

Compaq StorageWorks

RAID Array 3000 Pedestal
Storage Subsystem

Hardware User's Guide
EK-SMCPO-UG. B01

Compaq Computer Corporation
Houston, Texas

While Compaq Computer Corporation believes the information included in this manual is correct as of the date of publication, it is subject to change without notice. Compaq makes no representations that the interconnection of its products in the manner described in this document will not infringe existing or future patent rights, nor do the descriptions contained in this document imply the granting of licenses to make, use, or sell equipment or software in accordance with the description. No responsibility is assumed for the use or reliability of firmware on equipment not supplied by Compaq or its affiliated companies. Possession, use, or copying of the software or firmware described in this documentation is authorized only pursuant to a valid written license from Compaq, an authorized sublicensor, or the identified licensor.

Commercial Computer Software, Computer Software Documentation and Technical Data for Commercial Items are licensed to the U.S. Government with Compaq's standard commercial license and, when applicable, the rights in DFAR 252.227 7015, "Technical Data-Commercial Items."

© 1999 Compaq Computer Corporation.
All rights reserved. Printed in U.S.A.

Compaq, the Compaq logo, DIGITAL, DIGITAL UNIX, DECconnect, HSZ, HSG, StorageWorks, VMS, OpenVMS Registered in the United States Patent and Trademark Office.

UNIX is a registered trademark in the United States and other countries exclusively through X/Open Company Ltd. Windows NT is a registered trademark of the Microsoft Corporation. Sun is a registered trademark of Sun Microsystems, Inc. Hewlett-Packard, TACHYON, and HP-UX are registered trademarks of the Hewlett-Packard Company. IBM and AIX are registered trademarks of International Business Machines Corporation. All other trademarks and registered trademarks are the property of their respective owners.

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the manuals, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense. Restrictions apply to the use of the local-connection port on this series of controllers; failure to observe these restrictions may result in harmful interference. Always disconnect this port as soon as possible after completing the setup operation. Any changes or modifications made to this equipment may void the user's authority to operate the equipment.

Warning!

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

Achtung!

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

Attention!

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

JAPAN

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

USA

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

Contents

Revision Record	xi
------------------------------	-----------

About This Guide	xiii
-------------------------------	-------------

1 Product Overview

1.1	Product Description	1-1
1.2	Pedestal Features	1-4
1.3	Pedestal Cabinet	1-4
1.4	Pedestal Components	1-6
1.4.1	StorageWorks Building Blocks (SBBs)	1-6
1.4.2	RAID Array Controller	1-6
1.4.3	Pedestal Power Supplies	1-7
1.4.4	Uninterruptable Power Supply	1-7
1.4.5	Environmental Monitor Unit (EMU)	1-8
1.4.6	UltraSCSI Bus	1-9
1.5	Connecting the Pedestal to a Host System	1-12
1.6	Specifications	1-17

2 RAID Array Controller

2.1	Controller Overview	2-1
2.2	Controller Features	2-2
2.3	Controller Reset and LED Indicators	2-6
2.4	Flexible RAID Set Configuration	2-7
2.5	Performance Enhancements	2-8
2.5.1	Custom Components	2-8
2.5.2	Efficient Write and Read Algorithms	2-8
2.6	RAID Levels Supported	2-10
2.6.1	RAID 0	2-11
2.6.2	RAID 1	2-13
2.6.3	RAID 0+1	2-13

2 RAID Array Controller (continued)

2.6.4	RAID 4	2-15
2.6.5	RAID 5	2-16
2.6.6	JBOD.....	2-18
2.7	System Parameters	2-18
2.8	Redundant Operation.....	2-20
2.8.1	Initialization	2-20
2.8.2	Message Passing	2-20
2.8.3	Failover	2-20
2.9	Environmental.....	2-21
2.9.1	Backup Power Management.....	2-21
2.9.2	Voltage Monitoring	2-22
2.9.3	Temperature Monitoring.....	2-24

3 Maintenance

3.1	Introduction.....	3-1
3.2	Pedestal Status and Power LEDs	3-2
3.3	SBB Status LEDs	3-2
3.4	Controller LEDs.....	3-3
3.5	EMU Error Reporting.....	3-4
3.5.1	EMU Error Conditions.....	3-4
3.6	Replacing Components (FRU's).....	3-5
3.6.1	Removing the Pedestal Door.....	3-5
3.6.2	Replacing an SBB.....	3-6
3.6.3	Replacing a Power Supply	3-7
3.6.4	Replacing the RAID Array Controller.....	3-9
3.6.5	Replacing the EMU Board.....	3-10
3.6.6	Replacing the UPS.....	3-12
3.7	Differential/Wide UltraSCSI Bus	3-12
3.7.1	Reconfiguring the SCSI Bus	3-13
3.8	Replacing the Controller Memory Cache Modules	3-17

4 Expansion Pedestal Option

4.1	Product Description	4-1
4.2	Expansion Pedestal Cabinet	4-2
4.3	Expansion Pedestal Components.....	4-4
4.4	Reconfiguring Base Pedestal UltraSCSI Bus.....	4-5

5 Second Controller Option

5.1	Introduction	5-1
5.2	Installing the Upgrade	5-2
5.2.1	Save Existing Configuration.....	5-4
5.2.2	Install Two SIMMs	5-6
5.2.3	Install the New Controller	5-8
5.2.4	Restoring Configuration	5-9
5.3	Configuring a Dual Controller Installation for a Single Serial Port	5-11

Figures

1-1	RAID Array 3000 Pedestal Enclosure (Drives Optional).....	1-2
1-2	Pedestal Front Panel Major Components.....	1-5
1-3	Pedestal Rear Panel Power Supplies	1-5
1-4	EMU Circuit Board Location.....	1-8
1-5	Pedestal Rear Panel Components.....	1-9
1-6	UltraSCSI Bus Configuration Switch.....	1-10
1-7	Slot Locations and SCSI ID Addresses	1-10
1-8	Slot Locations and SCSI ID Addresses	1-11
2-1	Bridging the Gap Between the Host and the Pedestal.....	2-2
2-2	Units Created from Storagesets, Partitions, and Disk Drives	2-3
2-3	Controller Front Panel	2-7
2-4	RAID 0 Write.....	2-12
2-5	Diagram of a RAID 1 Write.....	2-13
2-6	Diagram of a RAID 0+1 Write	2-14
2-7	Diagram of a RAID 4 Write.....	2-15
2-8	Diagram of a RAID 5 Write.....	2-17
3-1	Pedestal Status LEDs	3-2
3-2	Disk Drive Status LEDs.....	3-3
3-3	Removing Pedestal Door	3-5

Figures (continued)

3-4	Replacing an SBB	3-7
3-5	Replacing a Power Supply	3-8
3-6	Removing the Controller from the Pedestal	3-9
3-7	Remove EMU Connector Nuts	3-10
3-8	Remove Screw and Panel	3-14
3-9	Location of SCSI Bus Configuration Switch	3-15
3-10	Configuration Switch	3-15
3-11	Remove Controller	3-16
3-12	Release Locking Clips	3-17
3-13	Remove Installed Memory Modules	3-17
3-14	Install Replacement Modules	3-18
3-15	Pivot Module Down to Secure	3-19
4-1	Expansion Pedestal	4-1
4-2	Expansion Pedestal Slot Locations and ID Addresses	4-3
4-3	Rear Panel Power Supplies	4-4
4-4	EMU Circuit Board Location	4-5
4-5	Remove Side Cover from Base Pedestal	4-7
4-6	Remove SCSI Bus Terminator	4-7
4-7	Disconnect SCSI Cable	4-8
4-8	Connect SCSI Jumper	4-8
4-9	Remove Connector Knockout Plate	4-9
4-10	Connect SCSI Cable	4-10
4-11	Set Configuration Switch to 7	4-10
4-12	Configuration Switch	4-11
4-13	Reconfigured SCSI Bus Addresses	4-11
4-14	Connect SCSI Cable Between Pedestals	4-12
4-15	Power Cable Connections	4-13
4-16	Transfer Drives from Base to Expansion Pedestal	4-14
5-1	Second Controller Slot Position	5-2
5-2	Insert Module into SIMM Connector	5-3
5-3	Pivot Module Down to Seat	5-4
5-4	Remove Controller from Top Slot	5-4

Tables

1-1	RAID Array 3000 Part Numbers and Model Descriptions	1-3
1-2	Pedestal Specifications	1-12
1-3	Pedestal Physical and Power Specifications	1-13
2-1	Controller Specifications	2-4
2-2	LED/Reset Switch Interface	2-6
2-3	RAID Levels Supported.....	2-10
2-4	Pedestal RAID Set Restrictions.....	2-10
2-5	RAID 0+1 Example.....	2-14
2-6	System Parameters.....	2-19
2-7	Backup Power Management	2-21
2-8	Acceptable System Voltage Levels.....	2-22
2-9	Acceptable Termination Voltage Levels	2-23
2-10	Acceptable 12 Volt Levels.....	2-23
2-11	Acceptable External Temperature Voltage Levels	2-24
2-12	Acceptable Board Temperature Voltage Levels	2-24
3-1	Disk Drive SBB Status LEDs.....	3-3
3-2	SCSI Bus Length and External Cables	3-11
3-3	Assigned Slot Device Addresses in the Pedestal	3-11

Revision Record

This Revision Record provides a concise publication history of this guide. It lists the revision levels, release dates, and reasons for the revisions.

The following revision history lists all revisions of this publication and their effective dates. The publication part number is included in the Revision Level column, with the last entry denoting the latest revision. This publication supports the StorageWorks RAID Array 3000 Storage Subsystem.

Revision Level	Date	Summary of Changes
EK–SMCPO–UG. A01	January 1998	Original release.
EK–SMCPO–UG. B01	July 1999	Adds Pedestal-to-Host connection diagrams to Chapter 1 and a dual controller configuration procedure to Chapter 5.

About This Guide

This section identifies the audience of this guide and describes the contents (chapter-by-chapter) and structure. In addition, this section includes a list of associated documents and the conventions used in this guide.

Intended Audience

This guide is intended for installers and operators of the RAID Array 3000 pedestal storage subsystem. Installing the subsystem requires a general understanding of basic SCSI terminology and SBus product installation procedures.

Document Structure

This guide contains the following chapters:

Chapter 1: Product Overview

Product Overview provides an overview and a physical hardware description of the single pedestal RAID Array 3000 storage subsystem. It includes the major features, a brief description of the major components, and the specifications for the pedestal. The chapter also contains a series of pedestal-to-host cabling diagrams.

Chapter 2: RAID Array Controller

RAID Array Controller describes the major features and characteristics of the controller in the RAID Array 3000. It also lists the RAID levels supported by the subsystem and a brief description of each level. Redundant operation and environmental considerations (i.e. backup power management, voltages, and temperature) are covered at the end of the chapter.

Chapter 3: Maintenance

Maintenance describes how to interpret the status and power LEDs on the pedestal as a troubleshooting aid in case of a problem. The chapter also describes how to replace a Field Replaceable Unit (FRU) in the pedestal, how to re-configure the SCSI bus, and how to add or replace a memory SIMM module in the RAID controller.

Chapter 4: Expansion Pedestal Option

Expansion Pedestal Option describes the major features and characteristics of the pedestal expansion kit option. It also explains how to reconfigure the SCSI bus in the base pedestal to accommodate the added storage capability of the expansion pedestal.

Chapter 5: Second Controller Option

Second Controller Option describes how to install a second controller in the pedestal for redundancy. In addition, a procedure describing how to configure the pedestal for dual-controller operation when only one serial port is available on the host is included at the end of the chapter.

Associated Documents

In addition to this guide, the following documentation is useful to the reader:

Table 1 Associated Documents

Document Title	Order Number
<i>StorageWorks RAID Array 3000 Configuration and Maintenance Guide</i>	<i>EK-SMCS2-UG</i>
<i>Installation Instructions for the RAID Array 3000 Pedestal Expansion Kit Option Product Note</i>	<i>EK-SMCPM-PN</i>
<i>Installation Instructions for the RAID Array Replacement Controller –SWXRC-03</i>	<i>EK-SMCPL-PN</i>
<i>Command Console 2.1 for RAID Array 3000, User's Guide</i>	<i>AA-RBF2B-TE</i>

Conventions

This guide uses the following documentation conventions:

Table 2 Style Conventions

Style	Meaning
plain monospace type	Text
boldface type	For the first instance of terms being defined in text, or both.
<i>italic type</i>	<i>For emphasis, manual titles, chapter summaries, keyboard key names.</i>

Conventions (continued)

Table 3 Nomenclature Convention

RAID Advisory Board Description	RAID Array 3000 Usage
RAID 0	STRIPset
RAID 1	MIRRORset
RAID 0+1	STRIPED MIRRORset
RAID 4	STRIPED with a Fixed parity drive
RAID 5	STRIPED with a Floating parity drive

Support and Services

Who to contact in the Americas

Information and Product Questions: Local Sales Office / StorageWorks Hotline
1-800-786-7967

Installation Support: Contact the Compaq Distributor where the Storage Solution was Purchased / Local Compaq Sales Office.

Compaq Multivendor Customer Service (MCS):

Installation Contact the Compaq Customer Support Center (CSC).

Warranty Contact the Compaq Customer Support Center (CSC) for warranty service after solution is installed and operating.

