



Sun StorEdge™ 3310 SCSI RAID Firmware Guide™

Version 3.25

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-7296-11
October 2002, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright © 2002 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, California 92009, USA. All rights reserved.

Sun Microsystems, Inc. and Dot Hill Corporation may have intellectual property rights relating to technology embodied in this product or document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and other countries.

This product or document is distributed under licenses restricting its use, copying distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

Federal Acquisitions: Commercial Software - Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2002 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, Californie 92009, USA. Tous droits réservés.

Sun Microsystems, Inc. et Dot Hill Systems Corporation peuvent avoir les droits de propriété intellectuels relatants à la technologie incorporée dans ce produit. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats - Unis et les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.

Contents

Preface xi

How This Book Is Organized xi

Using UNIX Commands xii

Typographic Conventions xiii

Shell Prompts xiii

Related Documentation xiv

Technical Support xiv

Accessing Sun Documentation Online xv

Sun Welcomes Your Comments xv

1. Basic RAID Concepts and Planning 1-1

1.1 Logical Drive 1-2

1.2 Logical Volume 1-2

1.3 RAID Levels 1-3

1.3.1 NRAID (Disk Spanning) 1-4

1.3.2 JBOD/Expansion Unit (Single Drive Control) 1-5

1.3.3 RAID 1 (Disk Mirroring) 1-6

1.3.4 RAID (0+1) (Disk Striping with Mirroring) 1-7

1.3.5 RAID 3 (Disk Striping with Dedicated Parity Disk) 1-8

1.3.6 RAID 5 (Striping with Interspersed Parity) 1-9

- 1.4 Local and Global Spare Drives 1-10
 - 1.4.1 Local Spare Drive 1-10
 - 1.4.2 Global Spare Drive 1-11
 - 1.4.3 Having Both Local and Global Spares 1-12
- 1.5 Controller Defaults and Limitations 1-13
- 1.6 Battery Operation 1-14
- 1.7 RAID Planning Considerations 1-15
- 1.8 Basic Configuration 1-16
- 2. Accessing the Controller Firmware 2-1**
 - 2.1 Setting Up the Serial Port Connection 2-1
 - 2.1.1 Accessing the Firmware Application from a Solaris Host 2-2
 - 2.1.2 The Controller Firmware Initial Screens 2-4
 - 2.1.3 Main Menu 2-5
 - 2.2 Upgrading Firmware 2-5
- 3. Viewing and Editing Logical Drives 3-1**
 - 3.1 Introducing the Logical Drive Commands 3-2
 - 3.2 Default Logical Drives and RAID Levels 3-3
 - 3.3 Viewing the Logical Drive Status Table 3-4
 - 3.4 Creating Logical Drive(s) 3-5
 - 3.5 Changing a Logical Drive Controller Assignment 3-11
 - 3.6 Partitioning a Logical Drive 3-13
 - 3.7 Deleting a Logical Drive 3-16
 - 3.8 Deleting a Partition of a Logical Drive 3-17
 - 3.9 Assigning Logical Drive Name 3-18
 - 3.10 Rebuilding a Logical Drive 3-19
 - 3.11 Performing a Logical Drive Parity Check 3-20
 - 3.12 Adding a SCSI Drive to a Logical Drive 3-21

- 3.13 Copying and Replacing Drives with Drives of Larger Capacity 3-25
- 3.14 Expanding a Logical Drive 3-28
- 4. Viewing and Editing Logical Volumes 4-1**
 - 4.1 Understanding Logical Volumes (Multi-Level RAID) 4-2
 - 4.2 Creating a Logical Volume 4-5
 - 4.3 Expanding a Logical Volume 4-7
 - 4.4 Viewing a Logical Volume Status Table 4-8
- 5. Viewing and Editing Host LUNs 5-1**
 - 5.1 Planning for 128 LUNs 5-1
 - 5.2 Mapping Logical Drive Partitions to Host LUNs 5-3
 - 5.2.1 Example Host LUN Mappings 5-6
 - 5.3 Deleting a Host LUN Mapping 5-7
- 6. Viewing and Editing SCSI Drives 6-1**
 - 6.1 SCSI Drive Status Table 6-2
 - 6.2 Assigning a Local Spare Drive 6-4
 - 6.3 Creating a Global Spare 6-5
 - 6.4 Viewing Drive Information 6-6
 - 6.5 Viewing the Connected Drives 6-7
 - 6.6 Scanning Drives 6-8
 - 6.7 Deleting a Spare Drive 6-9
 - 6.8 Setting Slot Numbers 6-10
 - 6.8.1 Assigning a Slot Number to an Empty Slot 6-10
 - 6.8.2 Deleting a Slot Number 6-11
 - 6.9 Adding or Deleting Drive Entries 6-11
 - 6.9.1 Removing an Empty Drive Entry 6-12
 - 6.10 Identifying a Drive 6-13
 - 6.10.1 Flashing Selected SCSI Drive 6-14

- 6.10.2 Flashing All SCSI Drives 6-15
- 6.10.3 Flashing All But Selected Drives 6-15
- 6.11 Cloning Failing Drive 6-16
 - 6.11.1 Replacing After Clone 6-16
 - 6.11.2 Perpetual Clone 6-18
- 6.12 SCSI Drive Utilities (Do Not Use) 6-21
- 7. Viewing and Editing SCSI Channels 7-1**
 - 7.1 SCSI Channel Status Table 7-2
 - 7.1.1 SCSI Drive Channel Commands 7-4
 - 7.1.2 SCSI Host Channel Commands 7-5
 - 7.2 Configuring SCSI Channels as Host or Drive 7-5
 - 7.3 Permanent SCSI Drive Channel IDs 7-7
 - 7.4 Creating Additional Host IDs 7-9
 - 7.5 Deleting a Host Channel SCSI ID 7-10
 - 7.6 Drive Channel SCSI IDs (Do Not Change) 7-11
 - 7.7 Setting a SCSI Channel Termination (Do Not Change) 7-12
 - 7.8 Setting a Transfer Clock Speed 7-13
 - 7.9 Setting a Transfer Width 7-14
 - 7.10 Viewing and Editing SCSI Target / Drive Channel 7-15
 - 7.10.1 Slot Number 7-16
 - 7.10.2 Maximum Synchronous Transfer Clock 7-16
 - 7.10.3 Maximum Transfer Width 7-16
 - 7.10.4 Parity Check 7-17
 - 7.10.5 Disconnecting Support 7-17
 - 7.10.6 Maximum Tag Count 7-18
- 8. Viewing and Editing Configuration Parameters 8-1**
 - 8.1 Optimization Modes (Caching Parameters) 8-2

