUN is down.

⁸The configuration change (drive being marked failed/warning) could not be saved to EEPROM, but the update of the RAM information on the board and the DACSTORE information on the disk were successful. The new configuration information is used until a board power-up/reset. At that time, a diagnostic failure is detected that prevents the board from being used.

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
4C	00	1,4	Logical Unit Failed Self-Configuration
4E	00	1,B	Overlapped Commands Attempted
5A	00	6	Operator Request Or State Change Input (Unspecified)
5A	02	1,7	Operator Selected Write Protect
5A	03	6	Operator Selected Write Permit
5B	00	1,6	Log Exception
5B	01	1,6	Threshold Condition Met
5B	02	1,6	Log Counter At Maximum
5B	03	1,6	Log List Codes Exhausted
5C	00	1,6	RPL Status Change
5C	01	1,6	Spindles Synchronized
5C	02	1,6	Spindles Not Synchronized
80	00	4	Error Manager Detected Error
80	01	4	The Error Manager Was Invoked Without Any EM_Code(s) Loaded
80	02 ¹	4	The Error Manager Was Passed An Out-Of-Range Code
80	03 ¹	4	The Error Manager Was Passed A Code By The SCSI Driver, But No Error Handler Exists
80	04	4	Fatal Null Pointer
80	05	(6) ²	No AEN Code Or An Invalid AEN Code Was Loaded By The Application Software This error is detected when the AEN condition is requested by the command handler (either when the next command for this host/LUN is received or when a polled Request Sense is received), not when the AEN condition was detected. The illegal AEN code is loaded into the Error-Specific Information field in the sense data.
80	06	4	Maximum # Of Errors For This I/O Exceeded
80	07	4	Drive Reported Recovered Error Without Transferring All Of Data
81	00	4	Reconstruction Setup Failed
82	00	B	Out Of Heap
82	01	B	No Command Control Structures Available
82	02	B	No DAC Application Control Blocks Available
83	00	B	Reservation Conflict
84	00	4	Command Cannot Execute Because The LUN Has Been Downed

¹The FRU code in byte 14 of the sense data indicates where the error was detected.

²The Unit Attention sense key in parentheses (6) indicates that this code depends on the Sense Key for Drive Failures field in the EEPROM data.

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
84	01 ^{1,9}	4	Multiple Drives Have Been Downed
84	02 ^{1,9}	4	Mode Parameters For Drives in LUN Do Not Match
84	03 ⁹	4	Drive Channel Verification Failed
84	04 ⁹	4	SCSI ID Verification Failed
84	05 ⁹	4	Format In Progress
84	06 ⁹	4	Awaiting Format Command
84	08 ⁹	4	Wrong Drive Was Replaced
84	09 ⁹	4	Component Failure Affecting Multiple Channels
85	00	4	General Application Code Command Handler Error
85	01 ¹⁰	4	Drive Error
85	02 ¹⁰	4	Host Error
85	03 ¹⁰	4	Drive Type Mismatch Within LUN
85	04 ¹⁰	4	Operation Not Allowed During Reconstruction
85	05 ¹⁰	4	Data Returned By Drive Is Invalid
85	06 ¹⁰	4	Non-Failed Drive Unavailable For Operations
85	07 ¹⁰	4	Insufficient Rank Structures Available
85	08 ¹⁰	4	Full Format Required, But Not Allowed (Sub-LUNs)
85	09 ¹⁰	4	Drive Cannot Be Mode Selected To Meet LUN Parameters
85	0A ¹⁰	4	Data Recovery After Re-Assign Block Command Failed
85	0B	4	Drive Not Returning Required Mode Sense Page(s) Drives must support Mode Sense pages 3 and 4.
86	00	4	Command Cannot Execute Because The LUN Is In Degraded Mode
87	00	4	Code Download/Upload Error
87	01	4	Partial Download (Missing Application Code Segment)
87	02	(6) ²	Downloaded Code Cannot Be Saved To Disk Old application code is uploaded on next power-up /reset.
87	03	4	Code CRC Failure Check is performed on the disk read after the new code is saved to disk.
87	04	(6) ²	Upload Of Latest Version Of Code Failed An older version of code has been uploaded successfully.

¹The FRU code in byte 14 of the sense data indicates where the error was detected.

²The Unit Attention sense key in parentheses (6) indicates that this code depends on the Sense Key for Drive Failures field in the EEPROM data.

⁹These errors (84 XX) are returned on every media access command to the downed LUN until the error is corrected.

¹⁰The original CDB field in the array sense data shows the failing command. Debug information is shown in the Error Detection Point sense field.