Remedial Contact the Compaq Customer Support Center (CSC)

Note: A Service Contract is recommended when the equipment is out of warranty. Contact the local Compaq Sales Office.

Customer Support Center (CSC) 1 800-354-9000

Who to contact in Europe

Information and Product Questions:	Contact the Compaq Distributor or reseller
Installation Support and Installation:	Contact the Compaq Distributor or reseller from whom the Storage Solution was purchased.
For Warranty Service	See the Warranty Card packaged with the product.
For Remedial Service	Contact the Compaq Distributor or reseller from whom the Storage Solution was purchased. Note: A Service Contract is recommended when the equipment is out of warranty.

Who to contact in Asia Pacific

For all services, contact the Compaq Distributor or reseller from whom the equipment was purchased

1

Product Overview

This chapter provides an overall description of the RAID Array 3000 Storage System and its components. Examples of Host/Storage System connections and a list of technical and environmental specifications are included at the end of the chapter.

NOTE

This guide is the Hardware User's Guide. For configuration information, refer to the *Getting Started RAID Array 3000 for Windows NT –Intel Installation Guide* and the *StorageWorks Command Console (SWCC) 2.1 User's Guide*.

1.1 Product Description


The RAID Array 3000 storage subsystem (Figure 1–1) is a desk-side storage system (subsystem) offering the basic components required to create a user-designed storage array with two 16-bit, differential UltraSCSI bus host interfaces. The pedestal can accommodate up to seven 3½-inch storage devices. The devices, referred to as StorageWorks Building Blocks or SBBs, are disk drives from the StorageWorks family of storage devices. The release notes that accompanies the subsystem lists the software solutions and disk drives that are supported. A battery backup subsystem is included as part of the pedestal enclosure in the form of a free-standing UPS (Uninterruptable Power Supply). In case of a power failure, the UPS provides a temporary backup for cache while the subsystem flushes to disks.

The RAID Array 3000 offering also includes option kits designed to increase the storage capacity and enhance the performance of the subsystem. The first is an expansion pedestal (second enclosure) designed to increase the storage capacity of the subsystem to a maximum of 14 drives. The pedestal expansion kit option is described in detailed in Chapter 4 of this guide.

The second option allows the addition of a second RAID controller to the subsystem for redundancy. The second controller operates in conjunction with the installed controller to protect data during a malfunction. Chapter 5 describes how to install the redundant controller option and how to reconfigure the subsystem to accommodate it.

The RAID Array 3000 pedestal enclosure and its associated options are listed and described in Table 1-1. Figure 1-1 shows the pedestal with a full complement of drives (optional) for completeness.

Figure 1-1 RAID Array 3000 Pedestal Enclosure (Drives Optional)



3000-01A

The RAID 3000 pedestal is equipped with a dual-channel RAID controller that supports all of the UltraSCSI bus features. It also contains an Environmental Monitor Unit (EMU) board for environmental monitoring and error detection.

Table 1–1 RAID Array 3000 Pedestal Part Numbers and Model Descriptions

DIGITAL Part No.	Item Description
DS-SWXRA-GA	<p>RA3000 pedestal subsystem with one controller, 120 V.</p> <p>Includes: Seven-slot pedestal for wide UltraSCSI SBBs, one HSZ22 two-channel controller with 16 MB cache, Environmental Monitor Unit (EMU), two 204 watt power supplies with fans, five meter host SCSI cable (BN37A), BN38E-OB adapter, one 120-volt Un-interruptable power supply (UPS), and North American power cords. Disks are not included.</p> <p>Requires: Solutions Software Kit for platform, host adapter, and disks.</p> <p>Options: Second HSZ22 controller, seven-disk SBB expansion pedestal, and cache memory upgrade.</p>
DS-SWXRA-GC	<p>RA3000 pedestal subsystem with one controller, 230 V.</p> <p>Includes: Seven-slot pedestal for wide UltraSCSI SBBs, one HSZ22 two-channel controller with 16 MB cache, Environmental Monitor Unit (EMU), two 204 watt power supplies with fans, five meter host SCSI cable (BN37A), BN38E-OB adapter, one 230-volt Un-interruptable power supply (UPS), and North American power cords. Disks are not included.</p> <p>Requires: Solutions Software Kit for platform, host adapter, and disks.</p> <p>Options: Second HSZ22 controller, seven-disk SBB expansion pedestal, and cache memory upgrade.</p>
DS-HSZ22-AA	Second HSZ22 controller with 16 MB of cache and two 16 MB SIMMs for mirror cache, one BN37A-05 five-meter UltraSCSI cable, and one BN38E-OB adapter.
DS-SWXRA-GD	Expansion Pedestal (120/240 V) with slot space for seven additional UltraSCSI disk drives.
DS-SWXRA-GR	Single 204-watt power supply for RA3000, 120/230 V for on-site spare.

1.2 Pedestal Features

The major features of the pedestal are:

- Two differential 16-bit UltraSCSI host buses
- Seven 3½-inch disk drive SBB slots
- One dual-channel RAID array controller
- Second controller option for redundancy
- Expansion pedestal option allowing up to fourteen SBB slots in a dual-pedestal subsystem configuration
- Memory cache expansion option for the controller
- Redundant power provided by two fan-cooled universal ac input power supplies (50/60 Hz, 100 to 240 V ac)
- Cache backup provided by an external Uninterruptable Power Supply (UPS)
- Environmental monitor unit (EMU) for error detection
- The ability to hot swap SBBs without powering down the system

1.3 Pedestal Cabinet

The pedestal cabinet is a modular free-standing storage enclosure that is completely self-contained. It has two fan-cooled power supplies, an internal EMU circuit board, and a RAID array controller with front panel display and control.

Figure 1–2 shows the major components in the pedestal enclosure.

Figure 1–3 identifies the items on the rear panel power supplies. The characteristics of the pedestal cabinet are outlined below.

- The disk drive storage capacity is seven 3½-inch disk drive SBBs
- The subsystem slots are numbered 0 through 6 from top to bottom
- There are two 68-pin VHDCI female SCSI connectors on the rear panel which interconnect the host system to the RAID controller in the pedestal
- The rear panel also contains an alarm switch, a UPS monitor connector, an external fault condition connector, and a serial port connector (for controller configuration)
- The pedestal is equipped with an internal configuration switch which sets the SCSI ID addresses of the controller and the storage devices

Figure 1–2 Pedestal Front Panel Major Components





Figure 1–3 Pedestal Rear Panel Power Supplies



1.4 Pedestal Components

The major components in the pedestal subsystem include:

- Dual-channel RAID array controller
- Two 16-bit single-ended split SCSI buses
- Environmental Monitor Unit (EMU)
- Two universal 50/60 Hz, 120 or 240 Vac power supplies
- Separate free-standing backup power supply (UPS)

1.4.1 StorageWorks Building Blocks (SBBs)

The pedestal has seven 3½-inch disk drive SBB slots. The number of drives that make up each configuration of the array is left to the end user with a maximum of seven storage devices.

1.4.2 RAID Array Controller

The controller contains two Wide/UltraSCSI/differential host channels and two Wide/UltraSCSI/Single-Ended disk channels. In dual-controller configurations, the controllers support fully automatic and smooth controller failover.

The controller supports one or two standard 72-pin cache SIMMs of up to 64 MB. In a redundant controller setup, both controllers must have identical cache configurations and the total usable cache (per controller) will be half the amount installed. Thus, in a single controller setup the maximum usable cache is 128 MB while a redundant setup has a maximum usable cache of 64 MB (per controller).

The RAID Array controller contains the following features:

- Single PCB form factor for inclusion in the enclosure
- Support for dual hot-swap controller operation
- Dual differential Ultra-Wide SCSI host channels
- Dual single-ended Ultra-Wide SCSI disk channels
- RAID level 0, 1, 0+1, 4, 5, and JBOD support
- EMU support
- Cluster support for Windows NT
- 32 Logical Units (LUNs) per host channel (some operating systems may be limited to 8)
- Support for Hot and Warm spare disks
- UPS backed write caching

- Per LUN write cache/write back selection
- Configuration/Maintenance via RS-232 or host SCSI channel using SWCC (StorageWorks Command Console)
- Update of firmware via host channel

1.4.3 Pedestal Power Supplies

The Pedestal has two interchangeable, air-cooled, AC power supply modules located at the rear of the unit. The power supplies provide redundant power if one of the units should malfunction. Each supply provides +5 and +12 Vdc to power the RAID controller, EMU, and the storage SBBs in the pedestal. In addition, each unit contains a high-speed fan for pedestal cooling. The upper power supply is designated as "A" and the lower as "B". The unit contains an ac power receptacle, a power status LED, a fan, and a latching slider switch to secure the supply in the pedestal.

1.4.4 Uninterruptable Power Supply (UPS)

The UPS is separate and free-standing unit designed to protect the pedestal from problems associated with poor quality AC power or a complete loss of AC power. The UPS is connected between the AC outlet and the line input of the pedestal power supply "B" to provide battery backup power.

The major features of the UPS include Cell Saver Technology (doubles battery life and speeds recharge time), hot-swap batteries, and network surge protection. The front panel display has user controls (LEDs and control buttons) and the rear panel contains a COMM port, which provides UPS status to the EMU in the pedestal. The rear panel also contains the network surge protector, a reset button, and four power receptacles. An audible alarm is activated when input power fails, as a Low Battery Warning, or whenever the UPS is in need of servicing.

The UPS automatically recharges its battery when power is returned following a power failure. Recharge time is four to six hours depending on the energy requirements of your load and the length of the power outage.

Its own installation, operation, and service manual support the UPS. The manual describes the UPS in detail and is part of the documentation set enclosed with your subsystem.

1.4.5 Environmental Monitor Unit (EMU)

The EMU is an internal circuit board that monitors the operation of the pedestal. The EMU monitors power supply voltages, fans, temperatures which are reported to the user, and controls (turns on and off) the audible alarm and status LED on the front panel. The EMU also reports the subsystem status to the controller that reports to the host, and has the capability of exchanging signals with auxiliary devices and controllers.

The EMU is located internally in the top rear of the pedestal as shown in Figure 1-4). It is connected to the SCSI bus and powered by an internal cable. The following external components on the rear panel of the pedestal are part of the EMU (see Figure 1-5):

- An alarm switch (S1) that enables (up) or disables (down) the audible alarm
- A power monitor connector (UPS) allows the EMU to monitor and report the status of a battery backup power supply
- An External Fault Condition connector allows the EMU to monitor the status of a user-selected device

Figure 1-4 EMU Circuit Board Location





Figure 1–5 Pedestal Rear Panel Components

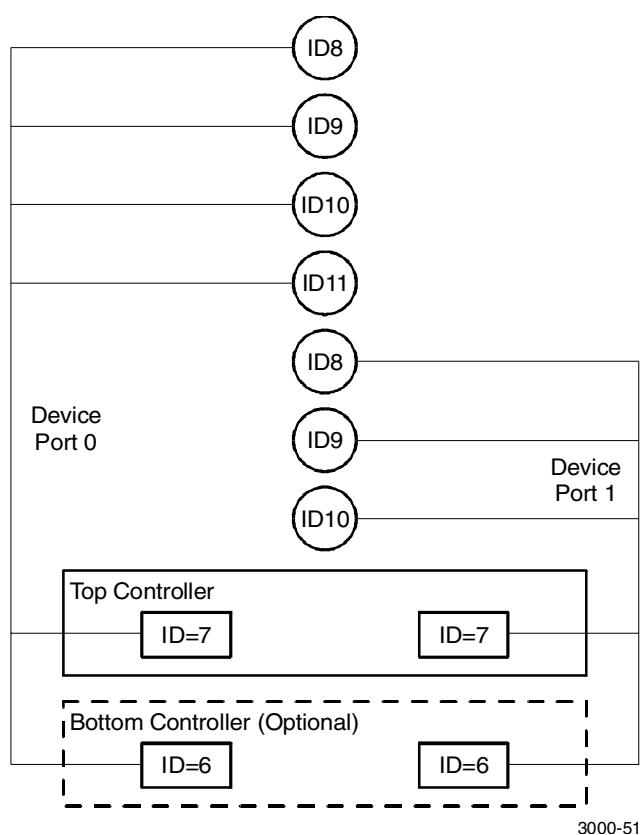
3000-23A

1.4.6 UltraSCSI Buses

The pedestal contains two, 16-bit, single-ended, wide UltraSCSI buses (factory-configured as a split bus) that connects the controller(s) to the disk drives.

The split-bus arrangement is divided into two-bus paths designated *device port 0* and *device port 1* as shown in Figure 1–6. Port 0 connects the controller to the upper four devices in the pedestal (IDs 8 through 11) and port 1 connects the controller to lower three devices (IDs 8 through 10). The device addresses on the bus are set at the factory by an internal configuration switch (see Figure 1–7).

Figure 1–6 UltraSCSI Bus Port and Default SCSI ID Assignments



When set to a specific position, the switch controls the addresses of each SBB slot. Figure 1–8 identifies the pedestal slot locations and their corresponding SCSI ID addresses for each device port.

The subsystem can be reconfigured for “through-bus” operation by modifying the bus and resetting the internal configuration switch. The information needed to reconfigure the bus from “split-bus” to a “through-bus” configuration is described in Chapter 4 of this guide (*Expansion Pedestal Option*).