8.1.1	Optimization for Random or Sequential I/O	8-3
8.1.2	Write-Back/Write-Through Cache Enable/Disable	8-4
8.2	Controller Failure Symptoms, Logical Drive Rebuild, and Replacement	8-5
8.2.1	General Controller Failure Response	8-5
8.2.2	Automatic Logical Drive Rebuild	8-6
8.2.3	Manual Rebuild	8-6
8.2.4	Concurrent Rebuild in RAID (0+1)	8-9
8.2.5	Identifying a Failed Drive for Replacement	8-9
8.2.6	Restoring Your Configuration (NVRAM) From a File	8-9
8.2.7	Determining Fatal Drive Failure	8-10
8.3	Controller Parameters	8-12
8.3.1	Controller Name	8-12
8.3.2	LCD Title Display - Controller Logo (Not Applicable)	8-13
8.3.3	Password Validation Timeout	8-13
8.3.4	Controller Unique Identifier	8-15
8.3.5	SDRAM ECC Function (Do Not Change)	8-15
8.4	Drive-side SCSI Parameters	8-16
8.4.1	SCSI Motor Spin-Up (Do Not Change)	8-17
8.4.2	SCSI Reset at Power-Up (Do Not Change)	8-18
8.4.3	Disk Access Delay Time	8-19
8.4.4	SCSI I/O Timeout	8-19
8.4.5	Maximum Tag Count (Tag Command Queuing)	8-21
8.4.6	SAF-TE and S.E.S. Enclosure Monitoring	8-22
8.4.7	Periodic Drive Check Time	8-22
8.4.8	Auto-Detect Failure Drive Swap Check Time	8-23
8.5	Disk Array Parameters	8-24
8.5.1	Rebuild Priority	8-24
8.5.2	Verification on Writes	8-25

- 8.6 Host-side SCSI Parameters 8-27
 - 8.6.1 Overview of SCSI Channel, SCSI ID and LUN 8-27
 - 8.6.2 Maximum Concurrent Host-LUN Connections 8-28
 - 8.6.3 Number of Tags Reserved for Each Host LUN Connection 8-29
 - 8.6.4 Maximum Queued I/O Count 8-29
 - 8.6.5 LUNs Per Host SCSI ID 8-30
 - 8.6.6 Cylinder/Head/Sector Mapping 8-31
- 8.7 Peripheral Device Type Parameters 8-33
- 8.8 Setting an IP Address 8-34

- 9. Viewing and Editing Peripheral Devices 9-1**
 - 9.1 Viewing Peripheral Device Controller Status 9-2
 - 9.2 Viewing Peripheral Device SAF-TE Status 9-2
 - 9.3 Setting Peripheral Device Entry 9-5
 - 9.3.1 Redundant Controller Mode (Do Not Change) 9-5
 - 9.3.2 Enable UPS Status 9-6
 - 9.4 Setting the UPS Power Fail Signal 9-7
 - 9.5 Viewing the Controller Voltage and Temperature Status Window 9-8

- 10. System Functions, Info, and Event Logs 10-1**
 - 10.1 System Functions 10-2
 - 10.1.1 Muting the Beeper 10-2
 - 10.1.2 Setting a New Password 10-3
 - 10.1.3 Changing Password 10-3
 - 10.1.4 Disabling the Password 10-4
 - 10.1.5 Resetting Controller 10-4
 - 10.1.6 Shutting Down the Controller (Restricted Use) 10-5
 - 10.2 Restoring Configuration (NVRAM) from a File 10-5
 - 10.3 Saving Configuration (NVRAM) to Disk 10-6

10.4 Viewing Event Logs on Screen 10-7

A. Firmware Specifications A-1

B. Parameter Summary Tables B-1

B.1 Optimizing Controller Parameters B-1

B.2 Keeping Specific Parameter Defaults B-3

C. Event Messages C-1

C.1 Controller Events C-2

Alerts C-2

Warnings C-2

Notifications C-3

C.2 SCSI Drive Events C-4

Warnings C-4

Notifications C-5

C.3 SCSI Channel Events C-6

Alerts C-6

Notifications C-6

C.4 Logical Drive Events C-7

Alerts C-7

Notifications C-8

C.5 General Target Alerts C-9

SAF-TE Device C-9

Controller on-board C-10

I²C Device C-11

SES Devices C-12

General Peripheral Device C-14

Glossary Glossary-1

Acronyms Glossary-1

Terminology Glossary-3

Preface

The *Sun StorEdge™ 3310 SCSI RAID Firmware Guide* gives an overview of RAID terminology, emphasizes the use of controller firmware and VT100 commands to configure and monitor the Sun StorEdge units. The guide also describes the available Command Line Interface (CLI) commands which can also be used for performing administrative tasks on the arrays.

This manual is written for Sun support and Sun-trained personnel who are already familiar with Sun's hardware and software products.



Caution – You should read the *Sun StorEdge 3310 SCSI Array Safety, Regulatory, and Compliance Manual* before beginning any procedure in this manual.

How This Book Is Organized

This book contains the following topics:

Chapter 1 introduces the RAID terminology concepts.

Chapter 2 describes how to access the array via the COM port and a terminal emulation program.

Chapter 3 describes the “view and edit Logical drives” command and related procedures.

Chapter 4 describes the “view and edit logical volumes” command and related procedures.

Chapter 5 describes the “view and edit Host luns” command and related procedures.

Chapter 6 describes the “view and edit scsi Drives” command and related procedures.

Chapter 7 describes the “view and edit Scsi channels” command and related procedures.

Chapter 8 describes the “view and edit Configuration parameters” command and related procedures.

Chapter 9 describes the “view and edit Peripheral devices” command and related procedures.

Chapter 10 describes the “system functions” command, array information, and event logs.

Appendix A provides the array firmware specifications.

Appendix B summarizes controller parameters for optimization and parameter defaults which should not be changed.

Appendix C lists the event messages.

The Glossary provides RAID terminology and definitions used throughout the product documentation

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2™ online documentation for the Solaris™ operating system
- Other software documentation that you received with your system

Typographic Conventions

TABLE P-1

Typeface*	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

* The settings on your browser might differ from these settings.

Shell Prompts

TABLE P-2

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Late-breaking news	<i>Sun StorEdge 3310 SCSI Array Release Notes</i>	816-7292
Installation and service	<i>Sun StorEdge 3310 SCSI Array Installation, Operation, and Service Manual</i>	816-7290
Best practices	<i>Sun StorEdge 3310 SCSI Array Best Practices Manual</i>	816-7293
Safety and compliance	<i>Sun StorEdge 3310 SCSI Array Safety, Regulatory, and Compliance Manual</i>	816-7291
Monitoring and configuration software	<i>Sun StorEdge 3310 SCSI Configuration Service User Guide</i>	816-7298
Reporting software	<i>Sun StorEdge 3310 SCSI Diagnostic Reporter User Guide</i>	816-7722
Command line interface tools	<i>Sun StorEdge 3310 SCSI CLI Version 1.0 User Guide</i>	816-7297

Technical Support

- For late-breaking news and troubleshooting tips, review the *Sun StorEdge 3310 SCSI Array Release Notes* located at:

www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3310

- For 24-hour access to web-based support solutions, visit the Online Support Center at:

www.sun.com/service/online

- To initiate or check on a USA-only service request, contact Sun support at:
1-800-USA4SUN

- To obtain international technical support, contact the sales office of a specific country at:

www.sun.com/service/contacting/sales.html

Accessing Sun Documentation Online

All Sun StorEdge 3310 SCSI Array online documentation is located at:

http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3310

You can order printed copies of the Sun StorEdge 3310 SCSI array manuals at:

<http://corppub.iuniverse.com/marketplace/sun>

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

<http://www.sun.com/documentation>

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

docfeedback@sun.com

Please include the part number (816-7296-11) of your document in the subject line of your email.