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
87	05	4	No Package Verification Partition Downloaded
87	06	4	ROM Partitions Required For Download Of Code Missing
87	07	4	Incomplete RAM Partitions
87	08	4	Incompatible Board Type For The Code Downloaded
87	09	4	Incompatible ROM Version For Support Of The Downloaded Code
87	0A	1	Download Of Microcode To A Failed Disk Completed Successfully The drive is still unusable until it is marked as replaced and has been reconstructed.
88	00	4	EEPROM Command Error
88	02	4	EEPROM Not Responding (There is an EEPROM on the board.)
88	03	4	EEPROM Not "Formatted" (certain key fields have not been set.)
88	04	4	Invalid EEPROM Offset (a Write to the write-protected maintenance area causes this error.)
88	05	4	EEPROM Soft Checksum Error (indicates that an update was interrupted)
88	06	4	EEPROM Hard Checksum Error (indicates a component may be going bad)
88	07	4	Maximum Write Count Exceeded
88	09	4	EEPROM Not Initialized
89	00	4	Error On Request Sense Command To A Drive
8A	00	5	Illegal Command For Pass-Through Mode The original CDB field in the array sense data shows the failing command.
8A	01	5	Illegal Command For Current RAID Level
8B	00	B	Write Buffer Command (For Code Download) Was Attempted While Another Command Was Active
8B	01	5	Write Buffer Command (For Drive Microcode Download) Attempted But The Enable Bit In The EEPROM Was Not On
8C	00	4	Destination Transfer State Machine Error
8C	01	4	Invalid Transfer Release Requester
8C	02	4	Invalid Transfer Requester
8C	03	4	Data Stripe/Parity Generation ASIC Configuration Error
8C	04	4	Data Transfer Request Error
8C	05	4	Invalid Transfer Pad Requester
8D	00	B,4	Destination Driver Data Transfer Did Not Complete

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
8E	00	B	Data Stripe/Parity Generation ASIC Error
8E	01	E	Parity/Data Mismatch
8E	02	B	Data Underrun
8F	00	B	Premature Completion Of A Drive Command (expected Data Transfer and received Good Status instead)
90	XX	4	DACSTORE Errors DACSTORE Error Bit Definitions Bits 0—3 = DACSTORE Identifier * Value of 0 = DACSTORE directory * Value of 1 = Disk store * Value of 2 = LUN store * Value of 3 = Controller store * Value of 4 = Log store * Value of 5 = High ID controller serial # store * Value of 6 = Low ID controller serial # store * Value of 7 = Redundant controller common store * Value of 8 = Boot block EEPROM store Bits 4—6 = Error Type * Value of 0 = Setup error * Value of 1 = Invalid directory data * Value of 2 = Drive error * Value of 3 = Invalid store data Bit 7 = Operation Type * Value of 0 = Read * Value of 1 = Write
91	00	5	Mode Select Errors
91	01	5	LUN Already Exists; Cannot Do "Add LUN" Function "80"
91	02	5	LUN Does Not Exist; Cannot Do "Replace LUN" Function "83" Or Any Logical Function
91	03	5	Drive Already Exists; Cannot Do "Add Drive" Function "80"
91	04	5	Drive Does Not Exist; Cannot Do Requested Action For It
91	05	5	Drive Cannot Be Deleted; It Is Part Of A LUN
91	06	5	Drive Cannot Be Failed; It Is Formatting
91	07	5	Drive Cannot Be Replaced; It Is Not Marked As Failed Or Replaced
91	08	5	Invalid Action To Take

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
91	09	5	Invalid Action With Multiple Sub-LUNs Defined (probably an attempt to change page 3—Format Device page)
91	0A	5	Invalid Reconstruction Amount
91	0B	5	Invalid Reconstruction Frequency
91	0C	5	Invalid LUN Block Size
91	0D	5	Invalid LUN Type
91	0E	5	Invalid Segment Size
91	0F	5	Invalid Segment 0 Size
91	10	5	Invalid Number Of Drives In LUN
91	11	5	Invalid Number Of LUN Blocks
91	12	5	Invalid RAID Level
91	13	5	Invalid Drive Sector Size
91	14	5	Invalid LUN Block Size/Drive Sector Size Modulo
91	15	5	No Disks Defined For LUN
91	16	5	Insufficient Rank Structures Available To Define LUN
91	17	5	Disk Defined Multiple Times For LUN
91	18	5	Sub-LUN Drives Not The Same As Those Used By Other Sub-LUNs On These Drives
91	19	5	Sub-LUN RAID Level Mismatch
91	1A	5	First Sub-LUN Defined For These Drives Has Not Yet Been Formatted; Second Sub-LUN Is Illegal
91	1B	5	Nonsub-LUN Drive Already Owned By Another LUN
91	1C	5	Sub-LUN Drive Already Owned By a Nonsub-LUN
91	1D	5	Drive Type Does Not Match The Drive Type Of The Other Drives In The LUN
91	1E	5	Drive Cannot Be Included In Rank Because Rank Is Full
91	1F	5	Ranks Have Different Number Of Disks Defined
91	20	5	Multiple Disks On Same Channel Within Same Rank
91	21	5	Mirrored Disks On The Same Channel
91	22	5	No Parity Disk Defined
91	23	5	No Data Disks Defined
91	24	5	Too Many Disks Defined
91	25	5	No Space Available For LUN—Sub-LUN Cannot Be Defined
91	26	5	Drive Status Cannot Be Changed To Good (drive cannot be revived through Mode Select)
91	27	5	Error In Processing A Subsystem Mode Page
91	28	5	Drive Inquiry Data Mismatch Between Drives In The LUN