Figure 1-7 UltraSCSI Bus Configuration Switch





Figure 1-8 Slot Locations and SCSI ID Addresses



1.5 Connecting the Pedestal to a Host System


This section illustrates how to connect four possible RA3000/Host configurations. The configurations are:

- Single host, single adapter configuration
- Single host, dual adapter/dual controller configuration
- Single host, single adapter/dual controller configuration (using a Y-cable)
- Dual host, single adapter/single controller configuration

NOTE


The following illustrations show the Expansion Pedestal Option as part of the storage system installation. See Chapter 4 of this guide for a full description of the Expansion Pedestal.

Figure 1–9 Single Host, Single Adapter Cabling Diagram




- 1 Host #1 VHDCI Connector
- 2 Host #0 VHDCI Connector
- 3 SCSI Device Bus Out Connector D0
- 4 SCSI Cable BN37A-05 (host adapter connection made using Technology Adapter cable BN38E-OB, not shown)
- 5 SCSI cable 17-04454-01
- 6 SCSI Device Bus In Connector D0

Figure 1–10 Single Host, Dual Adapter/Dual Controller Cabling Diagram



- 1 Host #1 VHDCI Connector
- 2 Host #0 VHDCI Connector
- 3 SCSI Device Bus Out Connector D0
- 4 (2) SCSI Cables BN37A-05 (host adapter connections made using two Technology Adapter BN38E-OB cables, not shown)
- 5 SCSI cable 17-04454-01
- 6 SCSI Device Bus In Connector D0


Figure 1–11 Single Host, Single Adapter/Dual Controller (Y-cable connection) Cabling Diagram



SHR-1313

- 1 Host #1 VHDCI Connector
- 2 Host #0 VHDCI Connector
- 3 SCSI Device Bus Out Connector D0
- 4 SCSI Y-Cable BN-21W-OB
- 5 (2) SCSI Cables BN37A-05 (connections to Y-cable made using two Technology Adapter BN38E-OB cables, not shown)
- 6 SCSI Cable 17-04454-01
- 7 SCSI Device Bus In Connector D0

Figure 1–12 Dual Host, Single Adapter/Single Controller Cabling Diagram



- 1 Host #1 VHDCI Connector
- 2 Host #0 VHDCI Connector
- 3 SCSI Cable BN37A-05 (host adapter connection made through Technology Adapter BN38E-OB cable, not shown)
- 4 SCSI Cable BN37A-05 (host adapter connection made through Technology Adapter BN38E-OB cable, not shown)
- 5 SCSI Device Bus Out Connector D0
- 6 SCSI Cable 17-04454-01
- 7 SCSI Device Bus In Connector D0

1.6 Specifications

Table 1–2 Pedestal Technical Specifications

Technical Specifications	
Item	Description
Cabinet	Pedestal with seven (7) disk SBB slots Expansion pedestal with an additional 7 slots
Controller	HSZ22
Controller cache	16 MB standard Upgrades to 128 MB for a two controller pair
Backup for cache	Standard Un-interruptable power supply (UPS)
Mirrored write-back cache	Yes
Device channels per controller	2
Maximum disks per device port	14
Dual active controllers	Yes, order HSZ22-AA and second solutions software kit
Host interface	UltraSCSI wide differential
Drive interface	UltraSCSI wide single-ended
RAID levels supported	0, 1, 0+1, 4, 5
Non-RAID disk support	Yes (JBOD)
Sustained I/O rate	4,400 I/O's per second per controller pair
RAID 5 sustained transfer rate	28 MB per second per controller pair
Maximum transfer rate	40 MB per second per controller pair
Redundant fans	Yes
Redundant power supplies	Yes
Global disk spares	Yes
Environmental Monitoring Unit	Yes, monitors power and temperature
Setup/control lines	One serial
Serviceability	Hot-swap components
RAID manager GUI support	StorageWorks Command Console (SWCC) 2.0 available for all platforms. NetWare requires serial line connection
Regulatory approvals	UL, CSA, TUV, FCC, CE MARK, C TICK, BCIQ, VCCI

Table 1–3 Pedestal Physical and Power Specifications

Physical Specifications	
Item	Dimension
Height	564 mm
Width	254 mm
Depth	494 mm
Rear Clearance (air exhaust)	305 mm
Front Clearance (door opening)	305 mm
Weight (no devices)	19.5 kg
Power Specifications	
Item	Rating
Input power	110-240 Vac, 50/60 Hz, single phase, 12A/6A
Heat dissipation	3070 BTUs/hr.
Temperature (optimal, minimum required)	18 to 24° C, 10 to 40° C
Altitude	Up to 2,400 m
Air quality	Not to exceed 500,000 particles/ft ³ for air at a size of 0.5 micron or larger
Total power per power supply (Total of +12 Vdc and +5 Vdc outputs)	204 Watts, maximum
Nominal output voltages	+5 Vdc @ 15 amps, maximum +12 Vdc @ 12 amps, maximum
Device startup time	4-second internal, minimum

2

RAID Array Controller

This chapter describes the major features and characteristics of the RAID array controller in the RAID Array 3000 subsystem. The number of devices supported by the controller may be limited by the enclosure.

2.1 Controller Overview

The RAID Array controller provides high performance, high-availability access to SCSI disk array subsystems along a UltraSCSI/Wide SCSI bus. With a modular hardware design and an intuitive configuration utility, the controller is designed to meet a wide range of storage needs.

The controller consists of a single 5 1/2" x 8" PCB mounted in a sheet metal subassembly. The package consists of the controller PCB, a 300-pin connector, mechanical insertion assists, and an LED/reset switch interface. All signals to the controller are routed through the backplane connector.

The unit is configured with two Ultra Wide, differential SCSI host channels capable of transferring data to and from the host at rates up to 40 MB/s. The host SCSI IDs are configurable via the Host Parameters and can support 32 deep tagged queuing. The controller is also configured with two Ultra Wide, single-ended SCSI disk channels capable of transferring data to and from the disk drives at rates up to 40 MB/s. Each channel can support up to 15 drives (14 in redundant controller subsystems).

The controller has two SIMM connectors for up to 128 MB of cache memory. The SIMM connectors form a mirrored pair when the controllers are configured in a redundant controller configuration; otherwise they are fully accessible by the controller. In a redundant controller setup, both controllers must have identical cache configurations and the total usable cache (per controller) will be half the amount installed. Thus, in a single controller setup the maximum usable cache is 128 MB while a redundant setup has a maximum usable cache of 64 MB (per controller).

There are two configurations for redundant pairs of controllers: *Active/Active Failover* mode and *Active/Passive Failover* mode. In *Active/Active Failover*, each controller in the redundant pair has one active SCSI host port and one passive SCSI host port. Redundancy Groups (Virtual LUNs) can be mapped only to one active host port and are not accessible from the passive port or the other controller (i.e. partitioned model).


In *Active/Passive Failover*, one controller in the redundant pair has both SCSI host ports active and the other controller is in a standby passive mode. Redundancy Groups (Virtual LUNs) can be mapped to either SCSI host port or to both as in the single controller model.

In both cases a single controller failure will not affect the subsystem because the surviving controller will take over.

2.2 Controller Features

The controller is the intelligent bridge between the host and the devices in the pedestal. From the host's perspective, the controller is simply another SCSI device connected to one of its I/O buses. Consequently, the host sends its I/O requests to the controller just as it would to any other SCSI device. Figure 2–1 shows the role of the controller between the host and the pedestal.

Figure 2-1 Bridging the Gap Between the Host and the Pedestal




From the pedestal's perspective, the controller receives the I/O requests from the host and directs them to the devices in the pedestal. Since the controller processes all the I/O requests, it eliminates the host-based processing that is typically associated with reading and writing data to multiple storage devices.

The controller does much more than simply manage I/O requests: it provides the ability to combine several ordinary disk drives into a single, high-performance storage unit called a *storage set*. Storage sets are implementations of RAID technology, also known as a “Redundant Array of Independent Disks”. Every storage set shares one important feature: whether it uses two disk drives or 14, each storage set looks like a single storage unit to the host.

You create storage units by combining disk drives into storage sets such as stripe sets, RAID sets, and mirror sets, or by presenting them to the host as single-disk units (see Figure 2–2).

Figure 2–2 Units Created from Storage sets, Partitions, and Disk Drives



- Stripesets (RAID 0) combine disk drives in serial to increase transfer or requests rates
- Mirrorsets (RAID 1) combine disk drives in parallel to provide a highly reliable storage unit
- RAID 4 provides striping with a fixed parity drive
- RAIDsets (RAID 5) combine disk drives in serial - just like stripesets - but also store parity data to ensure high reliability
- Stripe mirrorsets (RAID 0 + 1) combine mirrorsets in serial to provide the highest throughput and availability of any storage unit

Table 2–1 summarizes the features of the RAID 3000 controller.

Table 2–1 Controller Specifications

Item	Specifications
Environmental Monitoring	High Availability Fault Bus support via EMU
Processor	40 MHz, 32 bit LR33310 RISC CPU
System Bus Interface	Two, WIDE, Differential UltraSCSI channels
Configuration	Two RS232 serial ports
RAID Levels Supported	0, 1, 0+1, 4, 5
Drive Channels Supported	Two, WIDE, UltraSCSI single-ended channels
Number of Logical Drives (LUN's)	Up to 30 RAID sets, and up to 16 redundancy groups (LUNs) per RAID set
SCSI Channels	Two, UltraSCSI, 16-bit, single-ended
Metadata	20 blocks/disk (10240 bytes)
Largest Allowable Disk/RAID set/LUN	Two, 32 blocks (approximately 2.2 petabytes)
Non-RAID Device Support	Disk Drives (JBOD)
Drives Supported	StorageWorks 2, 4, and 9m GB SCSI and UltraSCSI drives.
Drive Reconstruct	Automatic with hot or warm spares
Disk Hot Swap	Yes
Disk Hot Spare (spinning)	Yes, global hot spare

Table 2–1 Controller Specifications (Continued)


Item	Specifications
Disk Warm Spare (not spinning)	Yes, global warm spare
Redundant Power Supplies	Yes
Redundant Controllers	Yes
Controller Failover	Yes, automatic
Controller Hot Spare	Yes (Active-Passive mode)
Controller Hot Swap	Yes
Cluster Support	Yes, Single (SCSI) bus cluster
Maximum number of units presented to host	64
Maximum host port transfer speed	20 MHz
Command Queuing	Yes, 64 commands (host and disk SCSI channels)
Heterogeneous Multi-Host Support	Yes
Mixed Drive Types	Yes
Configurable Reconstruct Time	Yes
Stripe Size (chunk size)	Variable
Write through Cache	Yes, User Selectable (default)
Write Back Cache	Yes, User Selectable (optional)
Write on Top	Yes
Write Gathering	Yes
Battery Backup for Cache	Yes, Pedestal Uninterruptable Power Supply (UPS)
Boot Capability	Bootable from RAID set (System dependent)
Number of Controllers /System	Two
Cache Support	Up to 256 MB (using two 128 MB industry-standard, 72-pin, 36-bit, 60 ns SIMMS)
FCC Rating	Class B
Environmental	
Temperature	5°C to 50°C operating, -40°C to +60°C non-operating
Relative Humidity	10% to 95% non-condensing (operating), 5% to 90%, non-condensing (non-operating)
Physical Size	8.55" deep, 5.03" wide, 1.6" high
Power Requirements	5 VDC @ 3 A, peak; - 12 VDC @ 1 A, peak

2.3 Controller Reset and LED Indicators

Figure 2–3 illustrates the front panel of the controller. All LEDs are numbered from left to right. The reset button (LED 0) flashes green about once every second (heartbeat) to indicate that the controller is operating normally. LEDs 1 through 4 display host and disk channel activity (amber). LED 5 (normally off) comes on red during a controller failure. The LED/Reset switch interface is defined in Table 2–2.

Table 2–2 LED/Reset Switch Interface

LED #	Name
0	Heart Beat Controller Reset Switch (green)
1	Host Channel 0 Activity LED (amber)
2	Host Channel 1 Activity LED (amber)
3	Disk Channel 0 Activity LED (amber)
4	Disk Channel 1 Activity LED (amber)
5	Fault LED (red)

Figure 2–3 Controller Front Panel

2.4 Flexible RAID Set Configuration

In addition to its flexible hardware design, the controller's firmware offers the user the flexibility to configure RAID sets in several different ways:

- RAID sets can comprise drives from any drive channel and SCSI ID.
- A RAID set can contain all the drives connected to the controller, a single drive, or any number of drives in between.
- The controller supports RAID Levels 0, 1, 0+1, 4, and 5. It also supports JBOD (Just a Bunch of Drives), allowing you to connect standalone disk drives (such as a system disk) to the controller without making them members of a RAID set.

- Each RAID set can be partitioned into smaller redundancy groups.
- The controller's host LUN Mapping feature makes it possible to map RAID sets differently to each host port. You make the same redundancy group appear on different LUNs to different hosts, or make a redundancy group visible to one host but not to another.
- Any drive may be designated as a hot or warm spare. Spares are global, meaning that in the event of a drive failure, the controller will search for the first available spare on any channel or SCSI ID and automatically begin rebuilding the failed drive's data.

2.5 Performance Enhancements

The controller employs a number of techniques to achieve as much performance as possible from its design.

2.5.1 Custom Components

To increase performance and reliability, the controller's core functions have been encapsulated in four custom ASIC (Application Specific Integrated Circuits) components as follows:

XOR ASIC: Used in the Exclusive -Or parity calculations employed by RAID levels 4 and 5.