Basic RAID Concepts and Planning

Redundant Arrays of Independent Disks, or RAID, offers major benefits in availability, capacity, and performance.

The Sun StorEdge arrays provide complete RAID functionality and enhanced drive failure management.

This chapter covers the following concepts and planning guidelines:

- “Logical Drive” on page 1-2
- “Logical Volume” on page 1-2
- “RAID Levels” on page 1-3
- “Local and Global Spare Drives” on page 1-10
- “Controller Defaults and Limitations” on page 1-13
- “RAID Planning Considerations” on page 1-15
- “Basic Configuration” on page 1-16

1.1 Logical Drive

Increased availability, capacity, and performance are achieved by creating logical drives. A logical drive is simply an array of independent physical drives. The logical drive appears to the host the same as a local hard disk drive does.

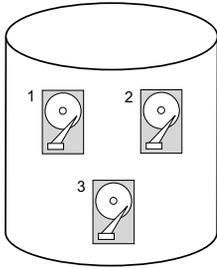


FIGURE 1-1 Logical Drive Including Multiple Physical Drives

The different methods by which logical drives are referred to as RAID levels.

1.2 Logical Volume

The concept of a logical volume is very similar to that of a logical drive. A logical volume is composed of one or several logical drives. The member logical drives can be composed of the same RAID level or each of different RAID levels.

The logical volume can be divided into a maximum of 8 partitions.

During operation, the host sees a non-partitioned logical volume or a partition of a partitioned logical volume as one single physical drive.

1.3 RAID Levels

Using a RAID array has the following advantages:

- Provides disk spanning by weaving all connected drives into one single volume.
- Increases disk access speed by breaking data into several blocks when reading/writing to several drives in parallel. With RAID, storage speed increases as more drives are added as the channel bus allows.
- Provides fault-tolerance by mirroring or parity operation.

Note – Drives on different channels can be included in a logical drive and logical drives of different RAID levels can be used to configure a logical volume.

RAID Level	Description
RAID 0	Striping without fault tolerance; provides maximum performance.
RAID 1	Mirrored or duplexed disks; for each disk in the array, a duplicate disk is maintained for fault tolerance. RAID 1 does not improve performance over that of a single disk drive. It requires 50% of total disk capacity for overhead.
RAID 3	One drive is dedicated to parity. Data is divided into blocks and distributed sequentially among the remaining drives. You need at least three physical drives for a RAID 3 logical drive.
RAID 5	Striping with fault tolerance; this is the best-suited RAID level for multi-tasking or transaction processing. In RAID 5, an entire transfer block is placed on a single drive, but there are no dedicated data or Error Correction Code (ECC) drives. The data and ECC are striped across each drive in the disk array, so that each drive contains a combination of data and ECC blocks. This allows data to be reconstructed on a replacement drive in the event of a single disk drive failure. The primary advantages of RAID 5 are: <ul style="list-style-type: none">• Provides fault tolerance• Increases performance through the ability to perform both read and write seeks in parallel• Low cost per usable megabyte of disk storage.• RAID 5 also requires at least 3 drives.

1.3.1 NRAID (Disk Spanning)

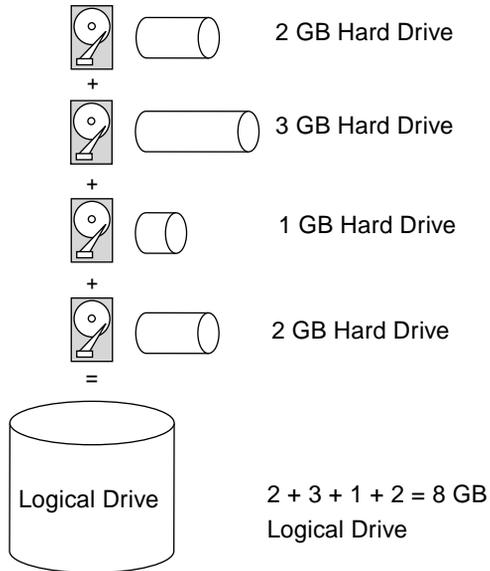


FIGURE 1-2 NRAID Configuration

NRAID stands for Non-RAID. The capacity of all the drives is combined to become one logical drive (no block striping). In other words, the capacity of the logical drive is the total capacity of the physical drives. NRAID does not provide data redundancy.

1.3.2 JBOD/Expansion Unit (Single Drive Control)

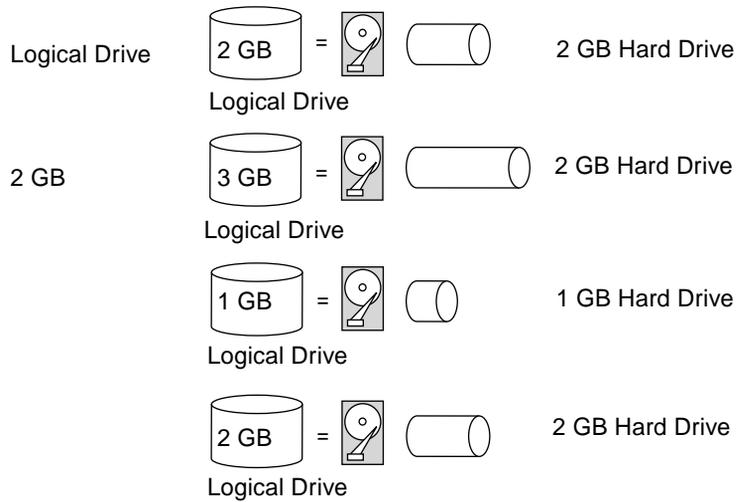


FIGURE 1-3 JBOD Configuration

JBOD/Expansion Unit stands for Just a Bunch of Disks. The controller treats each drive as a stand-alone disk, therefore each drive is an independent logical drive. JBOD/Expansion Unit is a disk array without a controller and does not provide data redundancy.

RAID 0 (Disk Striping)

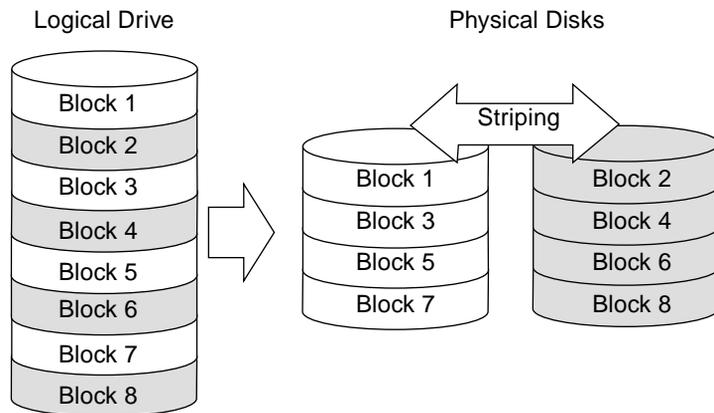


FIGURE 1-4 RAID 0 Configuration

RAID 0 provides the highest performance but no redundancy. Data in the logical drive is striped (distributed) across several physical drives.