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
91	29	5	Drive Capacity Mismatch Between Drives In The LUN
91	2A	5	Drive Block Size Mismatch Between Drives In The LUN
91	2B	5	Support Of TTD/CIOP Messages Is Not The Same For All Drives In The LUN
91	2C	5	Firmware Does Not Support Redundant Controller Options Selected
92	00	B	BUSY Status From Drives Could Not Be Cleared By Array Controller
93	00	4	Drive Vendor Unique Sense Data Returned See the drive sense area of the array sense for the drive ASC/ASCQ and sense key. Also refer to the drive documentation for further description.
94	00	5	Invalid Request Of A Controller In Redundant Controller Mode
95	00	1,B	A Drive Channel Was Reset Probable cause is the removal or replacement of a drive during a hot swapping operation.
95	01	1,B	An Extended Drive Channel Reset Has Been Detected Probable cause is a drive left partially removed or inserted.
96	00	6	Redundant Controller Not Supported By Current Firmware
96	01	4	Alternate Controller Not Supported By Current Firmware (however, the alternate controller has been detected)
B0	00	B	Command Timeout
B0	01	B	Watchdog Timer Timeout
B0	02	B	Software Loop Timeout
D0	00 ¹	4	SCSI Driver Timeout
D0	01 ¹	4	Disconnect Timeout
D0	02 ¹	4	Chip Command Timeout
D0	03 ¹	4	Byte Transfer Timeout
D1	00 ¹	4	Bus Errors
D1	01 ¹	4	CDB Transfer Incomplete
D1	02 ¹	4	Unexpected Bus Phase
D1	03 ¹	4	Disconnect Expected
D1	04 ¹	4	ID Message Not Sent
D1	05 ¹	4	Synchronous Negotiation Error
D1	06 ¹	4	Target Transfer Disable (TTD) Negotiation Conflict
D1	07 ¹	4	Unexpected Disconnect

¹The FRU code in byte 14 of the sense data indicates where the error was detected.

(continued on next page)

SCSI-2 Error Codes

A.2 SCSI Sense Keys and Additional Sense Codes

Table A-3 (Cont.) SCSI Error Codes

ASC	ASCQ	Sense Key	Description
D1	08 ¹	4	Unexpected Message
D1	09 ¹	4	Unexpected Tag Message
D1	0A ¹	4	Channel Busy
D2	00 ¹	4	Miscellaneous SCSI Driver Error
D2	01 ¹	4	Illegal C96 Chip Command
D2	02 ¹	4	Uncoded Execution Path
D3	00 ¹	4	Drive SCSI Chip Reported Gross Error
D4	00 ¹	4	Non-SCSI Bus Parity Error
D5	00 ¹	4	Miscellaneous Host-Related Errors
D5	01 ¹	4	Maximum Messages Received
D5	02 ¹	4	Message Reject Received on a Valid Message
D6	00 ¹	4	Source Driver Chip-Related Error
D7	00 ¹	4	Source Driver Programming Error
D8	00	4	An Error Was Encountered That Required The Data Pointers To Be Restored But The Host Is Non-Disconnecting And Does Not Support The Restore Pointers Message (indicated by an EEPROM option control bit or by host selection without sending the identify message)

¹The FRU code in byte 14 of the sense data indicates where the error was detected.

B

KZESA Host Adapter Error Codes

This appendix lists the error codes for the KZESA Host Adapter. Controller errors have priority over SCSI errors. Refer to Table B-1.

Table B-1 SCSI Status Codes for the KZESA Host Adapter

Error Description	Code
No error	00 H
Selection timeout	01 H
Command timeout	02 H
SCSI bus reset	03 H
Initial controller power-up	04 H
Unexpected bus phase	05 H
Unexpected bus free	06 H
Bus parity error	07 H
SCSI bus hung	08 H
Unexpected message reject	09 H
Reset stuck	0A H
Auto request sense failed	0B H
Controller RAM parity error	0C H
Abort message processed	0D H
Reset message processes	0E H

C

Total Call Concept (TCC)

The *Total Call Concept* does not apply to this product.

D

Supported Options

This appendix includes information regarding the disk drives and adapters supported by the system. Table D-1 describes the disk drives supported by the system.

Table D-1 Disk Drives Supported

Component	Description
RZ25	426 Mbyte SCSI Disk Drive
RZ26	1.05 Gbyte SCSI Disk Drive

Table D-2 describes the adapters supported by the system.

Table D-2 Adapters Supported

Component	Description
KZESA	EISA/SCSI Differential Host Adapter
DWZZA-VA	Single-ended to Differential/Wide SCSI Adapter

E

Host SCSI Cables

Table E-1 lists the SCSI cables used with specific host computer systems and the StorageWorks shelves. Be aware that this table does not address all possible applications.

Table E-1 SCSI Cable Selections

From	SCSI Adapter	To	SCSI Cable
DECpc Desktop Computers			
DECstation 425ST	KZESA	HSZ10-AA	BN21K, BN21P-0B, H879-AA
applicationDEC 4xxMP Systems			
applicationDEC 400MP	KZESA	HSZ10-AA	BN21K, BN21P-0B, H879-AA
applicationDEC 433MP	KZESA	HSZ10-AA	BN21K, BN21P-0B, H879-AA
MicroVAX 3100 Systems			
Model 30	Embedded	BA350-EA	BC09D
Model 40	Embedded	BA350-EA	BC09D
Model 80	Embedded	BA350-EA	BC09D
Model 90	Embedded	BA350-EA	BC09D
VAXstation 4000 Systems			
Model 60	Embedded	BA350-EA	BC09D
Model 90	Embedded	BA350-EA	BC09D
VAX 4000 Systems			
Model 100	Embedded	BA350-EA	BC09D

Glossary

ac distribution

The method of controlling ac power in a cabinet.

adapter

- (1.) A connecting device that permits the attachment of accessories or provides the capability to mount or link units.
- (2.) The device that connects an 8-bit differential SCSI bus to an 8-bit single-ended SCSI bus.

array

A set of multiple disk drives and a specialized controller, an *array controller*, which keeps track of how the data is distributed across the drives.

array channels

The SCSI-2 compliant buses on which the disk drives are located. Each array channel is an independent SCSI bus.

array controller

A device that exercises control over the SCSI bus (for example, an HSZ10-AA disk array controller).