DMA ASIC: Controls the data path hardware for the various I/O ports

CPU Interface ASIC: Supports the controller's MIPS R3000 RISC central processing unit.

Memory Controller ASIC: Controls the memory system and supports data movement on the internal bus at a maximum burst rate of 80 MB/second and a maximum sustainable rate of 60 MB/second.

2.5.2 Efficient Write and Read Algorithms

Standard RAID write operations that involve parity, such as those in RAID levels 4 and 5, require multiple, time-consuming steps:

1. Read data from the parity drive.
2. Read existing data from the target data drives.
3. Exclusive-Or the old parity, old data, and new data to generate new parity data.
4. Write the new parity data to the parity drive.
5. Write the new data to the target data drives.

The controller uses several techniques to streamline write operations and significantly improve performance. All the techniques use the controller's on-board cache, which can contain up to 64 MB of memory in the form of standard 72-pin, 60-nanosecond SIMMs.

NOTE

The controller will not operate without at least one 4 MB SIMM installed in its cache. Nor will it operate without either a backup or an un-interruptable power supply connected to the controller. Without a backup, data stored in the cache, but not yet written to the disk drives, would be lost in the event of a power interruption.

2.5.2.1 Write-Back Caching

When the host sends data to be written to a redundancy group the controller stores the data in its cache and immediately reports to the host it has completed the write. The controller eventually writes the data to the disk drives when the write can be done most efficiently, or when the controller must flush the cache to make room for other data or to prepare for a shutdown.

Write-back caching makes the host more responsive to the user, since the host does not have to wait for a lengthy RAID write before proceeding to another task.

2.5.2.2 Write Gathering

The controller will attempt to consolidate multiple writes destined for contiguous blocks and then write the entire data block in one operation. The controller stores the data in cache until it performs the write. Ideally, the controller will wait until it has gathered enough data to fill an entire stripe. This enables the controller to avoid reading from the parity and data drives before making the write. All it has to do is calculate parity from the data it already has in its cache, then write the data and parity to the drives. Even if the controller cannot accumulate enough data to fill a stripe, the consolidation of small writes can reduce the number of read/write operations that must take place.

2.5.2.3 Write On Top

If the host commands that data be written to disk, and data for that address is pending in the controller's cache, the controller writes the new data on top of the old in the cache. Only the new data is eventually written to the disk drives.

2.6 RAID Levels Supported

The RAID Array 3000 controller supports the following RAID levels:

Table 2–3 RAID Levels Supported

RAID Level	Description
0	Striping without parity
1	Mirroring
0+1	Striping and mirroring
4	Striping with fixed parity drive
5	Striping with floating parity drive
JBOD	“Just a Bunch of Drives”

NOTE

The controller stripes data in multi-block chunk sizes. Also, the controller does not support RAID level 3 or 0 with a one-block chunk size.

There are some restrictions you must adhere to when creating a RAID set on the RAID 3000 pedestal. The minimum and maximum number of drives required to support each RAID level is listed in Table 2–4.

Table 2–4 Pedestal RAID Set Restrictions

RAID Level	Min. No. ¹ of Drives	Max. No. ¹ of Drives
JBOD	1	1
0	2	56
1	2	56
0+1	4	16
4	3	56
5	3	56

¹Must be even number.

2.6.1 RAID 0

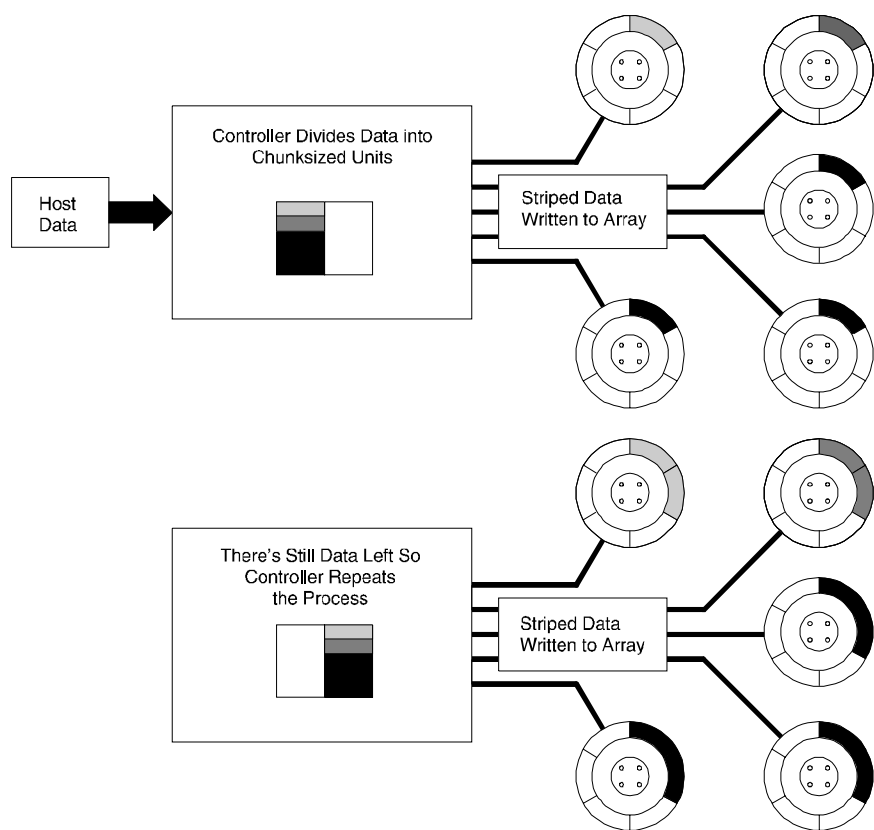
RAID 0 breaks up data into smaller chunks and writes each chunk to a different drive in the array. The size of each chunk is determined by the controller's chunk size parameter, which you set in the course of creating a RAID set.

The advantage of RAID 0 is its high bandwidth. By breaking up a large block of data into smaller chunks, the controller can use multiple drive channels to write the chunks to the disk drives. Furthermore, RAID 0 involves no parity calculations to complicate the write operation. Likewise, a RAID 0 read operation employs multiple drives to assemble a single, large data block. This makes RAID 0 ideal for applications such as graphics, video, and imaging that involve the writing and reading of large, sequential blocks. Figure 2–4 shows a diagram of a RAID 0 write.

CAUTION

The lack of parity means that a RAID 0-disk array offers absolutely no redundancy and thus cannot recover from a drive failure.

Figure 2–4 RAID 0 Write

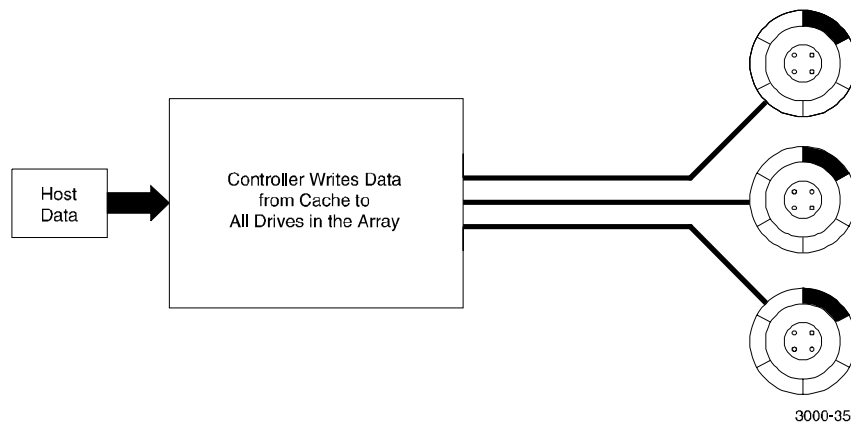


3000-34

2.6.2 RAID 1

RAID 1 (also known as mirroring or shadowing) takes data sent by the host and duplicates it on all the drives in an array. The result is a high degree of data availability, since you can lose all but one drive in the array and still have full access to your data. This comes at a price: a RAID 1 array requires multiple drives to achieve the storage capacity of a single drive. Figure 2–5 illustrates a RAID 1 write.


Figure 2–5 Diagram of a RAID 1 Write



A RAID 1 array will show up on the monitor as “degraded” when at least one drive fails, even if two or more members of the redundancy group remain in good working order. As long as at least two working drives remain in the array, you may continue to run the array in degraded mode without putting data in jeopardy.

2.6.3 RAID 0+1

RAID 0+1 combines RAID 0 (striping) with RAID 1 (mirroring). In RAID 0+1 write, the controller breaks up the data block from the host into smaller chunks, then writes the chunks to half the drives in the array, while writing duplicate chunks to the remaining drives.

Figure 2-6 Diagram of RAID 0+1 Write

In the event of a drive failure, a RAID 0+1 array will enter degraded mode and continue to operate by substituting the failed drive with its mirror.

When the controller creates a RAID 0+1 set, it first sorts the drives by channel number and SCSI ID. Then it stripes the data across every other drive and forms a mirrored pair with the first two drives, another mirrored pair with the second two drives, and so on. Table 2-5 describes how the controller uses the drives in a RAID 0+1 set.

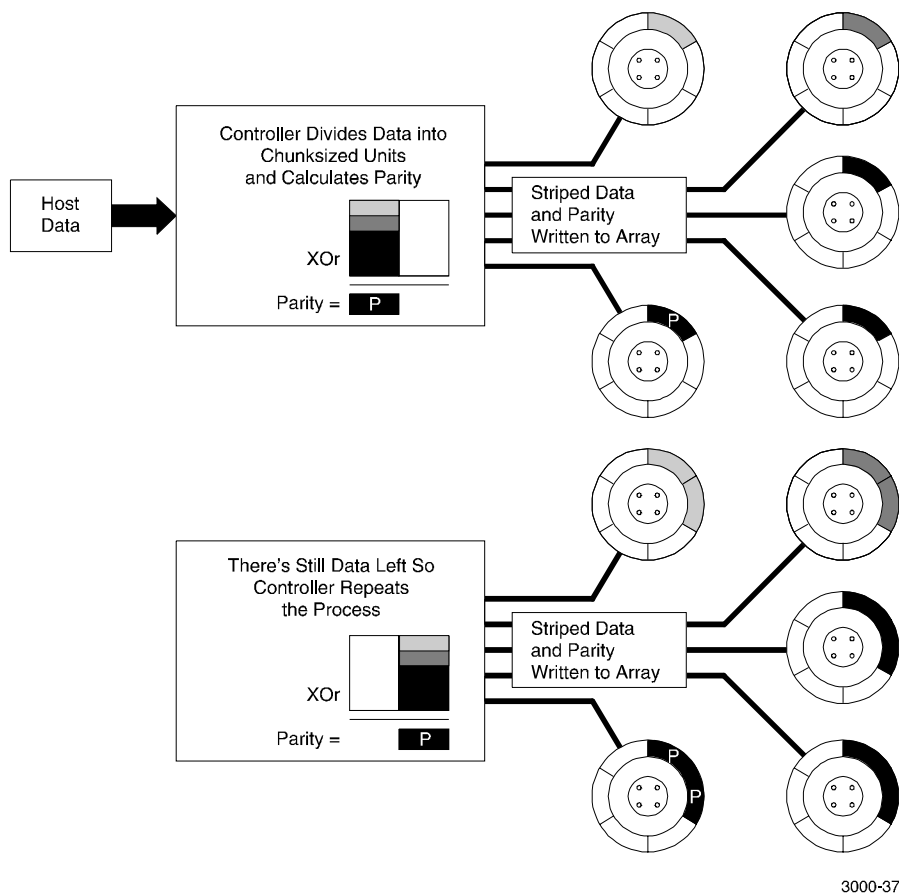
Table 2-5 RAID 0+1 Example

Drives Selected	Function
Channel 1, ID 0	First member of stripe set.
Channel 1, ID 1	Mirror of channel 1, ID 0
Channel 1, ID 2	Second member of stripe set
Channel 2, ID 0	Mirror of channel 1, ID 2
Channel 2, ID 1	Third member of stripe set
Channel 2, ID 2	Mirror of channel 2, ID 1

2.6.4 RAID 4

RAID 4 breaks up host data into chunks, calculates parity by performing an exclusive-or on the chunks, and then writes the chunks to all but one drive in the array and the parity data to the last drive. When the host requests data from the disk drives, the controller retrieves the chunks containing the addressed data, reconstitutes the data from the chunks, and passes the data to the host.

Figure 2-7 Diagram of a RAID 4 Write



In the event of a single drive failure, a RAID 4 array will continue to operate in degraded mode. If the failed drive is a data drive, writes will continue as normal, except no data will be written to the failed drive. Reads will reconstruct the data on the failed drive by performing an exclusive-or operation on the remaining data in the stripe and the parity for that stripe. If the failed drive is a parity drive, writes will occur as normal except no parity will be written. Reads will simply retrieve data from the data disks. There will be no deterioration in controller performance while a RAID set is in degraded mode.

In general, RAID 4 is best suited for applications such as graphics, imaging, or video that call for reading and writing large, sequential blocks of data. However, you may find that RAID 4 is preferable to RAID 5 even for applications characterized by many small I/O operations, such as transaction processing. This is due to the controller's intelligent caching, which efficiently handles small I/O reads and writes, and to the relatively less complex algorithms needed to implement RAID 4.