1.3.3 RAID 1 (Disk Mirroring)

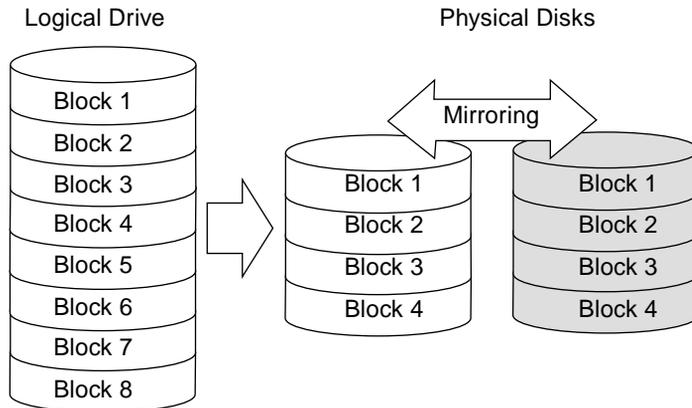


FIGURE 1-5 RAID 1 Configuration

RAID 1 mirrors the data stored in one hard drive to another. RAID 1 can only be performed with two hard drives. If there are more than two hard drives, RAID (0+1) will be performed automatically.

1.3.4 RAID (0+1) (Disk Striping with Mirroring)

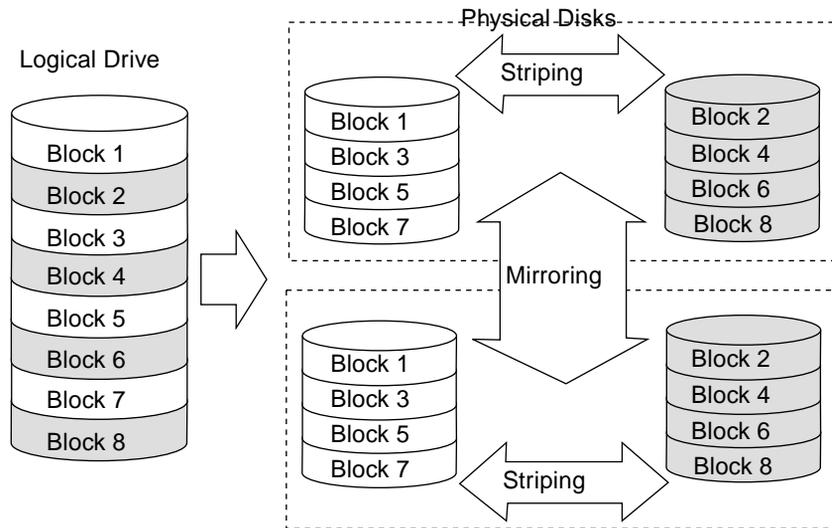


FIGURE 1-6 RAID 0+1 Configuration

RAID (0+1) combines RAID 0 and RAID 1, Mirroring and Striping. RAID (0+1) allows multiple drive failure because of the full redundancy of the hard drives. If there are more than two hard drives assigned to perform RAID 1, RAID (0+1) will be performed automatically.

Note – RAID (0+1) will not appear in the list of RAID levels supported by the controller. If you want to perform RAID 1, the controller will determine whether to perform RAID 1 or RAID (0+1). This will depend on the number of drives that has been selected for the logical drive.

1.3.5 RAID 3 (Disk Striping with Dedicated Parity Disk)

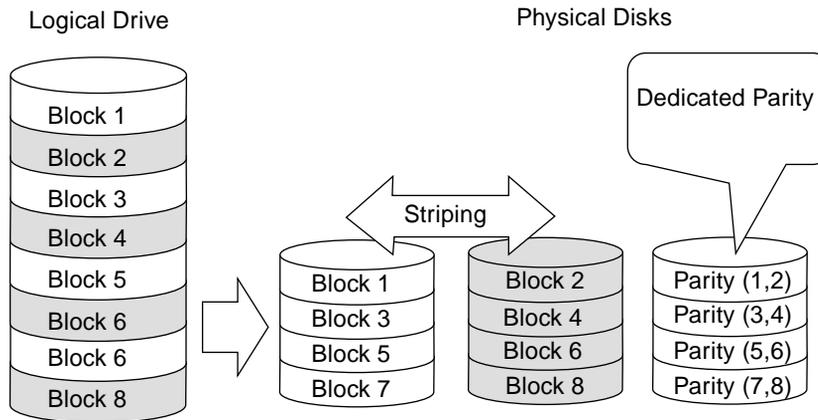


FIGURE 1-7 RAID 3 Configuration

RAID 3 performs block striping with dedicated parity. One drive member is dedicated to storing the parity data. When a drive member fails, the controller can recover/regenerate the lost data of the failed drive from the dedicated parity drive.

RAID 3

Minimum Disks Required	3
Capacity	N - 1
Redundancy	Yes

1.3.6 RAID 5 (Striping with Interspersed Parity)

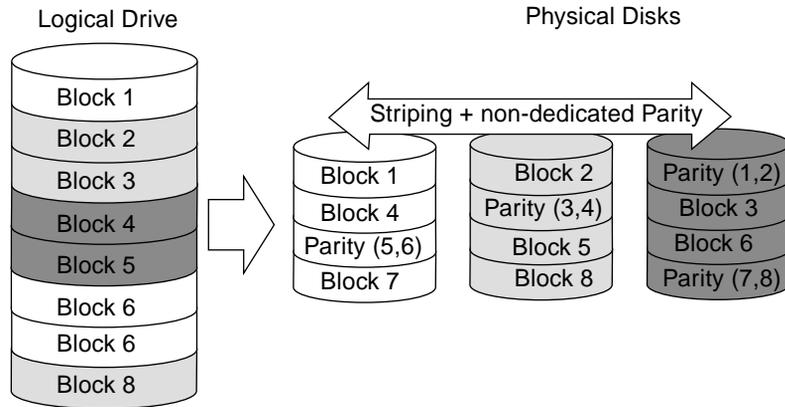


FIGURE 1-8 RAID 5 Configuration

RAID 5 is similar to RAID 3 but the parity data is not stored in one dedicated hard drive. Parity data is interspersed across the drive array. In the event of a failure, the controller can recover/regenerate the lost data of the failed drive from the other surviving drives.

RAID 5	
Minimum Disks Required:	3
Capacity:	N - 1
Redundancy:	Yes

1.4 Local and Global Spare Drives

The external RAID controllers provide both local spare drive and global spare drive functions. **The local spare drive is used only for one specified drive; the global spare drive can be used for any logical drive on the array.**

The local spare drive always has higher priority than the global spare drive; therefore, if a drive fails and both types of spares are available at the same or greater size to replace the failed drive, the local spare is used.