BA35X-VA

A collective reference to all versions of the vertical mounting kits—single and double.

Battery Backup Unit (BBU)

StorageWorks power unit option that provides sufficient power to prevent storage devices from losing data in the event of a shelf power unit failure.

Note

The BBU does not provide power for the operation of a storage device.
The BBU provides power only for protecting data.

CI

A Digital trademark for the Digital Computer Interconnect bus.

cold-swapping

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

See also warm-swapping and hot-swapping.

controller

A hardware line device that manages communications over a line. Controllers can be point-to-point, multipoint, or multiple line.

dc power system

The method for providing dc power in a cabinet.

double stand

A BA35X-VA vertical mounting kit composed of two single stands clipped together. This configuration can support one BA350-EA shelf.

See also single stand.

drive group

A set of 1 to 10 drives that have been configured into one or more logical units. A logical unit can be contained in only one drive group, and all the logical units in a drive group must have the same RAID level and be of the same drive type.

drive rank

Drive ranks represent a numbering scheme that provides information on the maximum number of drives on every array channel. A one-rank system indicates that there is a maximum of one drive per disk channel. A two-rank array indicates that there is a maximum of two drives per disk array channel. However, any channel can have zero for its maximum number.

DSSI

Digital Storage System Interconnect.

FD SCSI

The fast, differential SCSI bus with an 8-bit data transfer rate of 10 MB/s.

See also FWD SCSI and SCSI.

FWD SCSI

The fast, wide, differential SCSI bus with a 16-bit data transfer rate of 20 MB/s.

See also FD SCSI and SCSI.

H981X

A collective reference to the H9810 (short), H9811 (medium), and H9812 (tall) towers.

Heartbeat LED

The bottommost LED on the HSZ10-AA controller. Beats once per second.

host

The primary or controlling computer in a multiple computer network.

hot-swapping

A method of device replacement whereby the complete system remains online and active during device removal or insertion. The device being removed or inserted is the only device that cannot perform operations during this process.

See also cold-swapping and warm-swapping.

LUNs (logical units)

A *logical unit* is a grouping of drives that has its own device SCSI ID and number. Each logical unit has its own array parameters (RAID level, segment size, and so on). For most purposes, a logical unit is equivalent to an array.

mirrored

A copy of data on a disk or a set of disks. Refer to the description of RAID 1.

parity check/repair

The process of verifying and repairing parity information so that data can be maintained and reconstructed in the event of a drive failure. Parity Check/Repair functionality is provided by the DEC RAID utilities.

RAID

A redundant array of independent disks.

rank

The number of drives per channel. *See also* drive rank.

reconstruction

The process that restores a degraded RAID 1 or 5 logical unit to its original state after a single drive has been replaced.

redundancy

Also *data redundancy*. Data stored on another physical disk that can be used to recover data if the physical disk containing the data cannot be accessed.

SBB

System building block. A modular carrier plus the individual mechanical and electromechanical interface required to mount it into a standard shelf. Any device conforming to shelf mechanical and electrical standards is considered an SBB.

SCSI

Small Computer System Interface. This interface defines the physical and electrical parameters of a parallel I/O bus used to connect computers and a maximum of seven SBBs. The StorageWorks modular storage system implementation uses SCSI-2, which permits the synchronous transfer of 8-bit data at rates of up to 10 MB/s.

segment

A group of blocks that is continuous data which can be stored on a disk drive.

shelf array

A modular storage shelf that provides power, cooling, interconnects, and mounting for SBBs. Specific shelves are denoted by the prefix BA350 (that is, BA350-RA, BA350-SA, and so on). Shelves may be mounted in kits, towers, or cabinets.

single stand

A reference to the basic BA35X-VA vertical mounting kit with a capacity of one BA350-SA shelf.

See also double stand.

Small Computer System Interface

See SCSI.

stands

A collective reference to all versions of the vertical mounting kits—both single and double.

static storage device (SSD)

An electronic storage device such as the EZ51R-VA.

StorageWorks

The mnemonic for the Digital Storage/Modular Enclosure, a modular set of enclosure products that allows customers to design their own storage array. Components include power, packaging, and interconnections in a modular storage shelf into which SBBs (system building blocks) and array controller modules are integrated to form modular storage arrays. System-level enclosures to house the arrays and standard mounting devices for SBBs are also included.

striped

See the description of RAID 0 in the StorageWorks RAID Array 110 Subsystem User's Guide.

system building block

See SBB.

towers

A collective reference to the H9810 (short), H9811 (medium), and H9812 (tall) towers.

warm-swapping

A method of device replacement in which the complete system remains online during device removal or insertion. The system bus may be halted for a brief period of time, during device insertion or removal. No booting or loading of code is permitted except on the device being inserted.

See also cold-swapping and hot-swapping.

