

DEC RAID Subsystem

User's Guide

Order Number: EK-SZ200-UG-B01

Revision/Update Information: This new document supersedes the *DEC RAID Subsystem User's Guide*, Version 1.0.

**Digital Equipment Corporation
Maynard, Massachusetts**

May 1993

First Printing, November 1992

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.

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, DECserver, DECstation, DECwindows, OpenVMS, StorageWorks, 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, SmartCache, and SmartROM are registered trademarks of Distributed Processing Technology.
Intel is a registered trademark of Intel Corporation.
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 AT&T.
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	ix
1 Product Description	
1.1 Overview	1-1
1.2 Product Highlights	1-3
1.2.1 Array Features	1-3
1.2.2 Subsystem Features	1-3
1.3 Product Attributes	1-3
1.3.1 Data Reliability	1-3
1.3.2 Redundancy	1-4
1.3.3 Data Availability	1-4
1.3.4 Performance	1-4
1.3.5 Flexibility and Capacity	1-5
1.4 General Conclusions	1-5
2 DEC RAID Subsystem Component Descriptions	
2.1 Description of the SZ200	2-1
2.1.1 BA35X-VA Vertical Mounting Kit	2-1
2.1.2 BA350-EA Shelf	2-2
2.1.3 HSZ10-AA Controller (Disk Array Controller)	2-2
2.1.3.1 Features	2-3
2.1.3.2 Firmware	2-4
2.1.3.3 Software	2-4
2.1.4 System Building Blocks (SBBs)	2-4
2.2 Expansion Unit Description	2-5
2.3 DEC RAID Utilities	2-5
2.4 Enhanced SCSI Driver Support	2-5
2.5 SCSI Interconnects/Host Adapters	2-5
3 RAID Overview	
3.1 What Is a Disk Array?	3-1
3.1.1 Description of RAID 0	3-2
3.1.2 Description of RAID 1	3-3
3.1.3 Description of RAID 3	3-4
3.1.4 Description of RAID 5	3-5
3.2 Key Concepts	3-6
3.2.1 Array Channels	3-6
3.2.2 Logical Units (LUNs)	3-6
3.2.3 Drive Groups	3-7
3.2.4 Drive Ranks	3-7
3.2.5 Partitions	3-7

3.2.6	Reconstruction	3-8
3.2.7	Regeneration	3-8
4 SZ200 Base Configuration		
4.1	SCSI IDs for the DEC RAID Subsystem	4-1
4.2	HSZ10-AA Controller Location and SCSI Address	4-3
4.3	Physical Drive Map	4-4
4.4	Logical Drive Map	4-4
4.5	DWZZA-VA Bus Adapter	4-4
5 Basic Configuration Installation		
5.1	Connecting to a Host	5-1
5.1.1	Terminating the SCSI Bus	5-1
5.1.2	Connecting to a 16-Bit Differential Host/Adapter	5-2
5.1.3	Connecting to an 8-Bit Differential Host/Adapter	5-2
5.1.4	Connecting to an 8-Bit Single-Ended Host/Adapter	5-3
5.1.5	Maintaining Bus Continuity	5-4
5.2	Verifying Cables and Connectors	5-5
5.3	Powering On the Subsystem	5-6
5.4	Functional Verification	5-6
6 Operations		
6.1	DEC RAID Subsystem Monitoring Features	6-1
6.1.1	Monitoring Through the HSZ10-AA Controller	6-1
6.1.1.1	LUN Status	6-1
6.1.1.2	Drive Status	6-3
6.1.2	Monitoring Through the StorageWorks Shelf	6-4
6.2	User Monitoring Methods	6-4
6.2.1	Monitoring Operation Using the DEC RAID Utilities	6-4
6.2.2	Monitoring Operation Using LED Indicators	6-5
6.2.2.1	HSZ10-AA Controller LED Indicators	6-5
6.2.2.2	Power Supply LEDs	6-7
6.2.2.3	Drive SBB LEDs	6-8
6.3	When and How to Replace a Drive	6-10
6.4	When and How to Replace Power Supplies and Blowers	6-11
6.4.1	Replacing a Shelf Power Supply	6-11
6.4.2	Replacing a Blower	6-12
6.5	Reconstruction	6-13
6.6	Parity Check/Repair	6-14
6.7	Upgrading Software	6-15
7 Advanced Configurations		
7.1	Modifying Basic Configurations	7-1
7.2	Multiple Rank Configurations	7-2
7.3	Configuration Guidelines	7-2
7.3.1	Expansion Guidelines	7-3
7.4	Recommended Configurations	7-4

7.4.1	Expansion from a Single BA350-EA to a Single BA350-EA/Dual BA350-SA Configuration	7-5
7.4.1.1	Pre-Reconfiguration	7-5
7.4.1.2	Shelf Reconfiguration	7-6
7.4.1.3	Post-Reconfiguration	7-10
7.4.2	Expansion from a Single BA350-EA to a Single BA350-EA/Quad BA350-SA Configuration	7-11
7.4.2.1	Pre-Reconfiguration	7-11
7.4.2.2	Shelf Reconfiguration	7-12
7.4.2.3	Post-Reconfiguration	7-16
7.4.3	Expansion from a Single BA350-EA/Dual BA350-SA to a Single BA350-EA/Quad BA350-SA Configuration	7-17
7.4.3.1	Pre-Reconfiguration	7-17
7.4.3.2	Shelf Reconfiguration	7-18
7.4.3.3	Post-Reconfiguration	7-22
7.5	Custom Expansions	7-22

8 Error Handling/Troubleshooting

8.1	Before You Begin Troubleshooting	8-1
8.2	Using the Troubleshooting Table	8-1
8.3	If You Have Expanded Your DEC RAID Subsystem	8-3

A Physical Specifications

A.1	General Specifications	A-1
A.2	Power Unit Specifications	A-1
A.3	Environmental Stabilization	A-2
A.4	Environmental Specifications	A-3

B Supported Options

C Operating System Support

D System Connections

Glossary

Index

Figures

1-1	SZ200 DEC RAID Subsystem	1-1
1-2	Logical View of the DEC RAID Subsystem	1-2
2-1	SZ200 Base Unit	2-2
2-2	Physical Layout of the HSZ10-AA Controller	2-3
3-1	Conceptual Diagram of the DEC RAID Subsystem	3-1
3-2	Diagram of RAID 0	3-2
3-3	Diagram of RAID 1	3-3
3-4	Diagram of RAID 3	3-4

3-5	Diagram of RAID 5	3-5
3-6	Diagram of a Regular LUN	3-6
3-7	Drive Partition Diagram	3-8
4-1	DEC RAID Subsystem SCSI IDs	4-2
4-2	Physical Drive Map	4-3
5-1	Typical Y Cable Connection	5-2
5-2	BA350-EA and BA350-SA Combination Storage Array	5-3
5-3	Trilink Connector Midbus Connection	5-4
5-4	Trilink Connector End-Bus Connection	5-5
6-1	LUN State Diagram	6-2
6-2	LED Indicators	6-5
6-3	Power Supply Status LEDs	6-7
6-4	Shelf Status LEDs	6-9
6-5	Replacing a Drive	6-10
6-6	Replacing BA35X-MA Blowers	6-13
7-1	Single BA350-EA Shelf	7-6
7-2	Single BA350-EA to Single BA350-EA/Dual BA350-SA Shelf Reconfiguration	7-7
7-3	Single BA350-EA to Single BA350-EA/Dual BA350-SA Drive Reconfiguration	7-8
7-4	BA350-EA and BA350-SA Combination Storage Array	7-9
7-5	Single BA350-EA Shelf	7-12
7-6	Single BA350-EA to Single BA350-EA/Quad BA350-SA Shelf Reconfiguration	7-13
7-7	Single BA350-EA to Single BA350-EA/Quad BA350-SA Drive Reconfiguration	7-14
7-8	Single BA350-EA Shelf	7-18
7-9	Single BA350-EA/Dual BA350-SA Shelf to Single BA350-EA/Quad BA350-SA Shelf Reconfiguration	7-19
7-10	Single BA350-EA/Dual BA350-SA Shelf Configuration to Single BA350-EA/Quad BA350-SA Drive Reconfiguration	7-20

Tables

1	Conventions	x
5-1	Midbus Connections	5-4
5-2	End-Bus Connections	5-5
6-1	Logical Unit Status	6-2
6-2	Drive Status	6-3
6-3	Summary of HSZ10-AA Controller LED Codes	6-6
6-4	Shelf and Single Power Supply Status LEDs	6-8
6-5	Shelf and Dual Power Supply Status LEDs	6-8
6-6	Drive SBB Status LEDs	6-9
6-7	Blower Replacement	6-13
7-1	Expansion Paths	7-4
8-1	Troubleshooting Systems Problems	8-2
A-1	SZ200 General Specifications	A-1
A-2	StorageWorks Power Units	A-2

A-3	Thermal Stabilization Specifications	A-2
A-4	Environmental Specifications	A-3
B-1	Disk Drives Supported	B-1
B-2	Adapters Supported	B-1
C-1	Operating Systems Supported	C-1
D-1	System Connectivity	D-1

Preface

Purpose

This guide introduces the DEC RAID Subsystem storage solution. It describes available RAID features, packaging options, and configuration and operating instructions. This guide serves as complementary documentation to the StorageWorks™ documentation set.

Intended Audience

This guide is intended for customers of Digital Equipment Corporation, systems administrators, and Digital Service Engineers responsible for installing and maintaining the DEC RAID Subsystem products.

Structure

This guide is organized as follows:

Chapter 1, Product Description,
provides an overview of the DEC RAID Subsystem features and options.

Chapter 2, DEC RAID Subsystem Component Descriptions,
describes the DEC RAID Subsystem components including the base unit, expansion units, utilities, driver support, interconnect, and host adapters.

Chapter 3, RAID Overview,
presents background information on disk array levels and components.

Chapter 4, SZ200 Base Configuration,
describes the SZ200 base unit and its physical/logical mappings.

Chapter 5, Basic Configuration Installation,
provides installation instructions including the StorageWorks documentation references, host connections, and functional verification.