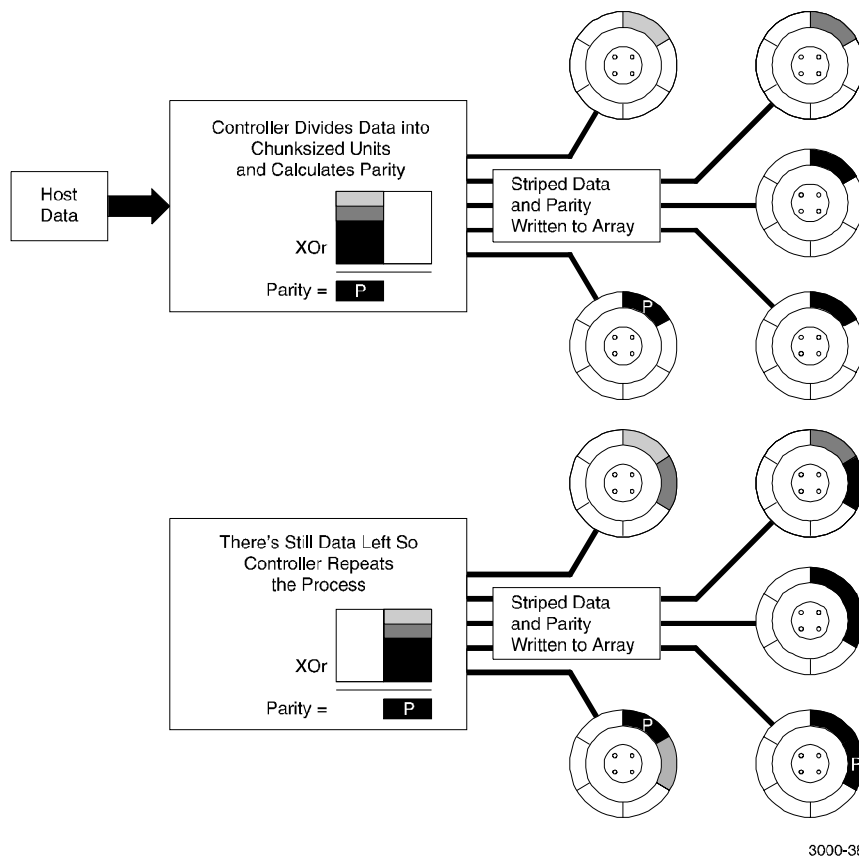
The benefits of RAID 4 disappear when you have many, small I/O operations scattered randomly and widely across the disks in the array. RAID 4's fixed parity disk becomes a bottleneck in such applications, as the following example illustrates. Let's say the host instructs the controller to make two small writes. The writes are widely scattered, involving two different stripes and different disk drives. Ideally, you would like both writes to take place at the same time, but RAID 4 makes this impossible, since the writes must take turns accessing the fixed parity drive. For this reason, RAID 5 is the better choice for widely scattered, small write operations.

CAUTION

RAID 4 can withstand a single failure and handle I/O activity without interruption in degraded mode until the failed drive is rebuilt. If a second drive fails while the RAID set is in degraded mode, the entire RAID set will fail.

2.6.5 RAID 5

RAID 5 addresses the bottleneck issue for barrages of widely scattered, small I/O operations. Like RAID 4, RAID 5 breaks up data into chunks, calculates parity, and then writes the chunks in stripes to the disk drives, saving one drive on each stripe for the parity data. Unlike RAID 4, however, RAID 5 changes the parity drive on each stripe. This means, for instance, that a write operation involving drive 2 on stripe 1 can conceivably take place at the same time as a write involving drive 3 on stripe 2, since they would be addressing different parity drives.

Figure 2–8 Diagram of a RAID 5 Write

RAID 5 handles drive failures in the same manner as RAID 4, except the parity is different for each stripe. The controller either uses the parity information on a stripe to reconstruct its data or simply reads the data as normal, depending on the location of the stripe's parity drive.

While RAID 5 is ideally suited for applications with many, small I/O operations, there is no reason why it cannot function equally well for applications with large, sequential I/O operations. This makes RAID 5 an excellent all-purpose RAID level.

CAUTION

RAID 5 can withstand a single failure and handle I/O activity without interruption in degraded mode until the failed drive is rebuilt. If a second drive fails while the RAID set is in degraded mode, the entire RAID set will fail.

2.6.6 JBOD

JBOD, which stands for “Just a Bunch of Disks”, makes it possible to connect one or standalone disk drives to the controller. A JBOD disk drive is not part of a redundancy group, even though the controller assigns a redundancy group number to the drive. This number becomes that logical unit number (LUN) that the host will use to address the drive.

One use for JBOD is to connect a system disk drive to the controller. The drive does not become part of a RAID set, but it is made available to the host on the same SCSI bus as the other devices controlled by the controller.

2.7 System Parameters

The system parameters (see Table 2–6) allow the user to customize certain aspects of the controller via the StorageWorks Command Console (SWCC) Graphical User Interface.

Table 2-6 System Parameters

Parameter	Description
Password Checking	Enables or disables password checking. When password checking is enabled, the controller will limit access to certain options unless the correct password is supplied.
Rebuild Rate	Determines how much of the controller's processing power is to be used during rebuild operations. A higher number speeds up the rebuild operation, but may impact I/O performance. A lower number will not impact performance, but will require longer to rebuild the RAID set.
Create Rate	Similar to Rebuild Rate, but determines the amount of processing power to be used while creating a RAID set.
UPS Installed	Determines whether the controller should enable the UPS warning signals (AC Lost and Two-Minute Warning). See "Backup Power Management" for more details.
Read-Ahead Enabled	Determines whether the controller should automatically pre-stage data into cache during read operations. Enabling read ahead improves sequential access times as it improves the chance of getting a cache hit on a read. However, it is not well suited for random access patterns since it relies on spatial locality.
Read-Ahead Limit	Limits the amount of cache blocks that will be read ahead when read ahead is enabled. This limit prevents the cache from being dominated by read-ahead data.
Validation Delay	Determines how long the controller will delay after spinning up drives before scanning them. The larger the number, the more time the drives have to initialize themselves.
12-Volt Sensor Enabled	Determines whether the controller should monitor the 12-volt input line.
Host Select WCE	Specifies whether the controller allows the host to set and/or clear the Write Cache Enable bit in the caching mode page. Prevents host systems from turning write caching off and degrading performance. See "Write Operations" for more details.
Drive Command Time-out	The number of seconds the controller will wait for a drive to re-select after disconnecting during the processing of a command. If a drive exceeds this time, the command will be retried. This value also indicates the maximum amount of time a drive can be on the SCSI bus. If a drive is on the bus for a greater length of time, the drive will be reset and the command will be retried.
Reset Propagation	Specifies whether the host channels should propagate resets on other host channels. This is useful in clustering environments where one cluster may need to silence the adapter of a failing peer cluster.

2.8 Redundant Operation

When operating in a redundant configuration, the two controllers are linked such that, in case of a failure, the surviving controller can access the other controller's cache memory and complete all operations that were in progress when the failure occurred. The controllers support two different configurations:

- **ACTIVE / ACTIVE:** One host port is active on each controller. The other port on each controller is passive and only used if the peer controller fails.
- **ACTIVE / PASSIVE:** Both host ports on one controller are active. The other controller's ports are both passive and only used if the primary controller fails.

When one controller fails, the survivor will process all I/O requests until the failed controller is repaired and powered on. The subsystem will then return to its previous state (i.e., ACTIVE / ACTIVE or ACTIVE / PASSIVE).

2.8.1 Initialization

During initialization, the firmware in the RAID 3000 verifies that both controllers have consistent configurations including identical memory cache and system parameters. If the controller setups are incompatible, the set is not bound and each controller operates in stand-alone mode.

2.8.2 Message Passing

Information is shared between the two controllers by a collection of messages passed through the backplane connectors. The messages provide configuration data as well as a heartbeat which is transmitted by each controller every 500 ms. If a controller does not receive a heartbeat within one second, it assumes the peer controller has become inoperable and begins failing over.

If the controllers cannot exchange messages due to communication problems over the backplane, they will break the connection and each controller will switch to a stand-alone mode.

2.8.3 Failover

Failover describes the process of transferring data from a failed controller to a survivor and completing any active tasks. When one controller begins the fail-over process, it sends a reset to the other controller, which prevents the failing unit from processing any more information and enables any host ports that are passive. It then downloads the failed controller's cache to its unused portion of cache and begins acting upon that data.

While downloading the data, the controller responds to I/O by disconnecting (if allowed) and waiting approximately three seconds before reconnecting and presenting a BUSY status. The delay is to prevent host operating systems from seeing too many errors and fencing off the controller.

2.9 Environmental

The controller incorporates a set of on board sensors to detect abnormal operating conditions that may affect data safety.

2.9.1 Backup Power Management

The controller must be connected to a backup battery and/or Uninterruptable power supply (UPS) to prevent the subsystem cache from being corrupted during unexpected losses of power. If no backup power supply is provided, the controller will remain off-line and reject all I/O requests with a status of Check Condition/Hardware Error.

If the backup power source indicates that power may be failing, the controller's first step is to sound an alarm, enter write through mode and begin flushing cache. If the backup power source reaches a critical state, the alarm frequency increases and all host channels are disabled to prevent new requests from interfering with the cache flush.

Table 2–7 shows how the controller reacts to the power supplies.

Table 2–7 Backup Power Management

Battery Status	UPS Status			
	OFFLINE (Not Installed)	NORMAL	WARNING (AC Lost)	SEVERE (2-Min. Warning)
OFFLINE (Not Installed)	OFFLINE	NORMAL	SEVERE	OFFLINE
NORMAL	NORMAL	NORMAL	WARNING	WARNING
WARNING	SEVERE	WARNING	WARNING	SEVERE
SEVERE	OFFLINE	WARNING	SEVERE	OFFLINE

NORMAL

- No Alarms
- Normal I/O

WARNING

- Slow Alarm
- Normal I/O

SEVERE

- UltraSCSI Alarm
- Flushes Cache
- Enters Write-Through Mode

OFFLINE

- UltraSCSI Alarm
- Flushes Cache
- No New I/O Allowed

2.9.2 Voltage Monitoring

2.9.2.1 System Voltage

The controller monitors the incoming system voltage levels and ensures they are satisfactory for controller operation. The acceptable voltage levels are listed in Table 2–8.

Table 2–8 Acceptable System Voltage Levels

State	Range	Action
Normal	4.80 - 5.25 V	Normal Operation
Low Warning	4.75 - 4.80 V	Alarm
High Warning	5.25 - 5.30 V	Alarm
Low Severe	Less than 4.75 V	Off-line
High Severe	Greater than 5.30 V	Off-line

2.9.2.2 Termination Voltage

The controller monitors the incoming termination voltage levels and ensures they are satisfactory for controller operation. The acceptable voltage levels are shown in Table 2–9.

Table 2–9 Acceptable Termination Voltage Levels

State	Range	Action
Normal	4.20 - 5.40 V	Normal Operation
Low Warning	4.00 - 4.20 V	Alarm
High Warning	5.40 - 5.55 V	Alarm
Low Severe	Less than 4.00 V	Off-line
High Severe	Greater than 5.55 V	Off-line

2.9.2.3 12 V Supply Voltage

The controller monitors the incoming 12-volt levels and ensures they are satisfactory for operation. The acceptable voltage levels are shown in Table 2–10.

Table 2–10 Acceptable 12 Volt Levels

State	Range	Action
Normal	10.80 - 13.80 V	Normal Operation
Low Warning	10.20 - 10.80 V	Alarm
High Warning	13.80 - 14.40 V	Alarm
Low Severe	Less than 10.20 V	Off-line
High Severe	Greater than 14.40 V	Off-line

2.9.3 Temperature Monitoring

2.9.3.1 External Temperature

The controller monitors the external operating temperature and ensures they are satisfactory for controller operation. The acceptable temperature levels are shown in the following table.

Table 2–11 Acceptable External Temperature Voltage Levels

State	Temperature	Action
Normal	< 41° C	Normal Operation
High Warning	< 41° C - 52° C	Alarm
High Severe	> 52° C	Off-line

2.9.3.2 Board Temperature

The controller monitors the on-board operating temperature and ensures they are satisfactory for controller operation. The acceptable temperature levels are shown in the following table.

Table 2–12 Acceptable Board Temperature Voltage Levels

State	Temperature	Action
Normal	< 52° C	Normal Operation
High Warning	52° C - 56° C	Alarm
High Severe	> 56° C	Off-line

3

Maintenance

This chapter describes how to interpret the status of the LEDs on the pedestal and use them as a troubleshooting aid during a pedestal malfunction. Both the pedestal LEDs and the LEDs on the major components are covered. The chapter also describes how to replace a Field Replaceable Unit (FRU) and how to reconfigure the SCSI bus.

3.1 Introduction

Troubleshooting the pedestal consists of monitoring the status of the external LEDs to determine if a major component is malfunctioning. The Field Replaceable Units (FRUs) in the pedestal are:

- Disk drives (SBBs)
- RAID array controller
- Power supplies
- Environmental Monitor Unit (EMU) circuit board
- Uninterruptable Power Supply (UPS)

The information in this chapter is divided into the following major sections:

- Pedestal status and power supply LEDs
- SBB status LEDs
- Controller LEDs
- EMU error reporting
- Replacing an FRU
- Reconfiguring the UltraSCSI bus
- Replacing the controller memory cache modules


3.2 Pedestal Status and Power LEDs

The pedestal is equipped with two front panel LEDs (see Figure 3–1) that monitor the following error conditions:

- A power supply fan that is not operating
- An over-temperature condition
- A dc power problem
- External fault conditions
- Controller faults

When the pedestal is operating properly, the green *power* LED is on and the amber pedestal *fault* LED is off. Also, the green power supply LED on the rear of each power supply is on.