Index

A

A cable

- part number, 11-5

ACE

- See Array Configuration Editor (ACE)

- creating a LUN, 4-9

- deleting a LUN, 4-7

- Drive window, 4-5

- exiting, 4-12

- exiting without a LUN 0, 4-9

- help, 4-6

- logical unit display, 4-5

- starting, 4-5

ACE logical units

- table of status, 4-30

Adapters description, D-1

Adding

- an HSZ10-AA controller, 8-7

- drive, 8-14

- logical units, 5-11

- logical units to existing volumes, 5-11

Additional sense code (ASC), A-1, A-3

Additional sense code qualifier (ASCQ), A-1, A-3

Advanced configurations, 10-1

AMD error messages, 12-10, 12-12

- console, 12-11

- log message, 12-11

- pop-up, 12-10

APC command, 4-33

APR command, 4-34

Array

- channels, 2-6

- definition, 2-1

- features, 1-3

Array Configuration Editor (ACE)

- creating a LUN, 4-9

- deleting a LUN, 4-7

- description, 4-24

- drive status, 4-28

- Drive window, 4-5

- exiting, 4-12

- exiting without a LUN 0, 4-9

- getting help, 4-6

- Logical Unit window, 4-5

- navigating in, 4-26

- quitting, 4-27

Array Configuration Editor (ACE) (cont'd)

- starting, 4-5, 4-24

- table of keys, 4-26 to 4-27

- window, 4-25

- window features, 4-4

Array controller

- features, 1-9

- SCSI-2 status codes, A-1

Array Controller Configuration Window, 4-16

- features, 4-17

- figure, 4-16

Array Monitor Daemon, 12-7

- description, 12-7

- error messages, 12-7

- how it works, 12-7

Array parity

- checking/repairing, 12-19

- checking/repairing for the StorageWorks RAID

- Array 110 Utility for Novell NetWare, 12-1

Array Parity Check Facility (APC), 4-32

Array Parity Repair Facility (APR), 4-32

Array performance, 5-11

Array Status Monitor (ASM), 4-35

ASM command, 4-36

Audience, xiii

B

BA350-EA shelf

- description, 1-7

BA350 shelf

- part numbers, 11-1

BA35X-VA description, 1-7, 11-1

Basic configuration

- modifying, 10-1

Blocks

- per delay interval, 5-19

Blower replacement, 8-10

Bus connections

- end-bus, 3-5

- midbus, 3-3

Bus continuity

- maintaining, 3-3

C

Cabinets

- BA35X-VA, 1-7
- part numbers, 11-1

Cable options, 1-11

Cables

- A cable part number, 11-5
- for host, E-1
- part numbers, 11-4
- P cable, 3-1, 3-2
- P cable part number, 11-4
- SCSI Host, E-1
- verifying, 3-6
- Y cable, 3-2
- Y cable part number, 11-6

Capacity, 1-5

Changing

- RAID level, 5-27

Cold swapping, 8-9

Command

- Install for NetWare, 5-11

Commands

- APC, 4-33
- APR, 4-34
- ASM, 4-36

Component descriptions, 1-6

Configurations

- advanced, 10-1
- guidelines, 10-2
- multiple rank, 10-2

Configuring

- logical units, 5-12

Connecting

- to a 16-bit differential host/adaptor, 3-2
- to a host, 3-1
- to an 8-bit differential host/adaptor, 3-2
- to an 8-bit single-ended host/adaptor, 3-3

Connectors, verifying, 3-6

Controller

- part number, 11-2

Cost

- of single-drive failure, 1-6

Creating

- logical unit, 5-22

Creating a LUN

- with ACE, 4-9

Custom expansion, 10-22

D

Data

- availability, 1-5
- mirrored, 2-2
- reliability, 1-4
- segmented, 2-2

Delay interval, 5-19

Deleting

- a LUN with ACE, 4-7
- drive, 8-15
- logical unit, 8-11

Device names, 4-22

- for NetWare, 4-22
- for the StorageWorks RAID Array 110 Utility
- for MS-DOS, 4-22

Device numbers

- table of, 4-22

Disk array controller, 1-7

- features, 1-8

Disk array description, 2-1

Disk Array Services Window, 4-14

- features, 4-15

- figure, 4-14

Disk array subsystems descriptions, D-1

Disk drives

- part numbers, 11-2

Displaying

- drive status, 8-14
- parameters, 5-20

Documentation

- related, xv

Drive

- adding, 8-14
- deleting, 8-15
- failing, 8-15
- groups, 2-7
- how to replace, 8-8
- ranks, 2-8
- reconstructing, 8-16
- SBB LED indicators, 5-9
- status, 4-20, 5-3
- when to replace, 8-21
- window, 4-5

Drive failures

- multiple drive failures - RAID 1, 8-24
- multiple drive failures - RAID 5, 8-25
- on RAID 0, 8-23
- on RAID 1, 8-24
- on RAID 5, 8-25
- single drive failure, 8-24
- single drive failure - RAID 0, 8-23
- single drive failure - RAID 5, 8-25

Drive Map, 5-17

- changing, 5-17

Drive Matrix, 4-27

Drive status

- displaying, 8-14
- modifying, 8-14
- table of, 4-21

Drive window, 4-5, 4-27

- options, 4-28

DWZZA-VA, 2-16

- cable requirements, 11-5
- description, 3-3

DWZZA-VA (cont'd)