Chapter 6, Operations,
describes modes of operation and methods of monitoring the DEC RAID Subsystem.

Chapter 7, Advanced Configurations,
discusses advanced configurations and expansion of the base unit.

Chapter 8, Error Handling/Troubleshooting,
provides error handling and troubleshooting guidelines.

Appendix A, Physical Specifications,
describes the physical, environmental, and performance specifications for the DEC RAID Subsystem product.

Appendix B, Supported Options,
includes information regarding the disk drives and adapters supported by the DEC RAID Subsystem.

Appendix C, Operating System Support,
contains information regarding the operating systems supported by the DEC RAID Subsystem.

Appendix D, System Connections,
describes the type of connectivity for various systems to help you determine cabling requirements for the systems.

Glossary,
defines terms associated with RAID and SCSI technologies that are used throughout this guide and other related documentation.

Related Documents

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

Document Title	Order Number
<i>DEC RAID Utilities User's Guide</i>	EK-DECRA-UG
<i>DEC RAID OpenVMS VAX Utility User's Guide</i>	AA-PYZLA-TE
<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>HSZ10-AA Controller Site Preparation Guide</i>	EK-HSZ10-IN
<i>BA35X-VA Vertical Mounting Kit User's Guide</i>	EK-BA350SV-UG
<i>StorageWorks Shelf Building Block Subsystem Configuration Guide</i>	EK-BA350-CG
<i>BA350-EA Modular Storage Shelf User's Guide</i>	EK-350EA-UG
<i>BA350-SA Modular Storage Shelf User's Guide</i>	EK-350SA-UG

Documentation Conventions

The following is a list of the conventions used in this manual.

Table 1 Conventions

boldface type	Indicates the first instance of terms being defined in text.
<i>italic type</i>	Indicates emphasis and complete manual titles.

Product Description


This chapter contains the following information:

- Overview
- Product highlights
- Product attributes
- General conclusions

1.1 Overview

The DEC RAID Subsystem is a modular, integrated, end-user **RAID** (Redundant Array of Inexpensive Disks) solution based on the HSZ10-AA controller (disk array controller). Figure 1–1 shows the SZ200 DEC RAID Subsystem.

Figure 1–1 SZ200 DEC RAID Subsystem



SHR-XR3000-VIF

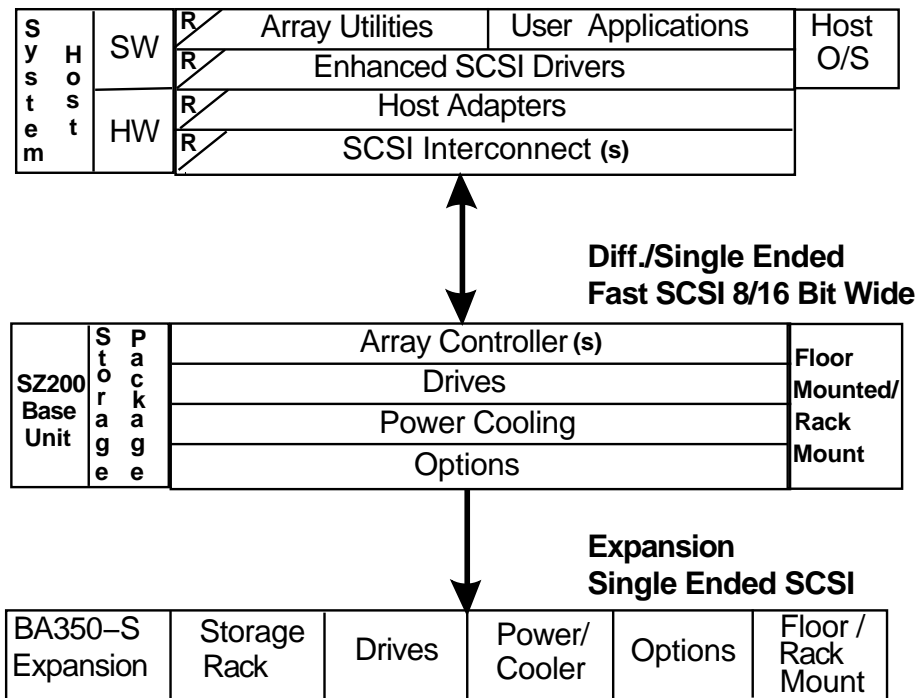
Product Description

1.1 Overview

The DEC RAID 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.

Figure 1-2 Logical View of the DEC RAID Subsystem



SHR-XR3002-GRA

Physically, the DEC RAID 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.

1.2 Product Highlights

This section includes the following information:

- Array features
- Subsystem features

1.2.1 Array Features

The following is a list of array features:

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

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

1.3 Product Attributes

This section includes information on the following attributes:

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

1.3.1 Data Reliability

Unlike a single SCSI disk drive solution, the DEC RAID 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 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.

Product Description

1.3 Product Attributes

In a DEC RAID 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.

1.3.2 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 DEC RAID 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.

1.3.3 Data Availability

One of the key attributes of the DEC RAID 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 whose livelihood depends upon minimal down time, these features are critical.

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, which 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 3 is most appropriate in 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.

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.

1.4 General Conclusions

The DEC RAID Subsystem provides a set of options from which users can choose. The options vary 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 1, 3, or 5) can provide.

Unlike multiple disks or RAID 0, the chance of losing data in a RAID 1, 3, or 5 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 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.

DEC RAID Subsystem Component Descriptions

This chapter contains a description of the DEC RAID Subsystem components as follows:

- SZ200 base unit
- Expansion units
- DEC RAID utilities
- Enhanced SCSI driver support
- SCSI Interconnects/Host adapters

2.1 Description of the SZ200

The SZ200 base unit is the basic building block for the DEC RAID Subsystem (and part of the StorageWorks family of products). The SZ200 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

For SZ200 physical specifications, refer to Appendix A. Figure 2–1 shows the SZ200 base unit and identifies its components.


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 SZ200 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.

DEC RAID Subsystem Component Descriptions

2.1 Description of the SZ200

Figure 2-1 SZ200 Base Unit



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 system building blocks (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.

Procedures for configuring and installing the BA350-EA shelf are discussed in the *BA350-EA Modular Storage Shelf User's Guide*.

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 levels 0, 1, 3, and 5. It supports expansion up to 35 devices (seven ranks).

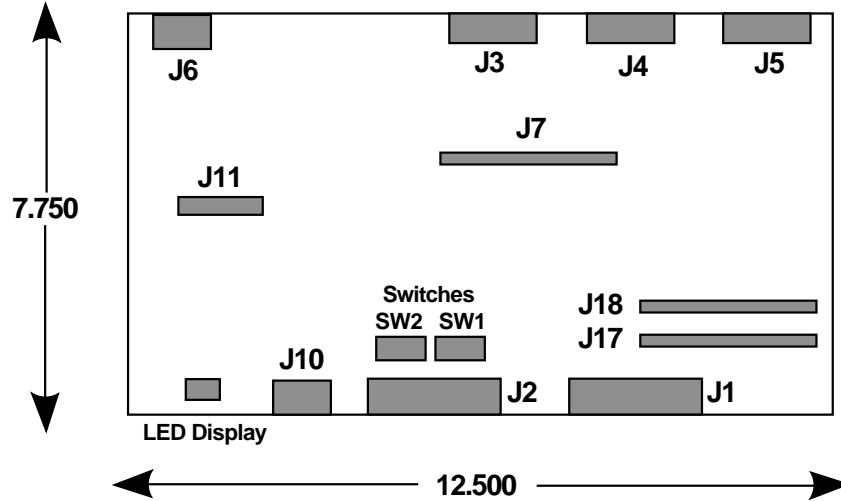
The host SCSI channel is a synchronous/asynchronous/fast, 8/16-bit differential. It supports SCSI-2 protocol. The host connection is made through a 68-pin, high-density connector (J1) as shown in Figure 2-2. A second connector (J2) is provided for daisy-chaining. SCSI term power is provided.

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

DEC RAID Subsystem Component Descriptions

2.1 Description of the SZ200

Figure 2-2 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

2.1.3.1 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
- Host SCSI target ID switch (4-bit)
- Status LEDs (8-bit)
- Self-test capability

DEC RAID Subsystem Component Descriptions

2.1 Description of the SZ200

Array 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 level 0, 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

2.1.3.2 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.

2.1.3.3 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.

2.1.4 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 removal 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.

Power SBBs

The SZ200 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 Subsystem Shelf Building Block Configuration Guide*.

Disk SBBs

The SZ200 base unit is configured with five 3-1/2 inch disk drives. For supported disk options, refer to Appendix B.

DEC RAID Subsystem Component Descriptions

2.1 Description of the SZ200

Adapter SBBs

The SZ200 unit provides a slot for a bus adapter option for connection to different host types. For example, the DWZZA-VA bus adapter, which is a single-ended differential SCSI converter, fits into the SZ200 unit.

Refer to Appendix B for supported adapter options.

2.2 Expansion Unit Description

Since the StorageWorks family of products is based on a building block concept, expansion of the SZ200 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 Subsystem Shelf Building Block Configuration Guide*.

2.3 DEC RAID Utilities

DEC RAID utilities are applications, stand-alone, 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 the *DEC RAID Utilities User's Guide*.

2.4 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

2.5 SCSI Interconnects/Host Adapters

The 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 included in Appendix B.

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

RAID Overview

This chapter contains information on RAID (Redundant Array of Inexpensive 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

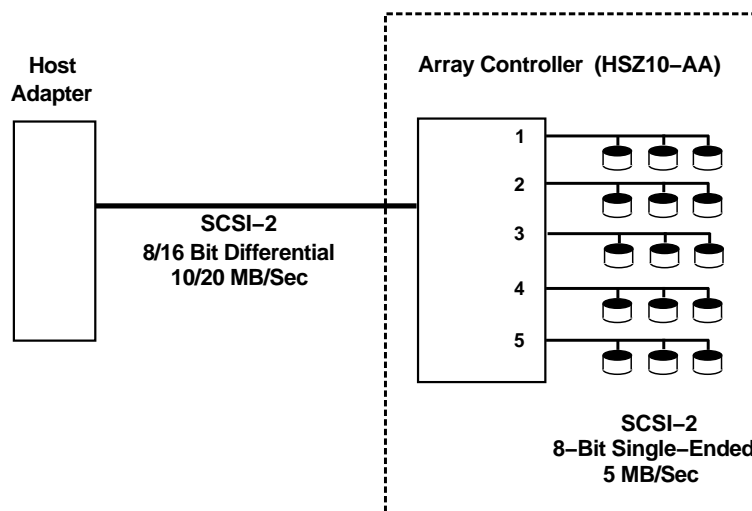
3.1 What Is a Disk Array?