If there is a failed drive in the RAID 5 logical drive, replace the failed drive with a new drive to keep the logical drive working. To identify a failed drive, refer to “Identifying a Drive” on page 6-13.



Caution – If, when trying to remove a failed drive, you mistakenly remove the wrong drive, you will no longer be able to access the logical drive because you have incorrectly failed another drive.

1.4.1 Local Spare Drive

A local spare drive is a standby drive assigned to serve one specified logical drive. When a member drive of this specified logical drive fails, the local spare drive becomes a member drive and automatically starts to rebuild.

A local spare drive always has higher priority than a global spare drive; that is, if a drive fails and there is a local spare and a global spare drive available, the local spare drive will be used.

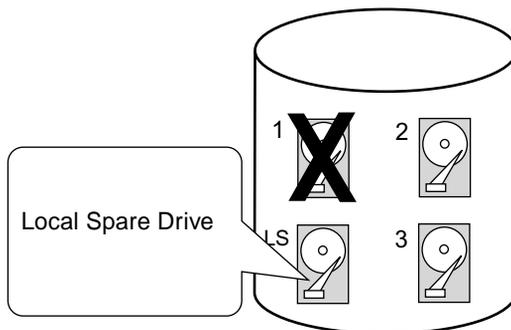


FIGURE 1-9 Local (Dedicated) Spare

1.4.2 Global Spare Drive

A global spare drive does not only serve one specified logical drive; it is available for all logical drives (see FIGURE 1-10). When a member drive from any of the logical drives fails, the global spare drive joins that logical drive and automatically starts to rebuild.

A local spare drive always has higher priority than a global spare drive; that is, if a drive fails and there is a local spare and a global spare drive available, the local spare drive will be used.

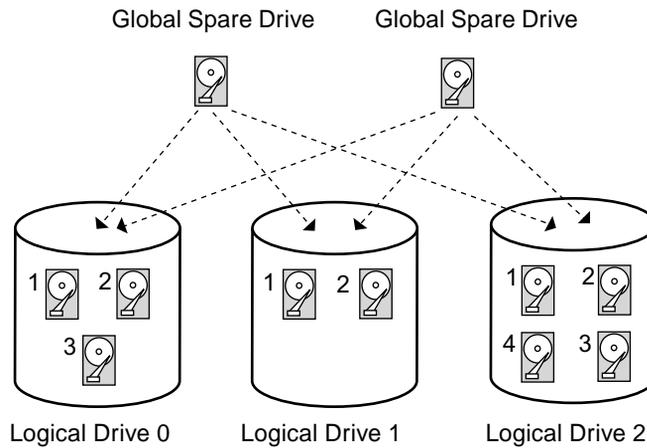


FIGURE 1-10 Global Spare

1.4.3 Having Both Local and Global Spares

In FIGURE 1-11, the member drives in Logical Drive 0 are 9 GB drives, and the members in Logical Drives 1 and 2 are all 4 GB drives.

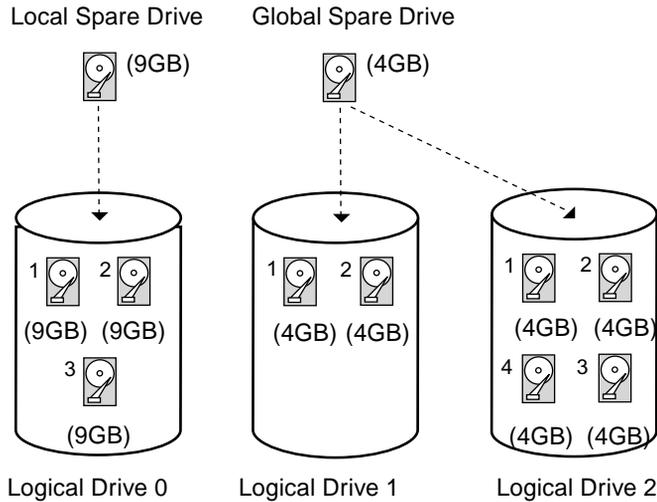


FIGURE 1-11 Mixing Local and Global Spares

A local spare drive always has higher priority than a global spare drive; that is, if a drive fails and both a local spare and a global spare drive are available, the local spare drive will be used.

In FIGURE 1-11, it is not possible for the 4 GB global spare drive to join logical drive 0 because of its insufficient capacity. The 9 GB local spare drive will aid logical drive 0 once a drive in this logical drive fails. If the failed drive is in logical drive 1 or 2, the 4 GB global spare drive will immediately give aid to the failed drive.

1.5 Controller Defaults and Limitations

The following controller functions describe the redundant controller operation.

- Both controllers must be exactly the same. Namely, they must operate with the same firmware version, the same size of memory, the same number of host and drive channels, etc. When a replacement controller is placed in a system, the firmware of the first controller automatically synchronizes (overwrites) the firmware of the second controller to be the same firmware.
- In redundant mode, the maximum number of disk drive IDs on a SCSI drive channel is 16. IDs 6 and 7 are used for host HBA connections.
- Both controllers must be initially configured as primary controllers. Upon bootup in a redundant configuration, the controllers autonegotiate and designate one controller as primary and the other controller as secondary.
- The two controllers behave as one primary controller. Once the redundant configuration takes effect, user configurations and settings can be done only on the primary controller. The secondary controller then synchronizes with the configuration of the primary controller, making the configurations of the two controllers exactly the same.

The two controllers continuously monitor each other. When a controller detects that the other controller is not responding, the working controller immediately takes over and disables the failed controller.

- It is necessary to connect all interfaces to both controllers so that the surviving controller can readily continue all services provided for the RAID system. For example, if you connect one controller to the Ethernet, you should also connect the second controller to the Ethernet.
- In an active-to-active configuration (standard configuration), you can assign any appropriate logical drives to either of the controllers, and then map the logical configurations to host channel IDs/LUNs. I/O requests from host computer are directed to the primary or the secondary controller accordingly. The total drive capacity can be grouped into several logical configurations and equally assigned to both controllers so that they share the workload.

The active-to-active configuration engages all array resources to actively maximize performance. Users might also assign all logical configurations to one controller and let the other act as a standby.

An active-to-standby configuration is an available configuration but it is not usually selected. By assigning all the logical configurations of drives to one controller, the other controller stays idle and becomes active only when its counterpart fails.

1.6 Battery Operation

A battery should be changed every 2 years if the unit is continuously operated at 25 degrees C and yearly at 35 degrees C or higher. The FRU shelf life for a battery is 3 years.

Battery Status

The battery LED (on far right side of the controller module) will be a yellow LED if the battery is bad or missing. The LED will be blinking green if the battery is charging and will be solid green when the battery is fully charged.

The initial firmware screen will also display the battery status at the top of the initial screen where BAT: status will be BAD, ----- (charging), or +++++ (fully charged).

Write-Back Versus Write-Through Cache Options

Unfinished writes are cached in memory in write-back mode. If power to the array is discontinued, data stored in the cache memory will not be lost. Battery modules can support cache memory for several days.