- features, 2-16
- figure, 3-3
- required, 3-3, 11-6

E

- Environmental specifications, 2-12
- Environmental stabilization, 2-11
- Error codes, 7-1, A-3
 - KZESA host adapter, B-1
- Error handling, 6-1
- Error log
 - example, 12-8
- Error messages
 - AMD format, 12-10
 - array component errors, 12-12
 - array device errors, 12-14
 - for StorageWorks RAID Array 110 Utility for Novell NetWare installation, 3-13 to 3-14
 - for the Parity Check/Repair utility, 12-21
 - for the RAID Monitor, 12-24
 - for the RAID status utility, 12-28
 - from the Array Monitor Daemon, 12-7
 - status change errors, 12-16
 - tables of, 7-1
 - types, 12-10
- Errors
 - reported by OpenVMS VAX, 7-7
 - table of, 7-4 to 7-7
 - reported by the StorageWorks RAID Array 110 Utility for OpenVMS VAX, 7-3
 - reporting, 7-3
- Exiting
 - from ACE, 4-12
 - from ACE without a LUN 0, 4-9
- Expansion paths
 - custom expansion, 10-22
 - from 1 BA350-EA/2 BA350-SA to 1 BA350-EA/4 BA350-SA, 10-17
 - from 1 BA350-EA to 1 BA350-EA/2 BA350-SA, 10-5
 - from 1 BA350-EA to 1 BA350-EA/4 BA350-SA, 10-11
 - recommended, 10-4
 - table of, 10-4
 - types of, 10-3
- Expansion unit description, 1-10

F

- Failing
 - drive, 8-15
- Fan
 - part number, 11-3
- Features
 - array, 1-3
 - of ACE, 4-4

Features (cont'd)

- subsystem, 1-4
- Field descriptions, 7-11
- Filler panels
 - part numbers, 11-3
- Firmware, 1-9, 11-6
- Flexibility, 1-5

G

- Getting started, 4-3

H

- Hardware, 11-1
 - requirements, 11-6
- Help
 - with ACE, 4-6
- Hot swapping, 8-9
- HSZ10-AA controller
 - adding a controller, 8-7
 - description, 1-7
 - features, 1-8
 - LED codes, 5-5
 - LED indicators, 5-5
 - location and SCSI address, 2-15
 - monitoring through, 5-1
 - replacing, 8-1

I

- Installation, 3-1
 - before you install, 3-10
 - error messages for NetWare, 3-13 to 3-14
 - running installation program, 3-10
 - StorageWorks RAID Array 110 Utility for MS-DOS for MS-DOS, 3-15
 - StorageWorks RAID Array 110 Utility for Novell NetWare, 3-8
 - StorageWorks RAID Array 110 Utility for SCO UNIX, 3-15
 - using an MS-DOS partition, 3-10
 - without an MS-DOS partition, 3-11
- Install command
 - for NetWare, 5-11

K

- Kits, 11-8
 - MS-DOS, 3-15, 11-8
 - NetWare, 11-9
 - OpenVMS VAX, 11-8
 - SCO UNIX, 11-10
- KZESA host adapter error codes, B-1

L

LEDs

- drive SBB, 5-9
- HSZ10-AA controller, 5-5
- HSZ10-AA controller codes, 5-5
- monitoring through, 5-5
- power supply, 5-7
- power supply and shelf, 5-8
- shelf and power supply, 5-8
- types of, 5-5

Logical block size, 5-18

Logical drive map, 2-16

Logical unit

- adding to existing volumes, 5-11
- adding to NetWare, 5-11
- configuration summary, 5-12
- configuring, 5-12
- creating, 4-9
- creating from an existing drive, 5-25
- creating from spare drives, 5-22
- definition, 2-2
- deleting, 4-7, 8-11
- description, 2-6
- display, 4-5, 4-29
- displaying parameters, 5-20
- displaying parameters procedure, 5-20 to 5-22
- modifying parameters, 5-20
- modifying parameters procedure, 5-20 to 5-22
- parameters, 5-15
- parameters, table of, 5-16
- procedure for restoring, 8-19 to 8-21
- restoring, 8-19
- restoring RAID 0, 8-27
- restoring RAID 1, 8-29, 8-31
- restoring RAID 5, 8-33, 8-35
- size, 5-18
- status, 4-20, 4-30
- status, table of, 4-20, 8-22 to 8-23

Logical Unit window, 4-29

- options, 4-31
- options list, 4-31

LUN

See Logical unit

LUN status, 5-2

M

Menu

- Array Controller Configuration, 4-16
- Disk Array Services, 4-14

Messages

- for the RAID Manager, 12-16 to 12-18
- tables of, 7-1

Mirrored data, 2-2

Modifying

- drive status, 8-14
- logical units, 5-20
- parameters, 5-20

Monitoring

- operations with LED indicators, 5-5
- operations with the StorageWorks RAID Array 110 utilities, 5-4
- through the HSZ10-AA controller, 5-1
- through the StorageWorks shelf, 5-4
- user monitoring methods, 5-4

MS-DOS, 11-7

- installation, 3-15
- installation with a partition, 3-10
- installation without a partition, 3-11
- kit, 3-15, 11-8