This section gives a brief definition of disk arrays and a description of the four RAID levels currently supported by the DEC RAID 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 3-1 shows a conceptual model of the DEC RAID Subsystem.

Figure 3-1 Conceptual Diagram of the DEC RAID Subsystem



SHR-XR3004-GRA

RAID Overview

3.1 What Is a Disk Array?

Data for a given file is segmented by dividing the data into segments and storing it to the different drives in the array, rather than the data being written to a single drive. A **segment** is a group of blocks that is continuous data which 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.

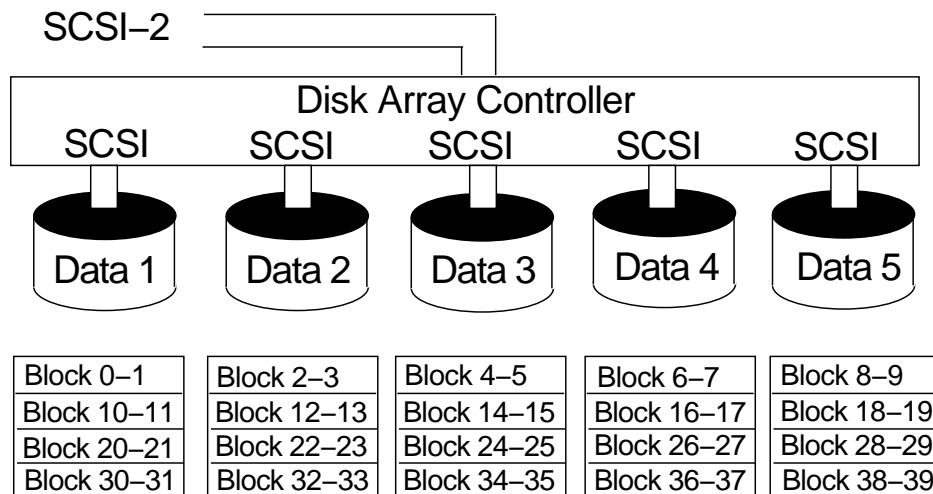
Arrays can also provide data redundancy, so that no data is lost if a single drive in the array fails. Depending on the RAID level, data is either **mirrored** or **striped**. The RAID levels currently supported are RAID 0, 1, 3, and 5. RAID 1, 3, and 5 offer data redundancy.

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 (LUNs)** – see the following sections for a description of arrays and logical units). Each logical unit has its own characteristics (RAID level, logical block size, logical unit size, and so on).

3.1.1 Description of RAID 0

RAID 0 stores data across the drives in the array one segment at a time, as shown in Figure 3–2.

Figure 3–2 Diagram of RAID 0



SHR–XR3005–GRA

In Figure 3–2, 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.)

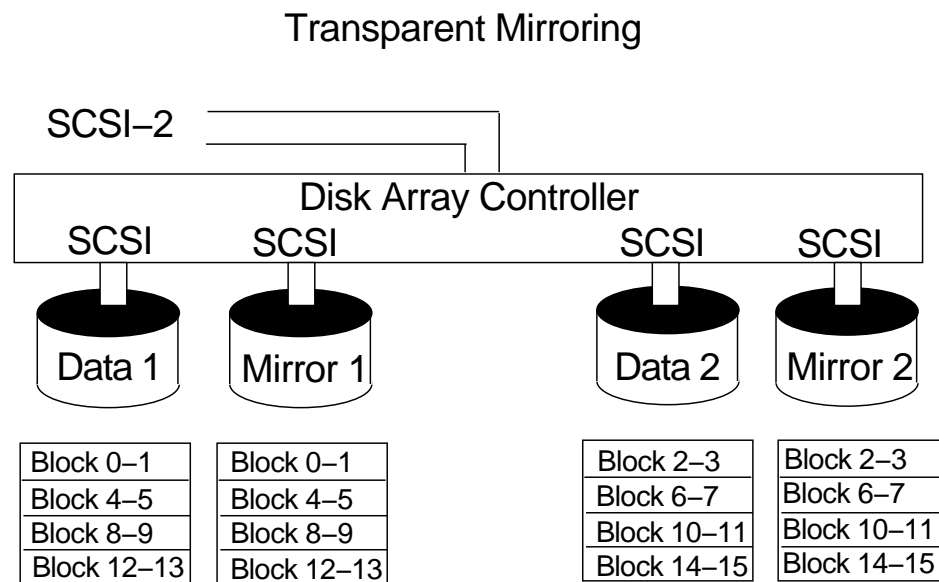
The host system treats a RAID 0 array like a standard hard drive. RAID 0 errors are reported in the same way as normal drive errors, and the recovery procedures are the same as those used on a standard hard drive. For example, if a drive fails, the array controller returns the same errors that occur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

3.1.2 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 3–3.

Figure 3–3 Diagram of RAID 1



SHR–XR3006–GRA

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 3–3, 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.

RAID Overview

3.1 What Is a Disk Array?

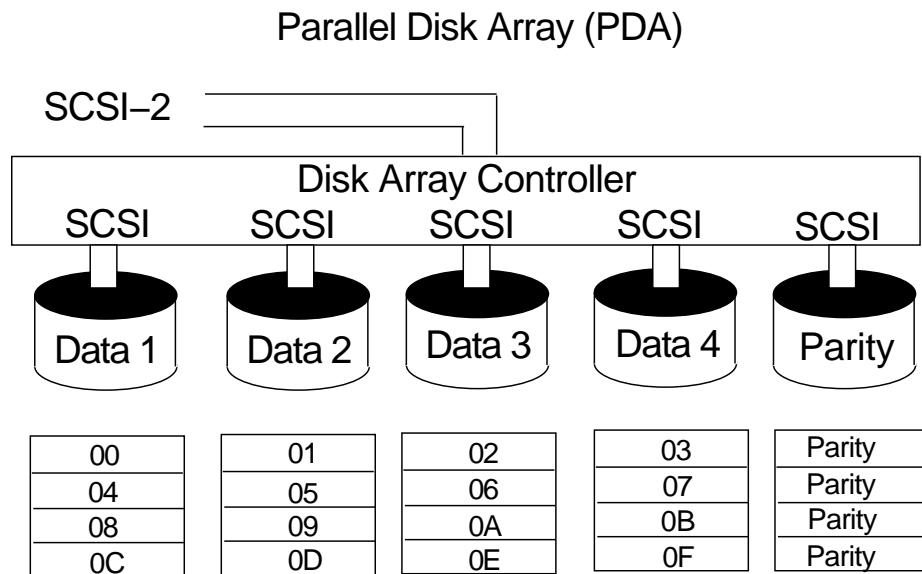
A RAID 1 array is unique in that you can have more than one failed drive and still use the data as long as there is only one failed drive in each mirrored pair. For example, if you have a six-drive RAID 1 logical unit, you have three pairs. If there is a drive failure in each pair, you can still access your data. It is still recommended that you replace the failed drives as soon as possible because if a second drive fails in a pair, all data on the logical unit is lost.

3.1.3 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 3–4, 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.

Figure 3–4 Diagram of RAID 3



SHR–XR3007–GRA

Note

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

block size of a three-drive logical unit is 1024 bytes, and 2048 bytes for a five-drive logical unit.

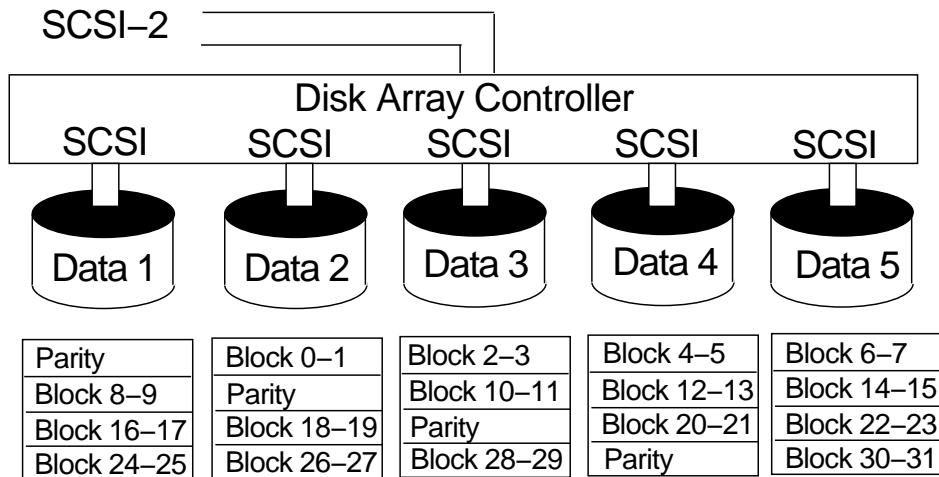
If a drive fails in RAID 3 array, you can continue to use the array normally since the array controller recalculates the data and parity blocks from the data and parity on the operational drives.

For example, if drive 1 in Figure 3–4 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.

3.1.4 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 3–5.

Figure 3–5 Diagram of RAID 5



SHR–XR3008–GRA

In Figure 3–5, 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 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 3–5 (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.