Figure 3-1 Pedestal Status LEDs




3.3 SBB Status LEDs

Each disk drive SBB in the pedestal has two status LEDs – a green *activity* LED and an amber *fault* LED as shown in Figure 3–2. When the pedestal is operating properly, the activity LED is flashing indicating normal disk activity on the SCSI bus, and the *fault* LED is off. Table 3–1 summarizes the states of the SBB LEDs and recommends corrective actions when a LED indicates a fault condition.

Table 3–1 Disk Drive SBB Status LEDs

Activity LED	Fault LED	Indication
On	Off	Drive is operating properly.
Off	Off	Drive is inactive and operating normally. There is no fault.
On	On	Fault status: drive is defective. Recommend that you replace the device.
Off	On	Fault status: drive is inactive and not spinning. Recommend that you replace the device.
On	Flashing	Fault status: drive is active and slowing down due to the fault.

Figure 3–2 Disk Drive Status LEDs

ESP-15

3.4 Controller LEDs

The LEDs on the front panel of the controller monitor host and disk channel activity and a controller fault condition. The reset button/LED flashes green approximately once every second (heartbeat) to indicate that the controller is operating normally. Figure 2–3, Chapter 2, identifies the LEDs. Table 2–2 describes their functions.

3.5 EMU Error Reporting

The primary function of the EMU is to detect and report conditions that can cause the pedestal to malfunction and to report malfunctions. To accomplish this the EMU constantly monitors the following pedestal signals:

- +5 and +12 V dc
- ac input
- Power Supply present
- Total power
- Power OK (P_OK)
- Power Disabled
- Fan Speed (minimum and high speed)
- Fan exhaust temperature
- Shelf OK (S_OK)

The EMU can exchange signals with auxiliary devices and controllers.

3.5.1 EMU Error Conditions

The EMU reports error conditions and malfunctions using an audible alarm and a LED. The user-enabled audible alarm and the amber fault LED on the front of the pedestal are the only error indicators.

Whenever any of the following error conditions occur, the amber fault LED on the front of the pedestal will turn on. When alarm switch S1 on the rear panel of the pedestal is in the up (enabled) position, the audible alarm will sound whenever one of the following conditions occurs:

- Loss of ac power to one of the power supplies
- Failure of either power supply fan
- UPS not connected
- UPS power failure
- UPS output too low
- Temperature exceeds 50° C (123° F)
- One of the +12 V dc outputs is less than + 9.85 V dc
- One of the +5 V dc outputs is less than + 3.95 V dc
- Miscellaneous error condition
- A controller error condition exists
- External Fault

3.6 Replacing Components (FRU's)

This section describes how to replace an FRU in the RAID Array 3000 pedestal. The information is divided into the following subsections:


- Removing the pedestal door
- Replacing an SBB
- Replacing a power supply
- Replacing the controller
- Replacing the EMU board
- Replacing the UPS
- Replacing a controller memory cache module

3.6.1 Removing the Pedestal Door

Proceed as follows to remove the pedestal door (Figure 3–3):

1. Unlock and open the door to a 90° angle in relation to the closed position.
2. Carefully lift up on the door until the hinge pins are against the top of the mounting holes.
3. Pull the door straight out until the hinge pins clear the bezel.

Figure 3–3 Removing Pedestal Door



5000-04

3.6.2 Replacing an SBB

There are two methods for replacing a disk drive SBB with an *identical* SBB:


- **Hot swap** – This method requires that the SCSI controller support removing and installing SBBs while the bus is active. Hot swap is supported by the RAID Array 3000 controller
- **Cold swap** – Requires removing ac power from the pedestal and disabling the UPS to deactivate the bus

Perform the following procedure to replace an SBB in the pedestal (see Figure 3–4):

CAUTION
<ul style="list-style-type: none">• Always use both hands when removing or inserting an SBB.• Never remove a drive from the pedestal before it has completely spun down.• Take care not to slam the replacement SBB into the pedestal enclosure.

1. Unlock and open the front door.
2. Release the drive from the pedestal slot by squeezing the mounting tabs on the SBB together, but do not remove the SBB from the pedestal.
3. Wait 15-30 seconds for the drive to stop spinning before removing it from the pedestal slot.
4. Insert an identical model SBB in the slot and push it in until an audible click is heard indicating the SBB is fully seated (SBB front panel mounting tabs expand and engage the pedestal shelf).
5. Observe that the SBB status LEDs are operating.
6. Close the front door.

Figure 3–4 Replacing an SBB



3000-06

3.6.3 Replacing a Power Supply


You can replace a power supply without affecting pedestal operation using the following procedure:

CAUTION

When you remove a power supply, the airflow through the SBBs is interrupted. Always install the replacement power supply as quickly as possible to prevent overheating.

1. Grasp the power supply handle shown in Figure 3–5.
2. Slide and hold the locking latch to the left and pull the supply out using a short jerking motion.
3. Insert the replacement power supply into the pedestal and carefully align it.
4. Push the power supply in until it is fully seated and the locking tab engages.
5. Until the fan is operating at the proper speed, the green power supply LED remains off. When the fan reaches the proper speed (several seconds), the LED should come on.

Figure 3–5 Replacing a Power Supply




3.6.4 Replacing the RAID Array Controller

Perform the following procedure to replace the RAID Array controller (see Figure 3–6):

1. Grip the two locking latches on the front panel of the controller and pull them forward until the controller disengages from its mounting slot.
2. Remove the controller by sliding it forward and free of the pedestal.
3. Insert the replacement controller into the open slot, align the module into the card guides, and gently slide it into the pedestal until the connector engages the backplane connector in the pedestal.
4. Turn the two front panel latches inward to fully seat the controller in the pedestal slot.

Figure 3–6 Removing the Controller from the Pedestal



3000-14

3.6.5 Replacing the EMU Board

WARNING

Only qualified service personnel should replace the EMU. Dangerous voltages are exposed when the pedestal side panel is removed. Always power off the pedestal and remove the power cords before replacing the EMU.

Perform the following procedure to replace the EMU Board:

Tools required:

- Flat-blade or 3/16 (5 mm) hex-head screwdriver
- # 10 TORX-head screwdriver


1. Quiesce the host bus by shutting down the host system.

CAUTION

To allow the UPS to supply power while the cache is being flushed to disk, do not unplug the base pedestal from the UPS.

2. Unplug the base pedestal power cord from the wall outlet.
3. Unplug the UPS power cord from the wall outlet, the UPS will now signal the controller to flush the cache.
4. Wait until the UPS shuts down completely (this may take several minutes).
5. Power off the base pedestal and plug the UPS power cord into the wall outlet.
6. Remove the pedestal door.
7. Remove the left side cover from the pedestal (see Section 3.7.1).
8. Record the orientation of the color traces on the two ribbon cables connected at the rear of the EMU board and disconnect the cables.
9. Remove the mounting nuts from the UPS and external fault connectors on the rear panel using a hex-head or a flat-blade screwdriver (see Figure 3–7). Remove the board from the pedestal.

Figure 3-7 Remove Nuts from UPS and External Fault Connectors



3000-49

10. Align the connectors on the replacement EMU with the connector holes on the rear panel and replace the mounting studs to secure the board (*do not over-tighten*).
11. Reconnect the two ribbon cables to the rear connectors on the EMU board.
12. Replace the side cover on the pedestal and secure it with the TORX-head screw on the front bezel of the pedestal.
13. Replace the door on the pedestal.

3.6.6 Replacing the UPS

Proceed as follows to replace the UPS:

1. Ensure the UPS power switch is set to off.
2. Disconnect the pedestal power cord from the UPS.
3. Disconnect the UPS power cord from the wall outlet.
4. Disconnect the signal control cable from the UPS
5. Install the replacement UPS and reconnect the power cords and the signal control cable.
6. Set the UPS power switch to on and then power up the pedestal.

3.7 Differential/Wide UltraSCSI Bus

The reliability of data transfers on a SCSI bus depends on the following factors:

- The bus transmission rate
- The maximum SCSI bus length (a function of bus type, the transmission rate, and the use of SCSI bus converters)
- The total length of the SCSI bus (as measured from the host bus terminator to the subsystem terminator)
- MT/s (mega-transfers per second) is the repetitive rate at which words of data are transferred across a bus. The number of megabytes per second (MB/s) is determined by the bus width (8 or 16-bit) and the number of bytes per word (1 or 2, respectively)

Table 3–2 lists the maximum SCSI bus lengths and the longest SCSI cables recommended by DIGITAL.

NOTE

Because the bus length includes the cable plus the backplane distance inside the enclosure, the recommended maximum cable lengths listed *are not* the same as the maximum bus length.

Table 3–2 SCSI Bus Length and External Cables

Bus Speed	Rate		Bus Length		Longest DIGITAL Cable		
	MT/s	MB/s	Meters	Feet	Number	Meters	Feet
Fast	20	40	25	82	BN21K-23 BN21L-23	23	82

The SCSI bus in your pedestal is factory-configured as a split bus. One bus is designated bus **D0** and the other as bus **D1**. An internal eight-position step switch in the cabinet controls the SCSI bus device address configuration of the SBB slots in the pedestal. For the RAID Array 3000, the switch is preset at the factory to configuration “6”. This results in a device slot address assignment of 8 through 11 for bus **D0** and 8 through 10 for bus **D1**. Table 3–3 shows a listing of the device addresses for each bus and their corresponding pedestal slot location.

Table 3–3 Assigned Slot Device Addresses in the Pedestal

Slot #	0	1	2	3	4	5	6
Bus	D0				D1		
Device Address	8	9	10	11	8	9	10

3.7.1 Reconfiguring the SCSI Bus

WARNING

Only qualified service personnel should reconfigure the SCSI bus. Dangerous voltages are present within the subsystem. To prevent electrical shock, always turn the subsystem off and disconnect the power cords before removing the side panel.

If you want to reconfigure the SCSI bus (to add the Expansion Pedestal Option to your subsystem installation for example) you must reconfigure the bus. This involves powering down the subsystem installation, removing the left side panel, and reconfiguring the bus by changing the setting of the configuration switch. Adding the Expansion Pedestal Option to your installation is described in Chapter 4 of this guide.

The SCSI bus configuration switch selects the eight (0 – 7) SCSI bus configurations in the pedestal. Each bus configuration determines the slot device addresses (0 – 6) for both 8-bit and 16-bit devices. The pedestal is configured at the factory for split-bus operation (configuration switch set to “6”). To change the configuration of the bus, you must remove the left side panel to gain access to the backplane and the configuration switch.

NOTE

SCSI device addresses 6 and 7 are reserved for the RAID Array controllers. The top controller slot in the pedestal is assigned device address 7 and the bottom slot (redundant controller) is device address 6. The controller allows a maximum of 14 disk devices on each of its device buses.

To remove the side panel, proceed as follows:


1. Quiesce the host bus by shutting down the host system.

CAUTION

To allow the UPS to supply power while the cache is being flushed to disk, do not unplug the base pedestal from the UPS.

2. Unplug the base pedestal power cord from the wall outlet.
3. Unplug the UPS power cord from the wall outlet, the UPS will now signal the controller to flush the cache.
4. Wait until the UPS shuts down completely (this may take several minutes).
5. Power off the base pedestal and plug the UPS power cord into the wall outlet.
6. Disconnect the SCSI cable from the host connector on the rear panel.
7. Remove the front door on the pedestal.
8. Remove the cover screw from the left side panel (see Figure 3–8).
9. Grasp the handhold at the rear of the panel and pull the panel toward the rear and clear of the cabinet.

Figure 3–8 Remove Screw and Panel



10. Figure 3–9 shows the location of the switch. Figure 3–10 depicts the separate step switches to set the subsystem to the desired SCSI bus configuration (refer to the label at bottom of pedestal to cross reference slot SCSI ID addresses with switch settings).

NOTE

To decrease the configuration number: Press the upper switch to step the address down one address at a time (decrement) until the desired configuration number is displayed.

To increase the configuration number: Press the lower switch to step the address up one address at a time (increment) until the desired configuration number is displayed.

11. Close the side panel and replace the cover screw.

Figure 3–9 Location of SCSI Bus Configuration Switch





Figure 3–10 Configuration Switch



NOTE

Reconfiguring the bus also involves cable and bus terminator changes. Chapter 4 of this guide describes in detail how to reconfigure the bus.

3.8 Replacing the Controller Memory Cache Modules


The two memory cache modules in the RAID controller are replaced by removing the controller from the pedestal to gain access to the modules. Then, place the controller on a flat working surface and proceed as follows:

CAUTION

To prevent an electrical discharge from damaging the SIMMs, always wear an ESD wrist strap connected to a suitable ground when handling the memory chips.

1. Power down the pedestal first and then the UPS.
2. Grasp the latches on the front of the controller and pull them forward until the controller disengages from the pedestal (see Figure 3–11).
3. Remove the controller from the pedestal and place on a flat working surface.
4. Remove the two installed cache memory SIMM modules by releasing the locking clips at each end of the module until it disengages and snaps into an upright position as shown in Figures 3–12 and 3–13.