Write cache is not automatically disabled when the battery is offline due to battery failure or a disconnected battery. You can enable or disable the write-back cache capabilities of the RAID controller. To ensure data integrity, you may choose to disable the Write-Back cache option and switch to the Write-Through cache option through the firmware application (go to “view and edit Configuration parameters” and select “Caching Parameters”). The risk of data loss is remote.

1.7 RAID Planning Considerations

Here are answers to some questions that may help you plan your RAID array.

- How many physical drives do you have?

Depending on the Sun StorEdge array ordered, you will have 5 drives or 12 drives in your array.

- How many drives would you like to appear to the host computer?

Determine what capacity will be included in a logical configuration of drives. A logical configuration of drives is displayed to the host as a single physical drive. For the default logical drive configuration, refer to “Default Logical Drives and RAID Levels” on page 3-3.

- What kind of host application?

The frequency of read/write activities can vary from one host application to another. The application can be a SQL server, Oracle server, Informix, or other data base server of a transaction-based nature. Applications like video playback and video post-production editing require read/writes of huge files in a sequential order.

The RAID level setting depends on what is the most important for a given application – capacity, availability, or performance. Before revising your RAID level (prior to storing data), you need to choose an optimization scheme and optimize the controller for your application.

The controller optimization mode can only be changed when there is no existing logical configurations. Once the controller optimization mode is set, the same mode will be applied to all logical drives. Data stripe size is changed once the optimization method is changed. Therefore, you can not proceed with changing the optimization mode until data is backed-up, all logical drives deleted, and array restarted. Therefore, think twice before choosing an optimization mode for your controller.

The controller factory defaults guarantee the optimal performance for most applications.

- How many logical drives and at what RAID level?

A *logical drive* is a set of drives grouped together to operate under a given RAID level and appears as a single contiguous storage volume. The controller is capable of grouping drives into 8 logical drives, each configured on the same or different RAID levels. Different RAID levels provide varying degrees of performance and fault tolerance.

A *logical volume* can be created but is rarely created due to the greater scalability and ease of use with logical drives.

- Any spare drives?

Spare drives allow for the unattended rebuilding of a failed drive, heightening the degree of fault tolerance. If there is no spare drive, data rebuild has to be performed manually by replacing a failed drive with a healthy one.

Drives must be configured and the controller properly initialized before host computer can access the storage capacity.

1.8 Basic Configuration

1. If a hard drive was connected after the controller completes initialization, use the “Scan scsi drive” function (under the “view and edit scsi Drives” in the Main Menu) to let the controller recognize the newly added hard drive and to make it available to be configured as a member of a logical drive.
2. Optionally define any additional partitions for each logical drive. Refer to “Partitioning a Logical Drive” on page 3-13. *A logical drive can be divided into a maximum of 32 Partitions. A total of 128 partitions can be created in an array.*
 - Optionally add host SCSI IDs and more logical drives to create 128 LUNs. Refer to:
 - “Planning for 128 LUNs” on page 5-1
 - “Creating Additional Host IDs” on page 7-9
 - “Creating Logical Drive(s)” on page 3-5
3. Map each logical drive and storage partition to a host ID/LUN. Refer to “Mapping Logical Drive Partitions to Host LUNs” on page 5-3. The host adapter will recognize the system drives after re-initializing the host bus.
4. Optionally optimize controller's parameters for your applications. For details on optimization modes, refer to “Optimization Modes (Caching Parameters)” on page 8-2.
5. The last step is to save your configuration profile as a file or to the logical drive you created.

Note – The controller is totally independent from host operating environment. The host operating environment will not be able to tell whether the attached storage is a physical hard drive or the logical drives created by the RAID controller.

Accessing the Controller Firmware

The RAID controller firmware can be configured via a workstation running a terminal emulation program or on a VT-100 compatible terminal. Topics covered in this chapter include:

- “Setting Up the Serial Port Connection” on page 2-1
- “Upgrading Firmware” on page 2-5

2.1 Setting Up the Serial Port Connection

The RAID controller can be configured by means of a Solaris workstation running a VT-100 terminal emulation program or by a Windows terminal emulation program such as HyperTerminal.

Note – You can also monitor and configure a RAID array over an IP network with the Configuration Service program after you assign an IP address to the array. For details, refer to Appendix C in the *Sun StorEdge 3310 SCSI Array Installation, Operation, and Service Manual* and to the *Sun StorEdge 3310 SCSI Configuration Service User Guide*.

To set up the serial port connection, perform the following steps.

1. Use a serial cable to connect the COM port of the RAID array to serial port b on a Solaris workstation.

A DB9-to-DB25 serial cable adapter is provided to connect the serial cable to DB25 serial ports on workstations.

Make sure that a null modem cable can be attached to the host serial port. The null modem cable has serial signals swapped for connecting to a standard serial interface.

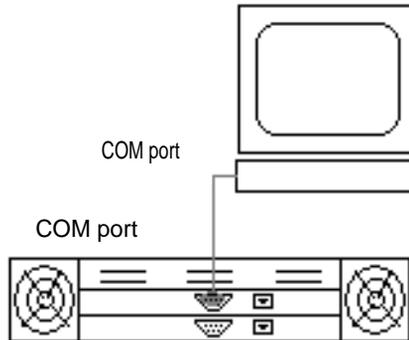


FIGURE 2-1 RAID Array COM Port Connected Locally to the COM Port of a Workstation or Computer Terminal

2. Power up the array.
3. Set the serial port parameters on the workstation as follows: 38400 baud, 8 bit, 1 stop bit, no parity.

2.1.1 Accessing the Firmware Application from a Solaris Host

See “Redefining the Baud Rate for the `tip` Command” on page 2-3 for setting serial port parameters on a Solaris host.

See “Using the `tip` Command for Local Access to the Array” on page 2-3 for instructions on using the Solaris `tip` command.

Redefining the Baud Rate for the `tip` Command

This section describes the method for using the Solaris `tip` command to access the array. Before you can use the `tip` command, you must redefine the baud rate.

The `tip` command uses the 9600 baud specified in the `/etc/remote` file on the Solaris host by default. The serial port of the array requires 38400 baud so you must edit the `etc/remote` file and revise your use of the `tip` command, to use the 38400 baud rate.

To edit the `/etc/remote` file for the 38400 baud rate, perform the following steps:

1. In the `/etc/remote` file, copy and edit the `hardware` line as follows:

Locate:

```
hardware: :dv=/dev/term/b:br#9600:el= . . .
```

Copy the `hardware` line into a blank line below the `hardware` line, rename the `hardware` term as `ss3000` and replace `#9600` with `#38400`. The editing will look like the following line:

```
ss3000: :dv=/dev/term/b:br#38400:el= . . .
```

2. Save the file in the `/etc` directory.

Use the replacement term `ss3000` with the `tip` command.