RAID Overview

3.2 Key Concepts

3.2 Key Concepts

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

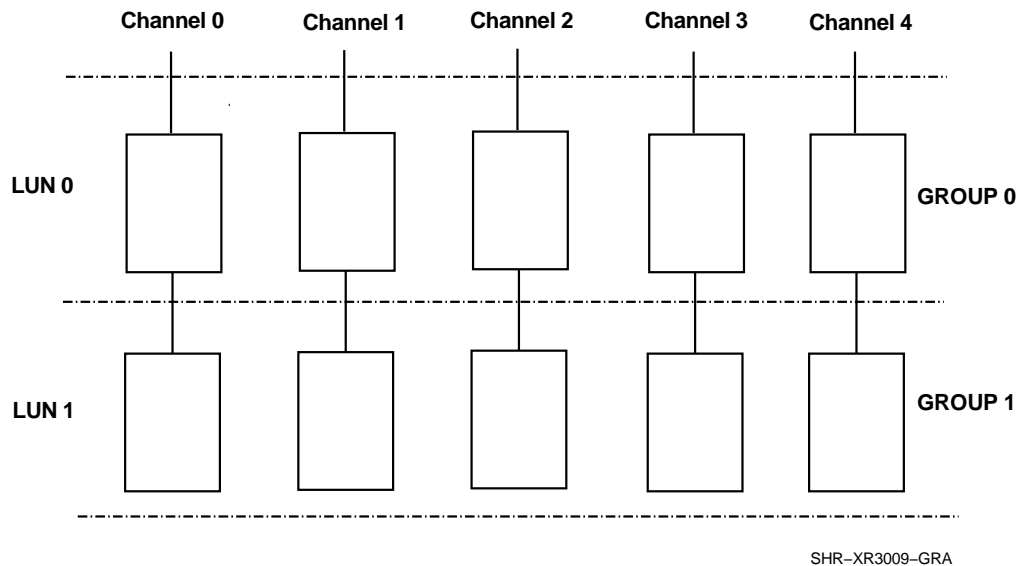
3.2.1 Array Channels

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

3.2.2 Logical Units (LUNs)

A logical unit (LUN) is a grouping of drives. A logical unit has its own device SCSI ID and LUN 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. See Figure 3–6.

Figure 3–6 Diagram of a Regular LUN



A logical unit by most operating systems is treated as a single disk drive, 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 logical unit can be in various modes of operation, depending upon the RAID level configuration.

Logical Unit modes of operation for nonredundant RAID level 0:

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

Logical Unit modes of operation for redundant RAID levels 1, 3, and 5:

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

3.2.3 Drive Groups

A **drive group** is a set of from 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 drive type. Operating systems treat the logical unit (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 logical unit in a drive group, so the use of the terms *logical unit* and *drive group* are synonymous.

3.2.4 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.

3.2.5 Partitions

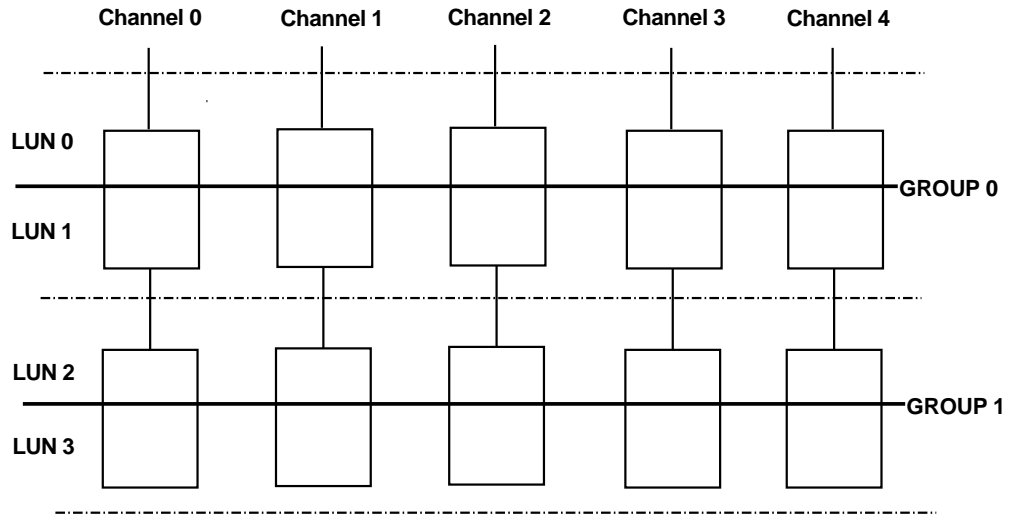
A **partition** is a portion of each drive in a rank. The partition allows multiple logical units on one rank of drives. This is useful if the host has a capacity limit per logical unit.

Figure 3–7 is a conceptual diagram of the drive partition.

RAID Overview

3.2 Key Concepts

Figure 3–7 Drive Partition Diagram



SHR–XR3010–GRA

3.2.6 Reconstruction

Reconstruction is done after a failed drive has been replaced in a RAID level 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.

3.2.7 Regeneration

Regeneration occurs when data cannot be accessed from a drive in a RAID level 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.

SZ200 Base Configuration

The SZ200 subsystem is preconfigured at the factory and ready for use. Preconfigured means that the disk array software has been installed, and default RAID level and settings have been configured. The five drives have been preconfigured to look as one logical unit.

The basic SZ200 configuration is as follows:

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

If this configuration does not meet your needs, it can be modified using the DEC RAID utilities. Refer to appropriate the *DEC RAID Utilities User's Guide* for your operating system. If you need to expand your configuration to include additional storage, refer to Chapter 7.


4.1 SCSI IDs for the DEC RAID Subsystem

Each device connected to the DEC RAID Subsystem is a 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. Small computer interface (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 DEC RAID Subsystem, the SCSI IDs are assigned on several different levels as shown in Figure 4-1.

SZ200 Base Configuration

4.1 SCSI IDs for the DEC RAID Subsystem

Figure 4–1 DEC RAID 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 4–1, 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 4–2. 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 4–2.

4.2 HSZ10-AA Controller Location and SCSI Address

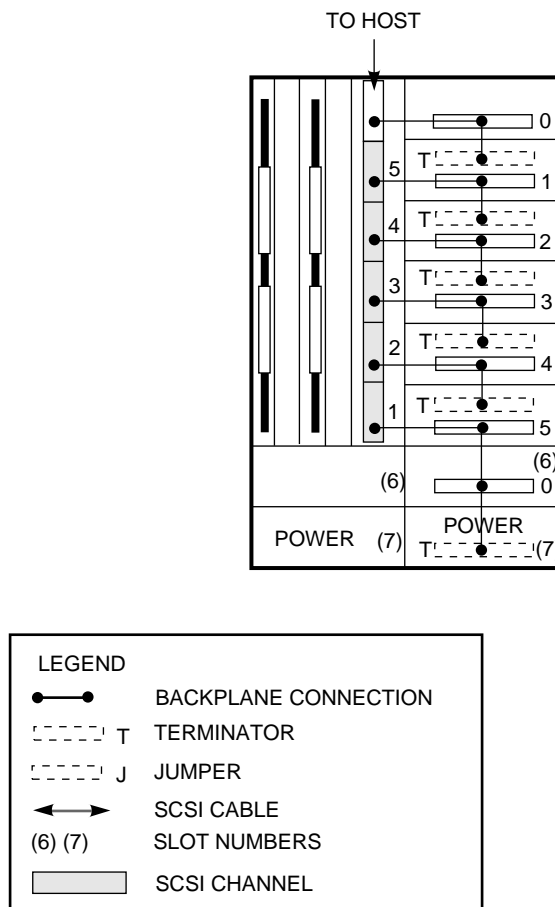
The HSZ10-AA controller must be mounted in the slot marked “Controller 1” in Figure 4–2. The position of this controller is important since its SCSI ID on the array channels is set by the backplane. In the slot position, Controller 1, the controller is set to SCSI ID 7.

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 4–2 Physical Drive Map



CXO-3589B-MC

SZ200 Base Configuration

4.3 Physical Drive Map

4.3 Physical Drive Map

The drive map defines the drives included in the logical unit. 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 SZ200 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.

4.4 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 logical unit and includes the following five drives:

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

4.5 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 an 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.

Basic Configuration Installation

This chapter discusses installation topics that are specific to the DEC RAID 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.

For a detailed discussion of the various cable options, refer to Chapter 3 of the *StorageWorks Subsystem Shelf Building Block Configuration Guide*.

This chapter discusses the following:

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

5.1 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/adapter used. Refer to the sections below for a description of connecting to 16-bit differential, 8-bit differential, and 8-bit single-ended hosts.

5.1.1 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.

Basic Configuration Installation

5.1 Connecting to a Host

5.1.2 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 DEC RAID Subsystem is at the end of the bus, a H879-AA bus terminator is needed. If the DEC RAID Subsystem is in the middle of the bus, use another SCSI-3 “P” cable to daisy-chain to another device.

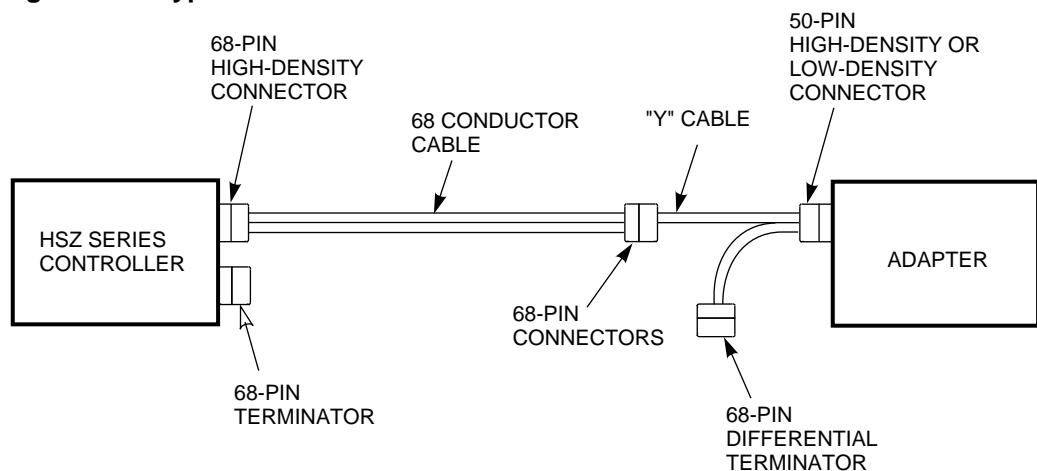
5.1.3 Connecting to an 8-Bit Differential Host/Adapter

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

In all cases, the SCSI-3 “P” cable will be 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 will be 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 5-1.