N

Naming conventions, 4-3

NetWare, 11-7

- adding logical units, 5-11
- device names, 4-22
- files for RAID Manager, 3-9
- installation, 3-8
- installation error messages, 3-13 to 3-14
- kit, 11-9
- logical unit configuration tasks, 5-13 to 5-14
- starting, 4-13
- using, 4-2
- using the Install command, 5-11

NetWare Utility

- description, 4-1

O

OpenVMS VAX, 11-7

- kit, 11-8
- logical unit configuration tasks, 5-15

OpenVMS VAX Utility

- description, 4-1
- using, 4-2

Operations, 5-1

Options

- from the Disk Array Services window, 4-14
- from the Drive window, 4-28
- from the Help window in ACE, 4-6
- from the Logical Unit window, 4-31

P

Parameters

- for logical units, 5-15

Parity

- array parity, 12-19
- array parity for the StorageWorks RAID Array 110 Utility for Novell NetWare, 12-1
- check, 4-32, 4-33

Parity (cont'd)

- repair, 4-32, 4-34

Parity check/repair

- functions, 4-12, 4-32

Parity Check/Repair

- automatic, 12-2
- automatic operation, 12-2
- automatic schedule changes, 12-6
- changing the schedule, 12-6
- description, 12-2
- error messages, 12-21
- functions, 12-2, 12-19
- manual, 12-4
- manual operation, 12-4
- options, 12-20
- output, 12-21
- running the check, 12-4 to 12-5
- SCO UNIX installation, 3-16
- starting, 12-19
- when to run, 12-2

Partitions, 2-8

Parts list

- adapter accessories, 11-3
- BA350 backplane SCSI components, 11-3
- BA350 shelves, 11-1
- BC09D-series SCSI transition cables, 11-6
- BN21H-series SCSI A cables, 11-4
- BN21J-series SCSI A cables, 11-4
- BN21K-series SCSI P cables, 11-4
- BN21L-series SCSI P cables, 11-4
- BN21M-series SCSI transition cables, 11-5
- BN21N-series SCSI transition cables, 11-5
- BN21P SCSI transition cables, 11-5
- BN21P SCSI Y cables, 11-6
- cabinets, 11-1
- controllers, 11-2
- disk drives, 11-2
- filler panels, 11-3
- power cords, 11-3
- power supply, 11-2
- replacement fans, 11-3
- SBBs, 11-2
- StorageWorks products, 11-1
- tape drives, 11-2

P cable

- part number, 11-4
- required, 3-1, 3-2

Performance, 1-5

- notes, 5-11

Physical drive map, 2-16

Power cords

- part numbers, 11-3

Powering on

- the subsystem, 3-6

Power supply

- how to replace, 8-9
- LED indicators, 5-7
- part numbers, 11-2

Power supply (cont'd)

- specifications, 2-10

Power supply and shelf

- LED indicators, 5-8

Power unit

- specifications, 2-10

Product

- attributes, 1-4
- description, 1-1
- highlights, 1-3
- overview, 1-1

Purpose, xiii

R

- RAID 0 description, 2-2

- RAID 1 description, 2-3

- RAID 3 description, 2-4

- RAID 5 description, 2-5

RAID level

- changing, 5-27
- description, 5-16

RAID Manager

- device names, 4-22
- files for NetWare, 3-9
- messages, 12-16 to 12-18
- Parity Check/Repair, 3-16
- table of tasks, 4-19

RAID Monitor utility, 12-23

- error messages, 12-24
- functions, 12-23
- installation, 3-17
- invoking, 12-19, 12-23
- invoking the executable, 12-23
- options, 12-24
- using the shell script, 12-23

RAID overview, 2-1

RAID status utility, 12-26

- error messages, 12-28
- functions, 12-26
- installation, 3-17
- invoking, 12-26
- invoking with the executable, 12-26
- invoking with the shell script, 12-26
- options, 12-27

Reconstructing a drive, 8-16

Reconstruction, 2-9, 8-16

- description, 8-25
- rate, 8-26
- rate, table of, 8-26

Redundancy, 1-4

Regeneration, 2-9

Related documents, xv

Replacing

- an HSZ10-AA controller, 8-1
- blower, 8-10
- drive, 8-21
- power supply, 8-9

- Reporting errors, 7-3
- Request sense data format, 7-9
- Requirements
 - hardware, 11-6
 - software, 11-7
- Restoring
 - a LUN, 8-19
 - RAID 0 logical unit, 8-27
 - RAID 1 logical unit, 8-29, 8-31
 - RAID 5 logical unit, 8-33, 8-35
 - task summary for a LUN, 8-19
- RMINST.NLM, 3-8

S

- SBBs, 1-9
 - adapter SBBs, 1-10
 - disk SBBs, 1-10
 - part numbers, 11-2
 - power SBBs, 1-10
- SCO UNIX, 11-7
 - installation, 3-15
 - installing the RAID Monitor, 3-17
 - install the Parity Check/Repair utility, 3-16
 - invoking RAID Monitor, 12-23
 - invoking RAID status with the executable, 12-26
 - invoking the RAID Monitor executable, 12-23
 - invoking the RAID status utility, 12-26
 - kit, 11-10
 - Parity Check/Repair, 12-19
 - RAID Monitor, 12-23
 - RAID Monitor error messages, 12-24
 - RAID Monitor options, 12-24
 - RAID status functions, 12-26
 - RAID status utility, 3-17, 12-26
 - RAID status utility error messages, 12-28
 - RAID status utility options, 12-27
 - shell script, 12-19
 - starting Parity Check/Repair, 12-19
 - using the shell script, 12-26
- SCSI
 - host cable selection, E-1
- SCSI-2 status codes
 - for the array controller, A-1
- SCSI driver support, 1-10
- SCSI error codes, A-3
- SCSI ID
 - examples, 8-4
- SCSI interconnects/host adapters, 1-11
- SCSI sense keys, A-1
- SCSI status codes
 - for the KZESA host adapter, B-1
- Segment, 2-2
- Segment size, 5-18
 - segment zero, 5-19