Using the `tip` Command for Local Access to the Array

You can access the RAID COM port (DB9 connector on the controller module) locally with the following commands:

3. Connect the RAID array COM port to the serial port of a Solaris workstation (see FIGURE 2-1)

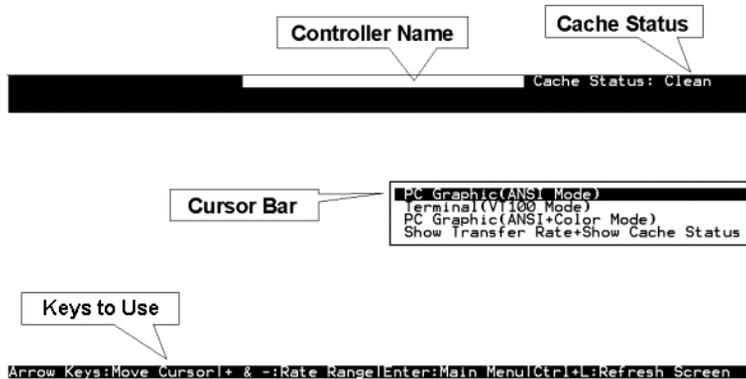
```
tip ss3000 /dev/ttyb
cntrl-1 (to refresh the screen)
```

2.1.2

The Controller Firmware Initial Screens

You will see the following initial controller screen when you first access the RAID controller firmware (via the controller COM port of Ethernet port).

To complete the connection to your management console, select the VT100 terminal mode or the appropriate mode for your communications software, and press Return.



Cursor Bar:	Move the cursor bar to a desired item, then press Return to select.
Controller Name:	Identifies the type of controller.
Transfer Rate Indicator	Indicates the current data transfer rate.
Gauge Range:	Use + or - keys to change the gauge range in order to view the transfer rate indicator.
Cache Status:	Indicates the current cache status.
PC Graphic (ANSI Mode):	Enters the Main Menu and operates in ANSI mode
(VT-100 Mode):	Enters the Main Menu and operates in VT-100 mode.
PC Graphic (ANSI+Color Mode):	Enters the Main Menu and operates in ANSI color mode.
Show Transfer Rate+Show Cache Status:	Press Return on this item to show the cache status and transfer rate.

2.1.3 Main Menu

After you have selected the mode on the initial screen and pressed Return on the initial screen, the Main Menu is displayed.

```
< Main Menu >
Quick installation
view and edit Logical drives
view and edit Logical Volumes
view and edit Host Luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
system functions
view system information
view and edit Event logs
```

Use the arrow keys to move the cursor bar through the menu items, then press Return to choose a menu, or Esc key to return to the previous menu/screen.

2.1.3.1 Navigate the Main Menu

Use the arrow keys to move the cursor bar through the menu items, then press Return to choose a menu, or ESC to return to the previous menu/screen.

- ← → ↑ ↓ To select options
- Return To go to a submenu or to execute a selected option
- [ESC] To escape and go back to the previous menu
- [Ctrl] [L] To refresh the screen information

2.1.3.2 Quick Installation

RESERVED for special use.



Caution – DO NOT USE this command unless directed by Technical Support.

2.2 Upgrading Firmware

For firmware upgrade procedures, refer to Chapter 8 in the *Sun StorEdge 3310 SCSI Array Installation, Operation, and Service Manual*.

Viewing and Editing Logical Drives

This chapter explains how to view and edit Logical drives. Topics covered include:

- “Introducing the Logical Drive Commands” on page 3-2
- “Default Logical Drives and RAID Levels” on page 3-3
- “Viewing the Logical Drive Status Table” on page 3-4
- “Creating Logical Drive(s)” on page 3-5
- “Changing a Logical Drive Controller Assignment” on page 3-11
- “Viewing the Logical Drive Status Table” on page 3-4
- “Partitioning a Logical Drive” on page 3-13
- “Deleting a Logical Drive” on page 3-16
- “Deleting a Partition of a Logical Drive” on page 3-17
- “Assigning Logical Drive Name” on page 3-18
- “Rebuilding a Logical Drive” on page 3-19
- “Performing a Logical Drive Parity Check” on page 3-20
- “Adding a SCSI Drive to a Logical Drive” on page 3-21
- “Copying and Replacing Drives with Drives of Larger Capacity” on page 3-25
- “Expanding a Logical Drive” on page 3-28

3.1 Introducing the Logical Drive Commands

On the Main Menu, the “view and edit Logical drives” command is used to stripe physical SCSI drives into a logical drive.

```
----- < Main Menu > -----
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
system Functions
view system Information
view and edit Event logs
```

For additional commands, highlight the first line or the desired logical drive (LG), and press Return. The following menu is displayed:

LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
S1	52D800AC	NA	RAID5	11000	GOOD	S	3	0	0	

View scsi drives
Delete logical drive
Partition logical drive
logical drive Name
logical drive Assignments
Expand logical drive
add Scsi drives
reGenerate parity
c0py and replace drive

Under this menu command, you can create multiple logical drives. Once logical drives are created, you can manipulate them by selecting the desired LG and pressing Return.

To create logical drives, you must map the logical drives to a host channel with the “view and edit Host luns” command. If this step is skipped, the logical drives will not be visible to the host.

3.2 Default Logical Drives and RAID Levels

A logical drive is a set of drives grouped together to operate under a given RAID level. Each controller is capable of supporting as many as 8 logical drives. The logical drives can have the same or different RAID levels.

- A logical drive can be further divided into a maximum of 32 partitions.
- The total number of partitions must not exceed 128 partitions per RAID array. To create a maximum of 128 partitions, refer to “Planning for 128 LUNs” on page 5-1.

For a 12-drive array, the RAID array will be pre-configured as follows:

- one RAID 5 logical drive, consisting of four physical drives
- one RAID 5 logical drive, consisting of five physical drives
- one global spare

For more information about spares and RAID levels, refer to “RAID Levels” on page 1-3 and “Local and Global Spare Drives” on page 1-10.

3.3 Viewing the Logical Drive Status Table

To check and configure logical drives, from the Main Menu select “view and edit Logical drives,” and press Return. The following screen displays the status of all logical drives.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
V	P0	64312D6F	NA	RAID0	208482	GOOD	S	6	-	0	
V	S1	76605A49	NA	RAID0	208482	GOOD	S	6	-	0	
V	2			NONE							
V	3			NONE							

TABLE 3-1 Parameters Displayed in the Logical Drive Status Window

Parameter	Description
LG	Logical drive number P0 : Logical drive 0 of the primary controller where P = primary controller and 0 = logical drive number S1 : Logical drive 1 of the secondary controller where s = secondary controller and 1 = logical drive number
ID	Logical drive ID number (controller-generated).
LV	The logical volume to which this logical drive belongs. NA indicated no logical volume.
RAID	RAID level
SIZE (MB)	Capacity of the logical drive in megabytes.
Status	Logical drive status.
INITING	The logical drive is now initializing.
INVALID	The logical drive was improperly created or modified. For example, the logical drive was created with “Optimization for Sequential I/O,” but the current setting is “Optimization for Random I/O.”
GOOD	The logical drive is in good condition.
DRV FAILED	A drive member failed in the logical drive.
FATAL FAIL	More than one drive member in a logical drive has failed.
REBUILDING	The logical drive is rebuilding.