Figure 5-1 Typical Y Cable Connection



CXO-3587B-MC


5.1.4 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 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 5-2.

Figure 5-2 BA350-EA and BA350-SA Combination Storage Array



SHR-XR3029-GRA

Basic Configuration Installation

5.1 Connecting to a Host

5.1.5 Maintaining Bus Continuity

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

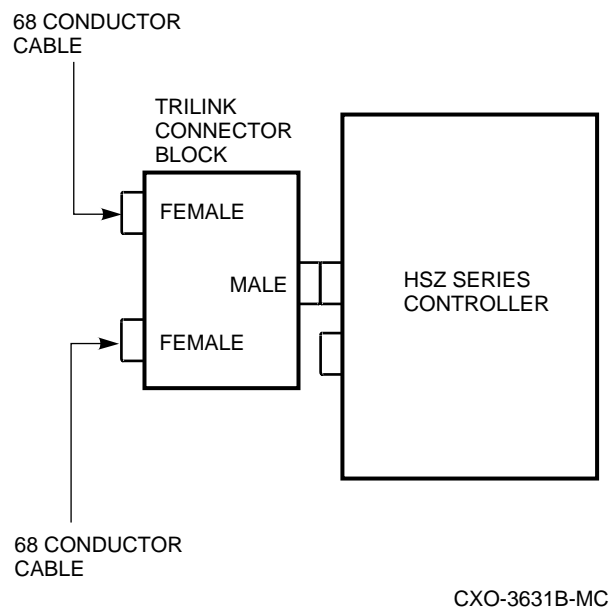
The tralink connector block (H885-AA) shown in Figure 5-3 is a small, metal block with two 68-pin, high-density, female connectors with 2-56 jack standoffs on one side and a single 68-pin, high-density, male connector with 2-56 jackscrews on the other side. Using the tralink connector, the HSZ10-AA controller can occupy either an end-bus or a midbus position, yet remain removable for service.

For midbus connections, follow the steps listed in Table 5-1 and refer to Figure 5-3.

Table 5-1 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 tralink connector block.
3.	Attach two BN21L cables to the other side of the tralink connector block.

Figure 5-3 Tralink Connector Midbus Connection

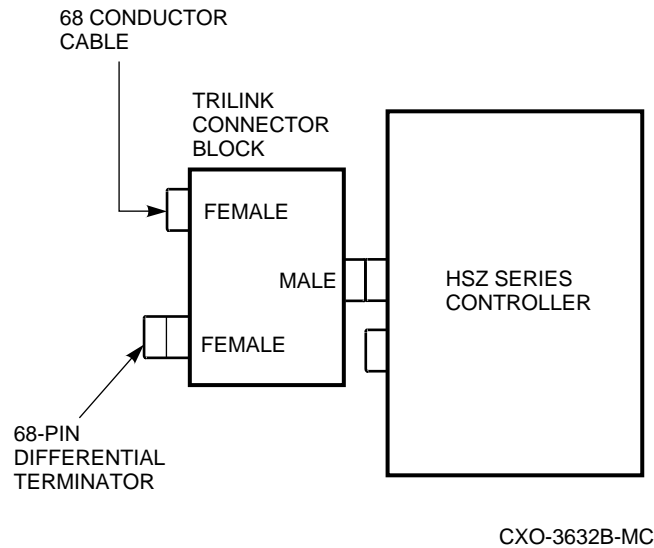


For end-bus connections, follow the steps listed in Table 5-2 and refer to Figure 5-4.

Table 5–2 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.

Figure 5–4 Trilink Connector End-Bus Connection



5.2 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.

Basic Configuration Installation

5.3 Powering On the Subsystem

5.3 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.

5.4 Functional Verification

After powering on the subsystem, you should 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*. See Section 6.2.2.1 for a description of the HSZ10-AA controller start-up procedures.

This chapter discusses the following:

- DEC RAID Subsystem monitoring features
- User monitoring methods
- When and how to replace a drive
- When and how to replace power supplies and blowers
- Reconstruction
- Parity/Check Repair

6.1 DEC RAID Subsystem Monitoring Features

The SZ200, 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 DEC RAID Subsystem provides self-monitoring through two components:

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

6.1.1 Monitoring Through the HSZ10-AA Controller

The HSZ10-AA controller monitors drive operation and logical unit status. If certain errors occur on a drive, the controller changes the status of a drive or logical unit. A change in drive or logical unit 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 Section 6.1.1.1; drive status is discussed in more detail in Section 6.1.1.2.

In addition, the HSZ10-AA controller is responsible for setting the fault indicators on the disk drive SBBs. This is done when the controller determines a drive has failed or the user has actively failed a drive through the DEC RAID utilities.

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


6.1.1.1 LUN Status

The logical unit or grouping of the drives can be in different states. There are four possibilities for status, as shown in Figure 6–1.

Operations

6.1 DEC RAID Subsystem Monitoring Features

Figure 6–1 LUN State Diagram



SHR–XR3015–GRA

Table 6–1 describes each logical unit status.

Table 6–1 Logical Unit Status

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

As detailed in Table 6–1, an *optimal* status is the desired condition for normal operation.

If a drive failure does occur in a redundant configuration (RAID 1, 3, or 5) the DEC RAID 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 case, the array controller retrieves the failed drive’s data from its mirrored drive whenever you read or write to a logical unit in the degraded mode. In the RAID 3 and 5 case, 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 data 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.

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.

6.1 DEC RAID Subsystem Monitoring Features

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.

6.1.1.2 Drive Status

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

Table 6–2 lists the status classifications for drives.

Table 6–2 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 logical unit. 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 DEC RAID utility. The drive must be replaced.
Spare	The drive is connected to an array but not configured into a logical unit.
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 logical unit.

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 as follows:

- 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 will be 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 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.

Operations

6.1 DEC RAID Subsystem Monitoring Features

- Unrecoverable errors during the read portion of a read/modify/write operation (RAID 5).
- The user marked a drive as failed using a SCSI command or DEC RAID utility.
- An error was reported during disk array format sequence.

6.1.2 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 6.4 for instructions in determining a power supply or blower fault condition.

6.2 User Monitoring Methods

There are two methods for monitoring the operation of the DEC RAID Subsystem:

- DEC RAID utilities
- LED indicators

6.2.1 Monitoring Operation Using the DEC RAID Utilities

The DEC RAID utilities are key tools to maintaining a disk array. They provide the user with the ability to monitor, configure, repair, and maintain the DEC RAID 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 logical unit (RAID levels 1, 3, and 5) after a drive failure.
- Check and repair array parity on logical units.
- Configure, reconfigure, and modify logical units.
- Change default array parameters.
- Change DEC RAID 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 DEC RAID 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* for the appropriate operating system.

6.2.2 Monitoring Operation Using LED Indicators

The DEC RAID 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 as follows:

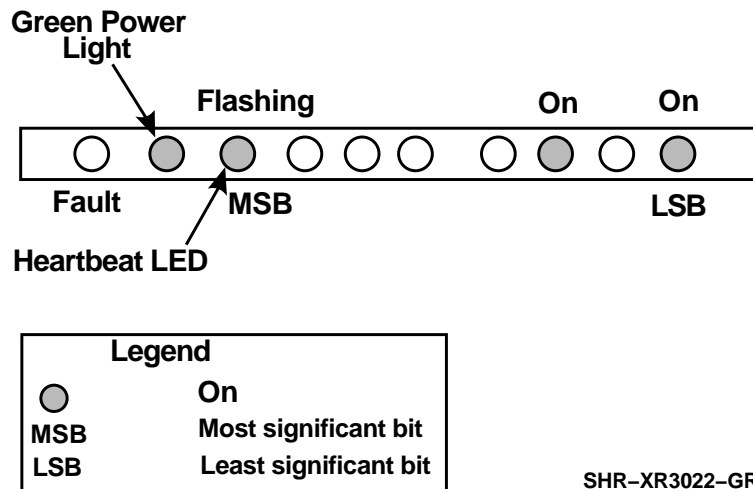
- HSZ10-AA controller LEDs
- Power supply SBB LEDs
- Drive SBB LEDs

6.2.2.1 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 2-2 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 6-2.

Figure 6-2 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 nothing else.

Table 6-3 shows the LEDs 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.

Operations

6.2 User Monitoring Methods

Table 6–3 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 turnaround diagnostics in progress.
88 XX	Cyclic	Error	SCSI channel data turnaround diagnostics failed.

(continued on next page)


Table 6–3 (Cont.) Summary of HSZ10-AA Controller LED Codes

Code MSB/LSB	S/F/C	State	Description
89	Solid	Normal	Subsystem turnaround diagnostics in progress.
89	Cyclic	Error	Subsystem turnaround 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.

6.2.2.2 Power Supply LEDs

Shelf and power supply status are displayed on the power supply LEDs shown in Figure 6–3. The upper LED displays the status of the shelf and the lower LED displays the status of the supply.

Figure 6–3 Power Supply Status LEDs



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, one of these fault conditions exists.

Operations

6.2 User Monitoring Methods

For a detailed explanation of the power supply LED codes, refer to Tables 6–4 and 6–5.

Table 6–4 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 6.4.

Table 6–5 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.

6.2.2.3 Drive SBB LEDs

Each drive SBB has two LED indicators, as shown in Figure 6–4, which display the SBB's status. These LEDs have three states as follows:

- On
- Off
- Flashing

Figure 6–4 Shelf Status LEDs



CXO-3614A-PH

- 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

On initial power-on of the DEC RAID 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 6–6 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 6.3 for instructions in replacing a drive SBB.

Operations

6.3 When and How to Replace a Drive

6.3 When and How to Replace a Drive

You need to replace a drive when the following occurs:


- DEC RAID utilities provide information stating that a drive is in a *warning* state.
- DEC RAID utilities provide information stating that a drive is in a *failed* state.
- 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.