Figure 3–11 Remove Controller



3000-14

Figure 3–12 Release Locking Clips





Figure 3–13 Remove Installed SIMM Modules



CAUTION

Ensure the “side 1” side of the two replacement SIMMs is facing toward you when installing the modules in the following step.

5. Installed the two replacement memory modules by aligning the module and connector pins (check alignment guide in center of module) and gently pivot the module the main controller board until it snaps into place (see Figures 3–14 and 3–15).
6. Replace the controller into the pedestal.
7. Power on the UPS and pedestal and check the activity LEDs on the front panel of the controller. The reset switch/LED should begin to flash at a half-second rate (heartbeat) and the host activity LEDs should flash.

Figure 3–14 Install Replacement Modules





Figure 3–15 Pivot Module Down to Secure




Expansion Pedestal Option

This chapter describes the major features of the Expansion Pedestal Option and how to connect the option to your RAID Array 3000-storage subsystem.

4.1 Product Description

The StorageWorks pedestal expansion option (Figure 4-1) is designed to expand the storage capacity of the RAID Array 3000 subsystem. When connected to the RAID Array 3000 base pedestal, the expansion option contains the basic components required to create a dual-pedestal storage array with a 16-bit, single-ended Ultra-SCSI bus. The option enables a user to add up to seven 3½-inch SBBs to create a 14-device storage array. The expanded array is controlled and operated in an identical fashion as the base RAID Array 3000 subsystem.

Figure 4-1 Expansion Pedestal



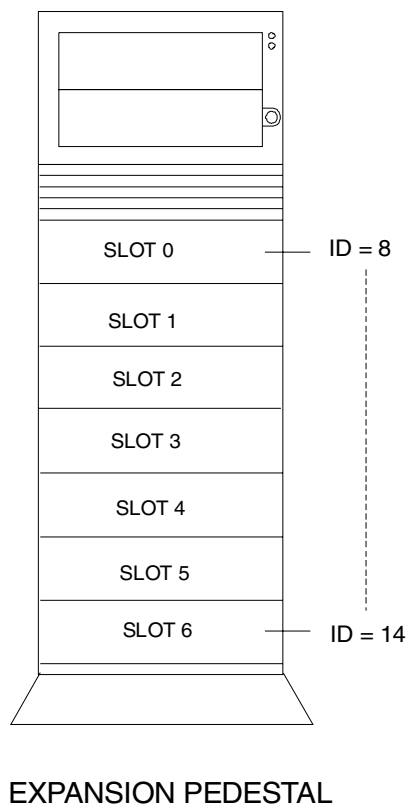
310EXP-02

4.2 Expansion Pedestal Cabinet

The expansion pedestal cabinet is a modular free-standing storage cabinet that is completely self contained with dual fan-cooled power supplies, an internal UltraSCSI single-ended extender module, and an internal EMU circuit board. The cabinet dimensions are the same as the subsystem base pedestal which houses the controller and is normally installed within one meter of the base cabinet to facilitate the cable connections between the two units. Figure 4–2 identifies the expanded subsystem's SBB slots and their corresponding SCSI ID addresses. Figure 4–3 identifies the items on the rear panel power supplies. The characteristics of the expansion pedestal cabinet are outlined below.


- The storage device capacity of the expansion pedestal is seven 3½-inch SBBs
- The pedestal slots are numbered 0 through 6 from top to bottom
- There are seven SCSI bus device addresses (target IDs) 8 through 14 which can be assigned to the 3½-inch SBBs
- There is a single 68-pin, VHDC female SCSI connector on the rear panel which interconnects the SBB expansion pedestal to the controller pedestal
- The rear panel contains an alarm switch and an external fault condition connector
- The expansion pedestal contains two interchangeable fan-cooled ac power supplies for redundant power
- The expansion pedestal is equipped with an internal configuration switch which selects one of the eight (0 through 7) subsystem configurations (set to the correct configuration setting at the factory to properly integrate the expansion pedestal to the controller pedestal)

Figure 4-2 Expansion Pedestal Slot Locations and ID Addresses



3000-19A

Figure 4-3 Rear Panel Power Supplies



3000-22

4.3 Expansion Pedestal Components

The expansion pedestal contains a 16-bit, wide/differential UltraSCSI bus, an Environmental Monitor Unit (EMU), a SCSI bus extender module, and two universal 50/60 Hz, 100 – 240 Vac fan-cooled power supplies.


The single-ended UltraSCSI bus is factory-configured as one continuous bus that runs along the backplane between the disk drive connectors and the internal cables. These cables connect the drives to the connectors located on the rear panel of the expansion pedestal. The device addresses on the bus are set at the factory by an internal configuration switch. When set to a specific position, the switch controls the addresses of each SBB slot.

The SCSI bus extender module extends the allowable electrical length of the bus to accommodate longer physical SCSI cable connections between the base and expansion pedestals.

The EMU (Figure 4-4) is an internal circuit board, which monitors the operation of the pedestal. The EMU monitors power supply voltages, fans, temperatures that are reported to the user, and controls (turns on and off) the audible alarm and status LED on the front panel of the pedestal. It is connected to the SCSI bus and powered by internal cabling. The following external components on the rear panel of the expansion pedestal are part of the EMU board:

- The alarm switch (S1) that enables (up) or disables (down) the audible alarm
- The External Fault Condition connector allows the EMU to monitor the status of a user-selected device

Figure 4-4 EMU Circuit Board Location



3000-37

4.4 Reconfiguring Base Pedestal UltraSCSI Bus

WARNING

Only qualified service personnel should reconfigure the base pedestal. Dangerous voltages are present within the subsystem. To prevent electrical shock, always turn the subsystem off and disconnect the power cords before removing the side panel.

The RAID Array 3000 base pedestal is factory-configured for split-bus operation. You must reconfigure the bus in the base subsystem from split-bus to a “through-bus” configuration prior to connecting the expansion cabinet to the base subsystem. The components needed to reconfigure the split-bus in the base subsystem are included with your pedestal expansion kit option.

CAUTION

To prevent ElectroStatic discharge (ESD) from damaging the controller, always wear an ESD wrist strap connected to a suitable ground whenever you handle any of the electronic components.

Perform the following procedure to reconfigure the SCSI bus in the base pedestal from a “split-bus” to a “through-bus” configuration:


1. Perform an inventory of the bus conversion items supplied with the pedestal expansion kit option. The items should consist of:
 - SCSI bus jumper cable 17-04166-03
 - SCSI cable 17-04454-01
2. Quiesce the host bus by shutting down the host system.

CAUTION

To allow the UPS to supply power while the cache is being flushed to disk, do not unplug the base pedestal from the UPS.

3. Unplug the base power cord from the wall outlet.
4. Unplug the UPS power cord from the wall outlet, the UPS will now signal the controller to flush the cache.
5. Wait until the UPS shuts down completely (this may take several minutes).
6. Power off the base pedestal and plug the UPS into the wall outlet.
7. Remove the side cover from the base pedestal (Figure 4–5).


Figure 4-5 Remove Side Cover from Base Pedestal



3000-06

8. Remove the bus terminator from backplane connector J11 (❶, Figure 4-6).


Figure 4-6 Remove SCSI Bus Terminator



SHR-1495

9. Remove the SCSI cable from device # 1 connector and backplane connector J16 (❶, Figure 4-7).


Figure 4-7 Disconnect SCSI Cable



SHR-1495

10. Connect jumper cable 17-04166-03 between the backplane connector J11 and the backplane connector J16 (see Figure 4-8).


Figure 4-8 Connect SCSI Jumper



3000-42

11. Remove the “knock-out” plate located above the **D1 OUT** label on the rear panel of the base pedestal (see Figure 4–9).


Figure 4–9 Remove Connector Knockout Plate



3000-45

12. Connect cable assembly 17-04454-01 between the **D1 OUT** bulkhead opening and the device # 1 backplane connector (see Figure 4–10). Secure the bulkhead connector by tightening the two 6-32 SEM screws.

Figure 4–10 Connect SCSI Cable



13. Set the bus configuration switch to "7" (see Figures 4–11 and 4–12). Figure 4–13 shows the reconfigured SCSI bus addresses of the expanded subsystem.

Figure 4–11 Set Configuration Switch to 7




Figure 4-12 Configuration Switch






Figure 4-13 Reconfigured SCSI Bus Addresses



14. Reinstall the side panel on the base pedestal.
15. Connect the HD68-to-HD68 SCSI cable from the **D1 OUT** connector on the base pedestal to the **D1 IN** connector on the rear of the expansion pedestal as shown in Figure 4-14.


Figure 4-14 Connect SCSI Cable Between Pedestals



3000-47

16. Make the power cable connections between the expansion pedestal, the UPS, and the ac power source (see Figure 4-15).

Figure 4–15 Power Cable Connections




CAUTION

If you want to maintain the existing RAID level configuration, three drives must be relocated from the base pedestal to specific slots in the expansion pedestal as described in step 13.

17. Transfer the bottom three disk drives (slots 4, 5, and 6) from the base pedestal to the top three slot locations (slots 0, 1, and 2) in the expansion pedestal as shown in Figure 4–16.

Figure 4-16 Transfer Drives from Base to Expansion Pedestal



18. Install the new drives in the remaining slots each pedestal to complete the installation.
19. Power up the UPS and the two pedestals and then proceed to StorageWorks *RAID 3000 Configuration and Maintenance Guide* to configure the expanded subsystem.

5

Second Controller Option

This chapter describes how to install a second RAID controller in the RAID Array 3000 pedestal. The second controller option adds a fail/safe feature to your storage subsystem. The chapter also contains a procedure describing how to configure the subsystem for dual-controller operation when only one serial port is available on the host.

5.1 Introduction

The controller option adds a second (redundant) controller unit to your subsystem to preserve the integrity of data should the first controller malfunction. The second controller is installed directly below the existing controller in the bottom controller slot of the RAID Array 3000 pedestal (see Figure 5-1).

The installation procedure consists of adding two SIMM memory modules to the redundant controller and, depending on the option kit you are installing (see inventory below), one or two SIMM modules to the existing controller. Following the memory upgrade, you must upgrade the firmware in both controllers, seat the devices in their respective controller slots and configure the subsystem to accommodate a redundant controller.

Perform an inventory of the items contained in the second controller kit option. The kit should contain the following:

RAID Array 3000 Controller

For option kit DS-HSZ22-AA, 3 SIMMs

For option kits DS-HSZ22-AB, 4 SIMMs


CD with new Firmware

Model Label

Warranty Card

This Guide

Figure 5-1 Second Controller Pedestal Slot Location



3000-48

5.2 Installing the Upgrade

CAUTION

To prevent an electrical discharge from damaging the SIMMs, always wear an ESD wrist or foot strap connected to a suitable ground when handling the memory modules.

NOTE

You can upgrade your firmware using the SCSI or network connection methods. These methods provide the fastest way to upgrade your firmware. To update your controller's firmware, proceed as follows.

5.2.1 Save Existing Configuration

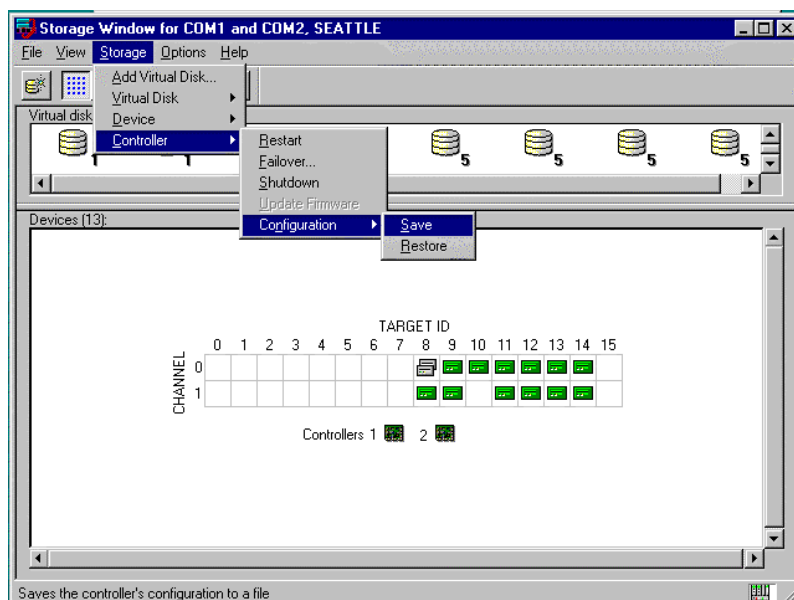
NOTE

You must save the subsystem configuration using the StorageWorks Command Console Client (SWCC) before performing the installation procedure. Otherwise, the RAID controller may lose your configuration.

To save your current configuration:

Select the **Storage** pull-down menu from the Toolbar (see Figure 5-2), click on **Controller**, select **Configuration**, and then **Save**.


Figure 5-2. Saving the Existing Configuration



A saved configuration screen with a “Save to File” field appears as shown in Figure 5-3. The example in Figure 5-3 shows the file name as `c:\config1`.

Enter your file name in the “Save to File” field and click on **Save**.

Figure 5-3. Saved Configuration



5.2.2 Update Firmware


CAUTION

If the systems disk is on the RA3000, firmware cannot be upgraded with SWCC. Establish a temporary system disk on a disk drive outside the RA3000 and proceed with these instructions or follow the alternative shown in Section 5.2.2.2.