- Sense keys, 7-11
- Shelf and power supply
 - LED indicators, 5-8
- Shelf description
 - BA350-EA, 1-7
- Shelf power supply
 - replacing, 8-9
- Shell script, 12-19, 12-23, 12-26
- Software, 1-9
 - kit numbers, 11-7
 - requirements, 11-7
 - upgrading, 12-1
- Specifications
 - environmental, 2-12
 - environmental stabilization, 2-11
 - for StorageWorks RAID Array 110 Subsystem, 2-10
 - general, 2-10
 - power supply, 2-10
 - power unit, 2-10
- Stand-Alone, 11-7
 - device names, 4-22
 - logical unit configuration tasks, 5-12 to 5-13
 - starting, 4-13
 - using, 4-2
- Stand-Alone Utility
 - description, 4-1
- Starting
 - ACE, 4-5
- Startup, 4-12
 - for the NetWare Utility, 4-13
 - for the Stand-Alone Utility, 4-13
- Status
 - drive, 5-3
 - LUN, 5-2
 - of ACE logical units, 4-30
 - of drives, 4-20
 - of logical units, 4-20
 - tables of, 7-2
- StorageWorks RAID Array 110 Subsystem
 - base configuration, 2-13
 - capacity, 1-5
 - component descriptions, 1-6
 - data availability, 1-5
 - data reliability, 1-4
 - firmware, 11-6
 - flexibility, 1-5
 - logical view, 1-2
 - monitoring features, 5-1
 - performance, 1-5
 - physical arrangement of the base configuration, 2-16
 - physical specifications, 2-10
 - product attributes, 1-4
 - product description, 1-1
 - product highlights, 1-3
 - product overview, 1-1
 - redundancy, 1-4

StorageWorks RAID Array 110 Subsystem

Operations, 5-1

StorageWorks RAID Array 110 utilities, 1-10

description, 4-1, 11-6

device names, 4-22

kits, 11-8

menu functions, 4-14

monitoring through, 5-4

moving around, 4-14

MS-DOS kit, 11-8

NetWare, 11-7

NetWare kit, 11-9

OpenVMS VAX, 11-7

OpenVMS VAX kit, 11-8

part numbers, 11-7

road map, 4-19

SCO UNIX, 11-7

SCO UNIX kit, 11-10

Stand-Alone, 11-7

starting, 4-12

table of tasks, 4-19

using menus, 4-14

StorageWorks RAID Array 110 Utility for

MS-DOS, 11-7

description, 4-1

device names, 4-22

installation, 3-15

logical unit configuration tasks, 5-12 to 5-13

MS-DOS kit, 3-15

starting, 4-13

using, 4-2

StorageWorks RAID Array 110 Utility for Novell

NetWare

adding logical units, 5-11

checking/repairing array parity, 12-1

description, 4-1

installation, 3-8

installation error messages, 3-13 to 3-14

logical unit configuration tasks, 5-13 to 5-14

starting, 4-13

using, 4-2

StorageWorks RAID Array 110 Utility for

OpenVMS VAX, 11-7

description, 4-1

logical unit configuration tasks, 5-15

using, 4-2

StorageWorks RAID Array 110 Utility for SCO

UNIX

installation, 3-15

installing the RAID Monitor utility, 3-17

installing the RAID status utility, 3-17

install the Parity Check/Repair utility, 3-16

invoking RAID Monitor, 12-19, 12-23

invoking RAID status with the executable,
12-26

invoking the RAID Monitor executable, 12-23

invoking the RAID status utility, 12-26

Parity Check/Repair, 12-19

StorageWorks RAID Array 110 Utility for SCO UNIX (cont'd)

RAID Monitor, 12-23

RAID Monitor error messages, 12-24

RAID Monitor options, 12-24

RAID status functions, 12-26

RAID status utility, 12-26

RAID status utility error messages, 12-28

RAID status utility options, 12-27

shell script, 12-19

starting Parity Check/Repair, 12-19

using the shell script, 12-26

StorageWorks shelf

monitoring through, 5-4

Subsystem

description, 1-6

Subsystem features, 1-4

Swapping

cold, 8-9

hot, 8-9

warm, 8-9

T

Tape drives

part numbers, 11-2

Tasks

table of, 4-19

Temperature

ranges, 2-11

Terminating

the SCSI bus, 3-2

Thermal stabilization, 2-11

Troubleshooting, 6-1

table of, 6-2

U

Utility

APC, 4-32

APR, 4-32

ASM, 4-35

V

Verifying

cables and connectors, 3-6

functionality, 3-7

Vertical mounting kit description, 1-7

Volumes

adding logical units, 5-11

W

Warm swapping, 8-9

Y

Y cable

figure, 3-2

part number, 11-6

required, 3-2