The following are steps you should take to replace a drive:

1. If the drive is not yet in a failed state, fail the drive using the DEC RAID utilities.
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 6–5.

Figure 6–5 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 6.5 for more detailed information.

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.

6.4 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 Section 6.2.2.2.

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, you simply disconnect the power cable from that power supply.

There are three swapping methods for replacing power supplies as follows:

- **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.

6.4.1 Replacing a Shelf Power Supply

Warning

Always use both hands when removing or replacing an SBB, to fully support its weight.

To replace either a shelf primary or redundant power supply, complete the following procedure:

Operations

6.4 When and How to Replace Power Supplies and Blowers

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 6-5.
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 6-4 or to Table 6-5).
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 device for normal operation indications.

6.4.2 Replacing a Blower

Each shelf has two blowers mounted on the rear. Connectors on the backplane provide the +12 Vdc 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.

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.

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 6-7.

Then replace the shelf in the cabinet and apply power.


To replace a blower, follow the steps in Table 6-7:

6.4 When and How to Replace Power Supplies and Blowers

Table 6–7 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 6–6, 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 module, 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 6–6 Replacing BA35X–MA Blowers



CXO-3515A-PH

6.5 Reconstruction

By default, the HSZ10-AA controller automatically initiates the reconstruction process after you replace a drive in a degraded RAID 1 or 5 logical unit. Through the DEC RAID utilities, you may 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 logical unit 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 logical unit. The controller then writes this data to the replaced drive. Note that although RAID level 1 does not have parity, the array controller can reconstruct data on a RAID 1 logical unit by copying data from the mirrored disk.

Operations

6.5 Reconstruction

Note

Reconstruction applies only to a degraded RAID 1 or 5 logical unit with a single-drive failure.

Once reconstruction is initiated (either by you or by the array controller), the HSZ10-AA 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 logical unit drives.
- Writes the recalculated data and parity to the new drive.

Once reconstruction is started, it can take place while the logical unit is in use. You do not need to shut down the DEC RAID Subsystem. You can see the access to all the drives in the logical unit 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.

6.6 Parity Check/Repair

It is important for the user to perform periodic parity checks on the redundant arrays, especially if there is a power failure or an abnormal shutdown.

Parity Check/Repair is the process of verifying and repairing parity information so that data can be maintained and reconstructed if there is a drive failure. Parity Check/Repair functionality is provided by the DEC RAID utilities.

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 actually have *parity* either, but parity check compares data on the mirrored drives. Parity Check/Repair cannot be performed on a degraded logical unit.

Parity Check/Repair performs the following functions:

- Scans the logical unit and checks the array parity for each block in the logical unit. On a RAID 1 logical unit, Parity Check compares the data on each mirrored pair of disks, block by block.
- Repairs any array parity errors found during the parity check.

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). Data cached in buffers will be lost and cannot be reconstructed if you do not have a UPS. This is one of the reasons you should always maintain backup files, even with a redundant array.

DEC RAID Parity Check/Repair utility can be run at the following times:

- Automatic parity check/repair will guarantee the data integrity of the logical unit so that you can reconstruct the data on the array if a drive fails. This procedure is either set at a time determined during the installation of the DEC RAID Utility, or it can be set up by means of a script that runs at a set time.
- Manually after an abnormal server or array shutdown. As the result of such a shutdown, required array parity may not have been updated, resulting in potential data corruption.

For more detailed information, refer to the *DEC RAID Utilities User's Guide* for the appropriate operating system

6.7 Upgrading Software

The disk array software can be upgraded using a DEC RAID 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.

Advanced Configurations

To do advanced configurations, you need to use one of the RAID Manager utilities. The following sections outline the configuration limits and steps needed for reconfigurations. All information is given in a generic sense, and 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

7.1 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 Logical Unit 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 logical unit = 1 to 5.
1	Number of drives allowed per logical unit = 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 mirrored pair cannot be on the same channel.
3	Supports 3 or 5 drives.
5	Number of drives allowed per logical unit = 3 to 5. Each drive must be on separate channel.
ALL	Drives within a RAID level must be of same vendor type and capacity.

Advanced Configurations

7.1 Modifying Basic Configurations

Note

It is recommended that you use the maximum number of drives when you create a logical unit.

7.2 Multiple Rank Configurations

The DEC RAID 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 level 5 configurations and the other two ranks are RAID level 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 7.3 to understand the RAID configuration and expansion path restrictions.

7.3 Configuration Guidelines

A key feature of the DEC RAID 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 logical unit = 1 to 10.
1	Number of drives allowed per logical unit = 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 mirrored pair cannot be on the same channel.
5	Number of drives allowed per logical unit = 3 to 5. Each drive must be on separate channel.
ALL	Drives within a RAID level must be of same vendor type and capacity.

Recommendations

It is recommended that you configure a logical unit 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

Advanced Configurations 7.3 Configuration Guidelines

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 level 5 and 0 within logical unit 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

7.3.1 Expansion Guidelines

There are three recommended packaging expansion paths as follows:

- Single BA350-EA to Single BA350-EA/Dual BA350-SA
Allows expansion from 1 rank to a maximum of 3 ranks of drives.
- Single BA350-EA to Single BA350-EA/Quad BA350-SA
Allows expansion from 1 rank to a maximum of 6 ranks of drives.
- Single BA350-EA/Dual BA350-SA to Single BA350-EA/Quad BA350-SA
Allows expansion from 4 ranks to a maximum of 6 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

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.

Advanced Configurations

7.3 Configuration Guidelines

Refer to Table 7–1 for a listing of the possible expansion paths available to you.

Table 7–1 Expansion Paths

Present	Future	BA350-EA	BA350-SA	Custom Reconfiguration Needed	Refer to:
1 Rank	2 Ranks	1	1	Yes	Section 7.5
1 Rank	2 Ranks	1	2	No	Section 7.4.1
1 Rank	3 Ranks	1	2	No	Section 7.4.1
1 Rank	4 Ranks	1	4	No	Section 7.4.2
1 Rank	5 Ranks	1	4	No	Section 7.4.2
1 Rank	6 Ranks	1	4	No	Section 7.4.2
2 Ranks	3 Ranks	1	3	Yes	Section 7.5
2 Ranks	4 Ranks	1	4	No	Section 7.4.3
2 Ranks	5 Ranks	1	4	No	Section 7.4.3
2 Ranks	6 Ranks	1	4	No	Section 7.4.3
3 Ranks	4 Ranks	1	4	No	Section 7.4.3
3 Ranks	5 Ranks	1	4	No	Section 7.4.3
3 Ranks	6 Ranks	1	4	No	Section 7.4.3
4 Ranks	5 Ranks	1	4	No	Section 7.4.3
4 Ranks	6 Ranks	1	4	No	Section 7.4.3
5 Ranks	6 Ranks	1	4	No	Section 7.4.3

7.4 Recommended Configurations

The following sections describe the recommended expansion paths. Each section describes the following:

- Shelf requirements
- Purpose of the expansion
- Pre-reconfiguration tasks
- Reconfiguration tasks
- Post-reconfiguration tasks

7.4.1 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 (5 drives, 1 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:

- Pre-reconfiguration
- Shelf reconfiguration
- Post-reconfiguration

7.4.1.1 Pre-Reconfiguration

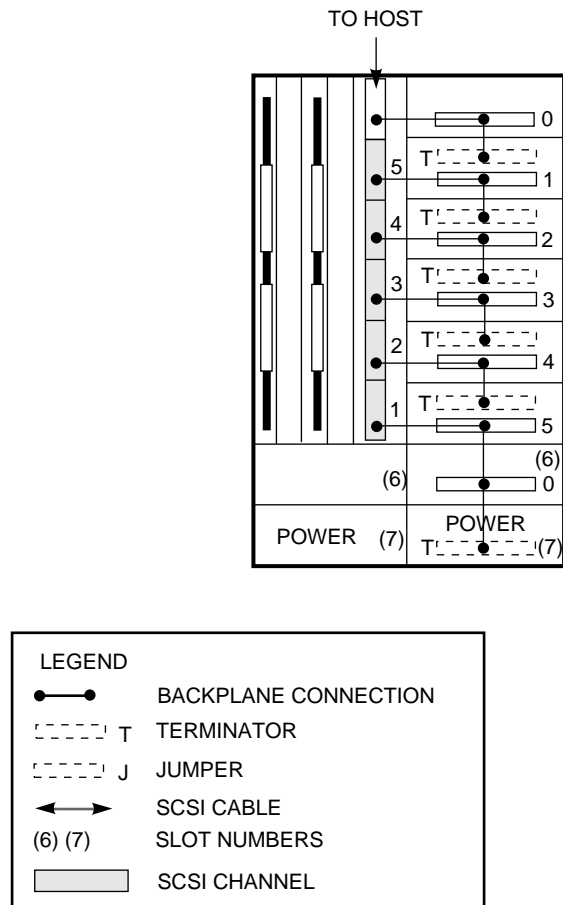
Before you begin to expand your configuration, perform the following steps:

1. Back up your data to other media.
2. Figure 7–1 shows your current configuration. Use Figure 7–1 as the guide for identification information. Label each of your current drives with array channel and SCSI ID information.

Advanced Configurations

7.4 Recommended Configurations

Figure 7–1 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 7–2 shows the physical layout of the reconfiguration. Figure 7–3 represents the logical view of the reconfiguration. Study Figure 7–2 and Figure 7–3 to be certain that you understand the cabling and relocation of drives. Refer to Chapter 5 to review connection information.

- Power down the DEC RAID Subsystem.

You are ready to begin the shelf reconfiguration described in Section 7.4.1.2.


7.4.1.2 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 7–2.

Advanced Configurations 7.4 Recommended Configurations

Figure 7–2 Single BA350-EA to Single BA350-EA/Dual BA350-SA Shelf Reconfiguration



CXO-3650B-MC

To begin, the BA350-EA shelf is reconfigured by removing some of the terminators and adding jumpers instead. 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.