5.2.2.1 Update Firmware Using SWCC


Start SWCC and choose **SCSI** or **Network Connection**. From the **Storage** menu (see Figure 5-4), select **Controller**, and then select **Update Firmware**.

Figure 5-4 Update Firmware Command



A window will appear (see Figure 5-5) asking you to specify the firmware file that you want to load. This file resides on the CD supplied with the controller kit option. You can easily identify Firmware software by its `.fdi` extension. Enter the `.fdi` file name (for example, `D:\firmware\xxx.fdi`), then click **Start Update**. Upon completion of the Firmware update, the system will automatically reboot and update the controller's firmware.

Figure 5-5. Firmware Update Dialog Box



5.2.2.2 Update Firmware Using Serial Interface

NOTE

Follow this procedure only when the host operating system is installed on the RA3000.

1. Properly shutdown your host system
2. Using either a laptop computer or a computer other than that attached to the RA3000, connect a serial cable between the COM port on that machine and the corresponding serial port (CTR1 or CTR2) for the controller on the RA3000.
3. Start a terminal emulator session. On Windows NT, we suggest the HyperTerminal emulator. Settings to be used are 9600 baud, 8 bits, No Parity, 1 stop bit, XON/XOFF.
4. To get the controller's attention, press Escape/& (escape key, shift key, number 7). The controller should respond with a banner stating DEC HSZ22 DEC Monitor Utility, followed by the firmware revision number. You will not actually be using this utility to upgrade the firmware, but knowing that it is operational makes the next step easier.

5. Reset the controller by power cycling the RA3000 subsystem. You should see a **"Flash Boot Utility..."** banner, followed by instructions to type CTRL/C to abort. Press CTRL/C to abort the load sequence. A **"FLASH Boot Utility Options"** menu should be displayed.
6. Choose menu item (2), Change serial baud rate. Select 38400. When presented with the "Please change your baud rate and press RETURN" message, do so from within HyperTerminal via the **File\Properties** menu, then choose **'Configure...** If there is no response from the utility after changing the baud rate, proceed with Step 6.
7. After changing the baud rate, you may have to close and re-enter HyperTerminal. Press the RETURN key after HyperTerminal restarts. You should see **the FLASH Boot Utility Option** menu again.
8. Select menu item **"1) Download new Firmware Image"**. Using the Transfer menu in HyperTerminal, choose **Send Text File** and send the firmware. You should see a **"Receiving code for System Version <ver>"** message, followed by a series of \ / - characters cycling at the end of the line. At 9600 baud, the download will take between 45-60 minutes. At 38400, it should take around 10-12 minutes. At completion, you'll see the **FLASH Programming complete** message, followed by the **FLASH Boot Utility Options** menu again.
9. Select item 9, **Restart Controller**. You'll be instructed to reset the baud rate back to 9600, which you'll again do from the HyperTerminal **File \Properties \Configure...** menu. As before, you'll probably have to exit and restart HyperTerminal to get any response. Press Enter when Hyperterminal restarts.
10. Your firmware should now be upgraded.


5.2.3 Shutdown RA3000

1. Ensure that both host ports are in a quiescent state (no I/O activity).
2. Shut down the Host System.
3. Issue a "shutdown" command from the SWCC Console to the pedestal controller.
4. Power **OFF** the pedestal or controller shelf as applicable.

5.2.4 Install Two SIMMs into Second Controller


1. Install two of the SIMM modules into the second controller (make sure all SIMM modules are of the same type) by aligning the connector pins and inserting the modules into the SIMM module connectors as shown in Figure 5-6.

Figure 5-6. Insert Module into SIMM Connector



2. Ensure the module is firmly seated and then gently pivot it toward the controller board until it snaps into place as shown in Figure 5-7.


Figure 5-7. Pivot Module Down to Seat



5.2.5 Replace Existing Controller

1. Replace the existing controller (see Figure 5-8 to remove) with the new controller.

Figure 5-8. Remove Controller from Top Slot



3000-14

NOTE

Ensure you install the new controller in the same slot as the existing controller removed in step 1 above. Do not leave the existing controller in the pedestal while performing the following step.

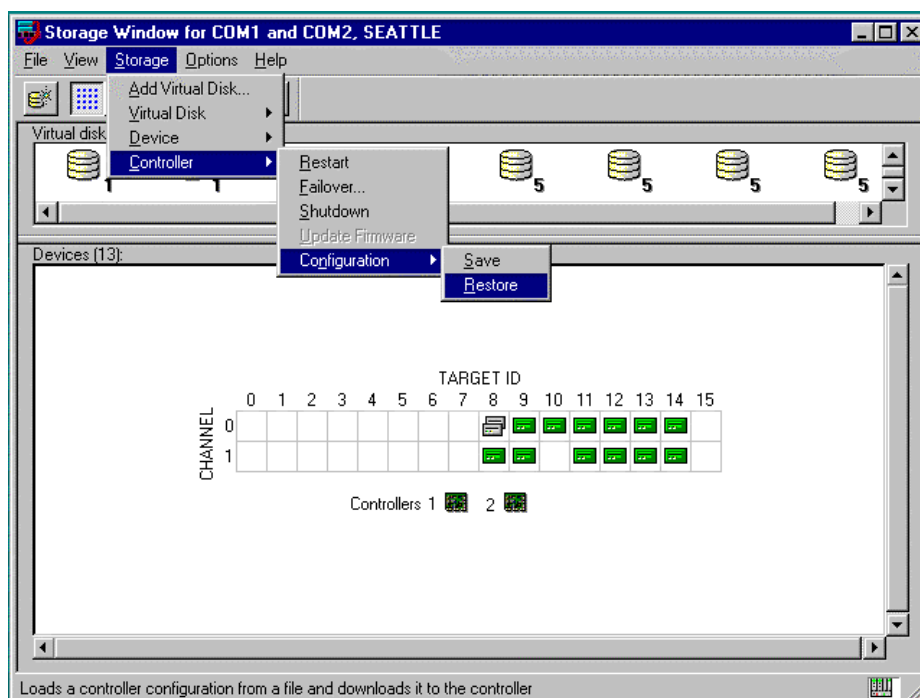
2. Power up the system and proceed to the following section to restore the configuration on the new controller.

5.2.6 Restore Configuration

To restore your configuration to the new controller:


1. Restart SWCC in the “Serial Mode” (refer to your *RA3000 Getting Started Guide* for instructions).
2. Select the **Storage** pull-down menu from the Toolbar, then choose **Controller**.
3. From the **Controller** pull-down menu (Figure 5-9), select **Configuration**, and then **Restore**.

Figure 5-9. Restoring Configuration to New Controller



The saved configuration screen with a “Restore from File” field appears as shown in the example of Figure 5-10.

Figure 5-10. Restored Configuration Example



NOTE

Restoring the configuration in the following step may take up to 5 minutes. Be patient.

4. Enter the file name that you saved in Section 5.2.1 and click on **Restore**.

5.2.7 Update Firmware on Second Controller

1. After configuration has been restored, update firmware on the second controller. Repeat the procedure in Section 5.2.2.1 or 5.2.2.2.
2. After the configuration has been restored and firmware has been updated, power down the system.

5.2.8 Install Original Controller

1. If installing option kit DS-HSZ22-AA, add the second SIMM to the original controller. If installing option kit DS-HSZ22-AB, replace the existing SIMM with the two remaining SIMMs supplied with the kit. (See Section 5.2.4 for instructions.)
2. Install the second controller into the pedestal (or controller shelf) and power up the system. The system is now in the dual-redundant operating mode.
3. Locate the second serial cable that came in the second controller kit. Connect this cable to the second controller serial port on the RA3000 and to COM 2 port and restart SWCC in serial mode. You now have two serial connections.
4. To check that the controller pair is in redundant mode, click on the **Controller** icon. Properties for both controllers should be visible. Firmware revision of both controllers must be identical. If either controller is shown in gray, refer to your *RA3000 Getting Started Guide* for more detailed instructions.

5.3 Configuring a Dual Controller Installation for a Single Serial Port

This section describes how to configure your RA3000 storage system for dual controller operation when only one serial port is available on the host. If required, refer to your Getting Started guide for SWCC installation instructions. The information in this section is divided as follows:

- Configuring the RA3000 Storage System for dual-controller operation
- Connecting to your dual controller storage system using a serial connection
- Verifying the Controller Operating Parameters

5.3.1 Configuring the RA3000 for the Active Mode of Operation

If you are upgrading from a single controller storage system, the installed controller will be in the active/passive mode of operation. To configure the storage system using a single serial connection, the dual controllers must be placed in active/active mode. Perform the procedure to change your storage system from “active/passive” to “active/active” mode of operation.

1. Shutdown the RA3000 storage system, install the second controller, and restart the system.
2. Connect the host serial port to the serial port connector on the top controller of the RA3000.
3. Open an HSZ22 monitor utility terminal session. The terminal settings should be: 9600 BAUD, 8 bits, no parity, 1 stop bit, XON/XOFF. If prompted for a password, enter “RAID”.
4. When the terminal session is displayed, press “ESCAPE/” (key sequence: Esc, Shift, 7). The controller should respond with banner listing *HSZ22 Monitor Utility* followed by the Firmware revision number.
5. Select *Setup Parameters* from the main menu.
6. Select *Rdnt Ctrlr Parameters* and change the value of Host I/O Channel 1 to *Passive*. Press Ctrl-Z to exit.
7. Transfer the serial cable to the serial port on the bottom controller and press “Ctrl-Z”. When the Monitor Utility screen is displayed, press any key to continue.
8. Select *Setup Parameters* and then select *Rdnt Ctrlr Parameters* from the main menu.
9. Both channels should be displayed as “Passive”. Select the host I/O channel that was **not** reconfigured in step 5 and change its value *Active*. Press CTRL-Z to exit.

10. Restart the controllers by power cycling the RA3000 storage system.
11. Recheck the *Rdnt Ctrlr Parameters*. The controller *Values* should display *Active/Passive*.
12. Transfer the serial cable to the serial port on the top controller on the RA3000 and press Ctrl-Z.
13. Check the *Rdnt Ctrlr Parameters*. The controller *Values* should display *Passive/Active*.

Your storage system is now properly configured for dual controller operation from a single serial host port. Proceed to the next section to configure the controllers.

5.3.2 Configuring the Dual Controllers

In order to create a virtual disk, a serial connection must be made to the controller. The serial connection provides a local connection to the RA3000 (HSZ22) controller. To configure the controller, perform the following steps:

1. Ensure the RA3000 Storage System is turned off.
2. Connect the serial cable between the host serial connector and the top controller's serial port connector.
3. Power on the RA3000 storage system.
4. Open the StorageWorks management utility.
5. Select the *Serial* button in the HSZ22 Management window.
6. Select the serial port available on the host's workstation and then click on the Connect button.
7. When the second connection window appears, select Cancel. Once the connection is established, a storage window will appear with the connected port identified at the top of the window.
8. From the Storage menu, select *Add Virtual Disk*. Select *JBOD* from the menu (you may select any RAID level you wish) and then select Next.
9. Select *drive* from the window *Devices Available to Create a New Virtual Disk* and then select Next.
10. Select Add and then select Next.
11. A window is displayed prompting for selection of the host port that the "JBOD" should be created on. There is a single connection to the controller so only one host port will be available. Select Next.
12. Choose a LUN number and enable the write-back cache.
13. Select Next and then Finish.

14. You have now created one Virtual Device (JBOD) on the top controller. Disconnect the serial cable from the top controller and reconnect the cable to serial port connector on the bottom (redundant) controller.
15. Close and reopen the StorageWorks window and repeat steps 4 through 13 above (Note: If the storage window connection is lost when the cable is moved, reopen the *HSZ22 StorageWindow*).
16. Once the second “Virtual Disk” has been created, close the StorageWorks window.
17. To complete dual-controller configuration, perform the following.
 - Ensure all cables and connections to your server are in place
 - Leave the RA3000 Rack Mount Storage System powered on
 - Reboot the server

Refer to your *Getting Started Guide* to configure the virtual disks as communications LUNs.

5.3.3 Connecting to Your Dual Controller Storage System

1. Open the StorageWorks management utility.
2. Refer to the *Getting Started Guide* for instructions describing how to connect the RA3000 to your host system.

5.3.4 Verifying the Controller Operating Parameters

For optimal dual- controller performance and load balancing, the host ports should be configured as listed below. Verify these parameters from the “Controller Properties” page:

- Host I/O channel 0 for Controller 1 (top) to Active
- Host I/O channel 1 for Controller 1 to Passive
- Channel 0 Target ID for Controller 1 to 0
- Channel 1 Target ID for Controller 2 to 1

NOTE

It is allowable to configure the host I/O and target IDs as follows:

Controller 1 (top) may be set to “Passive” while host I/O channel 1 for Controller 1 is set to “Active”. As long as each host channel value and channel target ID value is unique, the system is correctly configured for Active/Active mode.

1. Select the “Environment” tab from the *Controller Properties* window.
2. Ensure the cabinet and UPS settings are set to ‘normal’.
3. Select the Cache tab from the “Controller Properties” window. The SIMMs parameters must be identical between controllers.
4. Select *OK*. The system is now ready for dual-controller operation. Configuration of additional RAIDsets may now be performed as desired.