Advanced Configurations

7.4 Recommended Configurations

Refer to Figure 7–3. 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 7–2. As you continue with the reconfiguration, be certain to follow the pattern of jumpers and terminators shown in Figure 7–2.
2. Using Figure 7–3 as a guide, perform the following steps:

Figure 7–3 Single BA350-EA to Single BA350-EA/Dual BA350-SA Drive Reconfiguration



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 7–4. The top connector is used exclusively for connecting to host adapter SSBs (for example, the DWZZA-VA) in slot 0 and that connector is unnumbered in Figure 7–4. The other connectors are labeled as 5, 4, 3, 2, and 1.


Advanced Configurations

7.4 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 7-4.

Figure 7-4 BA350-EA and BA350-SA Combination Storage Array



SHR-XR3029-GRA

- 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.

Advanced Configurations

7.4 Recommended Configurations

- 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 7-2.

- c. At this point, the original five drives from the BA350-EA shelf must be inserted in locations with the same channel and SCSI ID as originally assigned to the drive. This will allow 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) must be relocated to the BA350-SA shelves. Refer to Figure 7-3.

- 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 7-3.

- 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 7-3.

- 2. Power on the DEC RAID 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 7.4.1.3.

7.4.1.3 Post-Reconfiguration

To verify your configuration, use DEC RAID Utilities appropriate for your operating system to perform the following steps:

- 1. Verify that original drives are all set to optimal condition.
- 2. Verify that new drives are displayed as “spare” drives in the proper array channel/SCSI ID positions.
- 3. Configure and format desired RAID level configurations.
- 4. Reboot the system.

7.4.2 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 (5 drives, 1 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:

- Pre-reconfiguration
- Shelf reconfiguration
- Post-reconfiguration

7.4.2.1 Pre-Reconfiguration


Before you begin to expand your configuration, perform the following steps:

1. Back up your data to other media.
2. Figure 7–5 shows your current configuration. Use Figure 7–5 as the guide for identification information. Label each of your current drives with array channel and SCSI ID information.

Advanced Configurations

7.4 Recommended Configurations

Figure 7-5 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.

- Figure 7-6 shows the physical layout of the reconfiguration. Figure 7-7 represents the logical view of the reconfiguration. Study Figure 7-6 and Figure 7-7 to be certain that you understand the cabling and relocation of drives. Refer to Chapter 5 to review connection information.

- Power down the DEC RAID Subsystem.

You are ready to begin the shelf reconfiguration described in Section 7.4.2.2.


7.4.2.2 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 7-6.

Advanced Configurations 7.4 Recommended Configurations

Figure 7-6 Single BA350-EA to Single BA350-EA/Quad BA350-SA Shelf Reconfiguration



To begin, the BA350-EA shelf is reconfigured by removing some of the terminators and adding jumpers instead. 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 7-7. 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).

Advanced Configurations

7.4 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 7-6. As you continue with the reconfiguration, be certain to follow the pattern of jumpers and terminators shown in Figure 7-6.
2. Using Figure 7-7 as a guide, perform the following steps:

Figure 7-7 Single BA350-EA to Single BA350-EA/Quad BA350-SA Drive Reconfiguration



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

- 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 7-5. The top connector is used exclusively for connecting to host adapter SSBs (for example, the DWZZA-VA) in slot 0, and that connector is unnumbered in Figure 7-5. The other connectors are labeled as 5, 4, 3, 2, and 1.

Advanced Configurations

7.4 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, 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 7–6.

- c. At this point, the original five drives from the BA350-EA shelf must be inserted in locations with the same channel and SCSI ID as originally assigned to the drive. This will allow 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) must be relocated each to a BA350-SA shelf. Refer to Figure 7–7.

- 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 7–7.
- 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 7–7.
- 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 7–7.
- 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 7–7.

Advanced Configurations

7.4 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 7-7.

3. Power on the DEC RAID 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 7.4.2.3.

7.4.2.3 Post-Reconfiguration

To verify your configuration, use DEC RAID Utilities appropriate for your operating system to perform the following steps:

1. Verify that original drives are all set to optimal condition.
2. Verify that new drives are displayed as “spare” drives in the proper array channel/SCSI ID positions.
3. Configure and format desired RAID level configurations.
4. Reboot the system.

7.4.3 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:

- Pre-reconfiguration
- Shelf reconfiguration
- Post-reconfiguration

7.4.3.1 Pre-Reconfiguration


Before you begin to expand your configuration, perform the following steps:

1. Back up your data to other media.
2. Figure 7–8 shows your current configuration. Use Figure 7–8 as the guide for identification information. Label each of your current drives with array channel and SCSI ID information.

Advanced Configurations

7.4 Recommended Configurations

Figure 7–8 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 7–9 shows the physical layout of the reconfiguration. Figure 7–10 represents the logical view of the reconfiguration. Study Figure 7–9 and Figure 7–10 to be certain that you understand the cabling and relocation of drives. Refer to Chapter 5 to review connection information.

- Power down the DEC RAID Subsystem.

You are ready to begin the shelf reconfiguration described in Section 7.4.3.2.


7.4.3.2 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 7–9.

Advanced Configurations 7.4 Recommended Configurations

Figure 7–9 Single BA350-EA/Dual BA350-SA Shelf to Single BA350-EA/Quad BA350-SA Shelf Reconfiguration



To begin, the BA350-EA shelf is reconfigured by removing some of the terminators and adding jumpers instead. 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 7–10. 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).

Advanced Configurations

7.4 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 7–9. As you continue with the reconfiguration, be certain to follow the pattern of jumpers and terminators shown in Figure 7–9.
2. Using Figure 7–10 as a guide, perform the following steps:

Figure 7–10 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 7–8. The top connector is used exclusively for connecting to host adapter SSBs (for example, the DWZZA-VA) in slot 0, and that connector is unnumbered in Figure 7–8. The other connectors are labeled as 5, 4, 3, 2, and 1.

Advanced Configurations

7.4 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, 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 7–9.

- c. At this point, the original drives in the BA350-EA and BA350-SA shelves must be inserted in locations with the same channel and SCSI ID as originally assigned to the drive. This will allow 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) must be 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 7–10.

- 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 7–10.

- 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 7–10.

Advanced Configurations

7.4 Recommended Configurations

3. Power on the DEC RAID 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 7.4.3.3.

7.4.3.3 Post-Reconfiguration

To verify your configuration, use DEC RAID Utilities appropriate for your operating system to perform the following steps:

1. Verify that original drives are all set to optimal condition.
2. Verify that new drives are displayed as “spare” drives in the proper array channel/SCSI ID positions.
3. Configure and format desired RAID level configurations.
4. Reboot the system.

7.5 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 7–1 for details on which expansion paths this includes.

To perform custom expansions, do the following:

Note

It is very important that you back up your data to other media, since this procedure will destroy your data.

1. Back up your data to other media.
2. Using DEC RAID utilities, perform the following steps:
 - a. Delete all configured LUNs. This makes all drives spares.
 - b. Reconfigure drive in array channel 5, SCSI ID 1 to be a 1 drive RAID 0 array.
3. Power down your DEC RAID Subsystem.
4. Refer to Table 7–1 for the appropriate section for reconfiguring your shelves.
5. Power on your DEC RAID Subsystem.
6. Using DEC RAID 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 7.4.1 through 7.4.3. 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 7.4.1 through 7.4.3 because you are using this customized expansion technique.

Error Handling/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 DEC RAID Subsystem

8.1 Before You Begin Troubleshooting

To determine where the problem with your DEC RAID Subsystem exists, follow these steps:

1. Turn off the DEC RAID 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 HSZ10-AA controller
 - Terminator
 - SCSI cable connected to host system
 - 50-pin SCSI cables to BA350-SA shelves (if applicable)
4. Turn the DEC RAID Subsystem back on.
5. Verify that all the drive SBBs 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 LEDs indicators flash and the controller begins its diagnostics.

8.2 Using the Troubleshooting Table

When the DEC RAID 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 8–1, follow these steps:

1. Note the symptoms of the problem with your DEC RAID Subsystem.
2. Check the *Symptom* column in Table 8–1 for a match.

Error Handling/Troubleshooting

8.2 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 8–1 Troubleshooting Systems Problems

Symptom	Possible Cause	Corrective Action
Drive SBB Fault Light is on.	Drive has failed.	Replace drive using instructions in Chapter 6, Section 6.3.
Drive SBB Fault and Activity Lights are on.	Drive has failed or is hung.	Replace drive using instructions in Chapter 6, Section 6.3.
Drive SBB Fault Light is flashing.	Drive has been failed and is spinning down.	Replace drive using instructions in Chapter 6, Section 6.3.
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 6, Section 6.2 and Section 6.4, for description of fault condition and resolution.
Both lights off on Power Supply.	Input power problem.	Check for proper connection of input power.
	Shelf and Power Supply Fault.	Refer to Chapter 6, Section 6.2 and Section 6.4, for description of fault condition and resolution.
HSZ10-AA controller LED indicators are off.	Controller is not properly installed.	Remove and reseat 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> .
DEC RAID Subsystem is not seen by host.	SCSI cable is not connected.	Check the SCSI cable both at the host and DEC RAID Subsystem ends for proper connection.
	Incorrect termination.	Check that both ends of SCSI bus are terminated correctly. Verify that proper termination scheme is being used (see Chapter 5).

(continued on next page)

Table 8–1 (Cont.) Troubleshooting Systems Problems

Symptom	Possible Cause	Corrective Action
	Duplicate SCSI IDs on bus.	Check SCSI ID settings on all devices connected to bus for duplication. Run bus scan console diagnostics if they are available on your host system.
	Defective HSZ10-AA controller.	Refer to Chapter 6, Section 6.2, for a description of the possible LED status. Determine whether controller has failed. Replace the controller if needed.
Software does not boot from the DEC RAID Subsystem.	A problem exists with the DEC RAID Subsystem.	Use any available system console diagnostics or the DEC RAID utilities to test DEC RAID Subsystem.
	A problem exists with the software installed on the DEC RAID Subsystem.	Refer to the <i>DEC RAID Utilities User's Guide</i> for help.

8.3 If You Have Expanded Your DEC RAID Subsystem

If you have attempted an expansion of your DEC RAID Subsystem, follow these steps to determine the problem:

1. Turn off the DEC RAID 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 7.
 - 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 HSZ10-AA controller
 - Terminator
 - SCSI cable connected to host system
 - 50-pin SCSI cables to BA350-SA shelves (if applicable)
5. Check that drive movement and drive addition have been done correctly. Refer to Chapter 7.
6. Turn the DEC RAID Subsystem back on.
7. Verify that all the drive SBBs LED indicator flash on initial power on.
8. Verify that the two LED indicators on the power supplies are lit.

Error Handling/Troubleshooting

8.3 If You Have Expanded Your DEC RAID Subsystem

9. Verify that the HSZ10-AA controller LEDs indicators flash and the controller begins its diagnostics.
10. Use the DEC RAID utilities to determine the status of the LUN and drive.

Physical Specifications

This appendix describes the physical, environmental, and performance specifications for the DEC RAID Subsystem product. For additional StorageWorks component specifications, refer to the *StorageWorks Subsystem Shelf Building Block Configuration Guide*.

A.1 General Specifications

Table A-1 SZ200 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 DEC RAID Subsystem is two power supplies. However, the power supply SBB is rated at 131 watts and sufficient power for a fully loaded shelf.

A.2 Power Unit Specifications

Each StorageWorks shelf requires either a primary ac or dc power unit. The power unit type is determined the 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 SBB data is maintained while the devices power down. See Table A-2 and *StorageWorks Shelf Building Block Subsystem User's Guide* for more information about the power units.

Physical Specifications

A.2 Power Unit Specifications

Table A–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 Vac	36–72 Vdc	N/A
Nominal input voltage	110 Vac	48 Vdc	12 Vdc
Autoranging feature	Yes	Yes	N/A
Output voltages	12 Vdc 5 Vdc	12 Vdc 5 Vdc	N/A
Output power†	131 W	131 W	N/A

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

A.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 A–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 A–3.

Table A–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

A.4 Environmental Specifications

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

Table A–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

B

Supported Options

This appendix includes information regarding the disk drives and adapters supported by the system. Table B-1 describes the disk drives supported by the system.

Table B-1 Disk Drives Supported

Component	Description
RZ25	426 Mbyte SCSI Disk Drive
RZ26	1.05 Gbyte SCSI Disk Drive

Table B-2 describes the adapters supported by the system.

Table B-2 Adapters Supported

Component	Description
KZESA	EISA/SCSI Differential Host Adapter
DWZZA-VA	Single-ended to Differential/Wide SCSI Adapter

C

Operating System Support

This appendix contains information regarding the operating systems supported by the DEC RAID Subsystem and part numbers associated with the DEC RAID Subsystem.

Table C-1 lists the operating systems.

Table C-1 Operating Systems Supported

Operating System	O/S Version with DEC RAID Subsystem Support
NOVELL	3.11
SCO UNIX	3.2.4
MS-DOS	3.3, 5.0
OpenVMS VAX	V5.5-2

D

System Connections

Table D-1 describes the type of connectivity for various system adapters to help you determine cabling requirements for systems. Refer to the appropriate sections in this manual based on the type of connectivity described in Table D-1.

Table D-1 is provided for informational purposes only. You should not assume that this is a complete listing of systems. Nor should you assume that the all systems listed here are fully qualified or supported. Please contact your local sales representative to verify support on a particular system or operating system platform.

Table D-1 System Connectivity

System	SCSI Adapter	Refer to:
MicroVAX 3100		
Model 30	Embedded SCSI	Section 5.1.4
Model 40	Embedded SCSI	Section 5.1.4
Model 80	Embedded SCSI	Section 5.1.4
Model 90	Embedded SCSI	Section 5.1.4
VAXstation 4000		
Model VLC	Embedded SCSI	Section 5.1.4
Model 60	Embedded SCSI	Section 5.1.4
Model 60	PMAZ	Section 5.1.4
Model 90	Embedded SCSI	Section 5.1.4
Model 90	PMAZ	Section 5.1.4
VAX System 4000		
Model 100	Embedded SCSI	Section 5.1.4
applicationDEC		
433MP	KZESA	Section 5.1.3
400XP	KZESA	Section 5.1.3
DECpc		
425ST	KZESA	Section 5.1.3

(continued on next page)

System Connections

Table D-1 (Cont.) System Connectivity

System	SCSI Adapter	Refer to:
433ST	KZESA	Section 5.1.3
450ST	KZESA	Section 5.1.3

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 inexpensive disks.

rank

The number of drives per channel. *See also* drive rank.

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 DEC RAID 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.

A

- Adapters description, B-1
- Advanced configurations, 7-1
- Array
 - channels, 3-6
 - definition, 3-1
 - features, 1-3, 2-4
- Audience, ix

B

- BA350-EA shelf description, 2-2
- BA35X-VA description, 2-1
- Basic configuration
 - modifying, 7-1
- Blower replacement, 6-12
- Bus connections
 - end-bus, 5-4
 - midbus, 5-4
- Bus continuity
 - maintaining, 5-4

C

- Cable options, 2-5
- Cables, verifying, 5-5
- Capacity, 1-5
- Cold swapping, 6-11
- Component descriptions, 2-1
- Configurations
 - advanced, 7-1
 - guidelines, 7-2
 - multiple rank, 7-2
- Connecting
 - to a 16-bit differential host/adaptor, 5-2
 - to a host, 5-1
 - to an 8-bit differential host/adaptor, 5-2
 - to an 8-bit single-ended host/adaptor, 5-3
- Connectors, verifying, 5-5
- Cost
 - of single-drive failure, 1-5
- Custom expansion, 7-22

D

- Data
 - availability, 1-4
 - mirrored, 3-2
 - reliability, 1-3
 - segmented, 3-1
- DEC RAID Subsystem
 - capacity, 1-5
 - component descriptions, 2-1
 - data availability, 1-4
 - data reliability, 1-3
 - flexibility, 1-5
 - logical view, 1-2
 - monitoring features, 6-1
 - performance, 1-4
 - physical specifications, A-1
 - product attributes, 1-3
 - product description, 1-1
 - product highlights, 1-3
 - product overview, 1-1
 - redundancy, 1-4
- DEC RAID utilities, 2-5
 - monitoring through, 6-4
- Disk array description, 3-1
- Disk array subsystems descriptions, B-1
- Documentation
 - related, x
- Drive
 - groups, 3-7
 - how to replace, 6-10
 - ranks, 3-7
 - SBB LED indicators, 6-8
 - status, 6-3
- DWZZA-VA, 4-4
 - features, 4-4

E

- Environmental specifications, A-3
- Environmental stabilization, A-2
- Error handling, 8-1
- Expansion paths
 - custom expansion, 7-22
 - from 1 BA350-EA/2 BA350-SA to 1 BA350-EA/4 BA350-SA, 7-17

Expansion paths (cont'd)

from 1 BA350-EA to 1 BA350-EA/2 BA350-SA,
7-5

from 1 BA350-EA to 1 BA350-EA/4 BA350-SA,
7-11

recommended, 7-4

table of, 7-4

types of, 7-3

Expansion unit description, 2-5

F

Features

array, 1-3

subsystem, 1-3

Firmware, 2-4

Flexibility, 1-5

H

Hot swapping, 6-11

HSZ10-AA controller

description, 2-2

features, 2-3

LED codes, 6-5

LED indicators, 6-5

location and SCSI address, 4-3

monitoring through, 6-1

I

Installation, 5-1

L

LED indicators

monitoring through, 6-5

LEDs

drive SBB, 6-8

HSZ10-AA controller, 6-5

HSZ10-AA controller codes, 6-5

power supply, 6-7

power supply and shelf, 6-8

shelf and power supply, 6-8

types of, 6-5

Logical drive map, 4-4

Logical units (LUNs)

description, 3-6

modes of operation, 3-6

LUNs

definition, 3-2

LUN status, 6-1

M

Mirrored data, 3-2

Monitoring

operations with LED indicators, 6-5

operations with the DEC RAID utilities, 6-4

through the HSZ10-AA controller, 6-1

through the StorageWorks shelf, 6-4

user monitoring methods, 6-4

O

Operating systems supported, C-1

Operations, 6-1

Overview, 1-1

P

Parity check/repair, 6-14

Partitions, 3-7

Performance, 1-4

Physical drive map, 4-4

Powering on

the subsystem, 5-6

Power supply

how to replace, 6-11

LED indicators, 6-7

specifications, A-1

Power supply and shelf

LED indicators, 6-8

Power unit

specifications, A-1

Product

attributes, 1-3

description, 1-1

highlights, 1-3

overview, 1-1

Purpose, ix

R

RAID 0 description, 3-2

RAID 1 description, 3-3

RAID 3 description, 3-4

RAID 5 description, 3-5

RAID overview, 3-1

Reconstruction, 3-8, 6-13

Redundancy, 1-4

Regeneration, 3-8

Related documents, x

S

SBBs, 2-4

adapter SBBs, 2-5

disk SBBs, 2-4

power SBBs, 2-4

- SCSI driver support, 2–5
- SCSI interconnects/host adapters, 2–5
- Segment, 3–1
- Shelf and power supply
 - LED indicators, 6–8
- Shelf power supply
 - replacing, 6–11
- Software, 2–4
- Software, upgrading, 6–15
- Specifications
 - environmental, A–3
 - environmental stabilization, A–2
 - for DEC RAID Subsystem, A–1
 - general, A–1
 - power supply, A–1
 - power unit, A–1
- Status
 - drive, 6–3
 - LUN, 6–1
- StorageWorks shelf
 - monitoring through, 6–4
- Subsystem features, 1–3
- Swapping
 - cold, 6–11
 - hot, 6–11
 - warm, 6–11

- SZ200
 - physical arrangement of the base configuration, 4–4
- SZ200 base configuration, 4–1
- SZ200 description, 2–1

T

- Temperature
 - ranges, A–2
- Terminating
 - the SCSI bus, 5–1
- Thermal stabilization, A–2
- Troubleshooting, 8–1
 - table of, 8–2

V

- Verifying cables and connectors, 5–5
- Verifying functionality, 5–6
- Vertical mounting kit description, 2–1

W

- Warm swapping, 6–11