TABLE 3-1 Parameters Displayed in the Logical Drive Status Window (Continued)

Parameter	Description
DRV ABSENT	One of the disk drives cannot be detected.
INCOMPLETE	Two or more member disk drives in this logical drive have failed.
O	Indicates the performance optimization set when the logical drive was initialized. This cannot be changed after the logical drive is created. S Optimization for Sequential I/O R Optimization for Random I/O
#LN	Total number of drive members in this logical drive.
#SB	Number of standby drives available for the logical drive. This includes local spare and global spare disk drives available for the logical drive.
#FL	Number of failed disk drive member(s) in the logical drive.
Name	Logical drive name (user configurable)

3.4 Creating Logical Drive(s)

The RAID array is already configured with one or two RAID 5 logical drives and one global spare. Each logical drive consists of a single partition by default.

This procedure is used to modify the RAID level and to add more logical drives, if necessary. In this procedure, you configure a logical drive to contain one or more hard drives based on the desired RAID level, and partition the logical drive into additional partitions.

Note – If you want to assign 128 partitions to 128 LUNs in an array, you need to have a minimum of four logical drives (each with 32 partitions).

For redundancy across separate channels, you can also create a logical drive containing drives distributed over separate channels. You can then partition the logical unit into one or several partitions.

A **logical drive** consists of a group of SCSI drives. Each logical drive can be configured a different RAID level.

A drive can be assigned as the *local spare drive* to one specified logical drive, or as a *global spare drive* that is available to all logical drives on the RAID array. Spares can be part of automatic array rebuild. A spare is not available for logical drives with no data redundancy (NRAID and RAID 0).

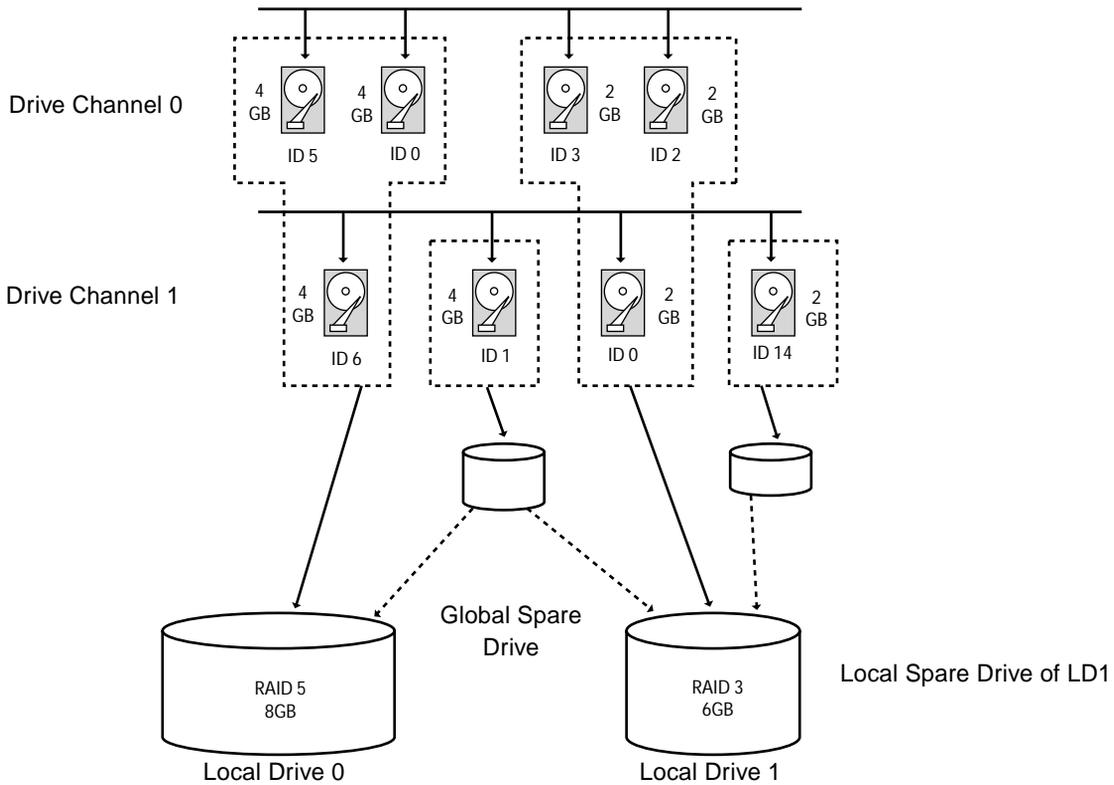


FIGURE 3-1 Allocation of drives in Logical Configurations

1. View the connected drives.

Before configuring disk drives into a logical drive, it is necessary to understand the status of physical drives in your enclosure.

- a. Use the arrow keys to scroll down to “view and edit Scsi drives” and press Return.

This displays information of all the physical drives that are installed.

Quick view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
view		0	0	70007	160MB	0	ON-LINE	SEAGATE ST373405LC
view		0	1	70007	160MB	0	ON-LINE	SEAGATE ST373405LC

- b. Use the arrow keys to scroll through the table. Check that all installed drives are listed here.

If a drive is installed but is not listed, it might be defective or might not be installed correctly, contact your RAID supplier.

When the power is on, the controller scans all hard drives that are connected through the drive channels. If a hard drive was connected after the controller completed initialization, use the “Scan scsi drive” function accessed with the “view and edit scsi Drives” command on the Main Menu, to let the controller recognize the newly added hard drive and configure it as a member of a logical drive.



Caution – Scanning an existing drive removes its metadata references and its assignment to any logical drive. All data on that drive will be lost.

2. Create a logical drive.

a. Scroll through the Main Menu and select “view and edit logical drive.”

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
v	0			NONE							
v	1			NONE							
v	2			NONE							
v	3			NONE							
s	4			NONE							
v	5			NONE							
	6			NONE							
	7			NONE							

b. Select the first available unassigned logical drive (LG) and press Return to proceed.

You can create as many as eight logical drives from drives on any SCSI bus.

When prompted to “Create Logical Drive?” select Yes and press Return to proceed.

A pull-down list of supported RAID levels is displayed.

3. Select a RAID level for this logical drive.

RAID 5 is used in the following example screens



For brief descriptions of RAID levels, refer to “Default Logical Drives and RAID Levels” on page 3-3. For more information about RAID levels, refer to Chapter 1 in the *Sun StorEdge 3310 SCSI RAID Firmware User Guide*.

4. Select your member drive(s) from the list of available physical drives.

The drives can be tagged for inclusion by positioning the cursor bar on the drive and then pressing Return. An asterisk (*) mark is displayed on the selected physical drive(s).

To deselect the drive, press Return again on the selected drive. The “*” mark will disappear.

Note – You must select at least the minimum number of drives required per RAID level. Each RAID level requires an even number of drives.

a. Use the up and down arrow keys to select more drives.

LG	ID	LU	RAID	Size<MB>	Status	o#	RAID 5	NAME
0	Slot	Ch1	ID	Size<MB>	Speed	LG_DRU	Status	Vendor and Product ID
1		2	0	9999	40MB	NONE	NEW DRU	SEAGATE ST31055W
2		2	1	9999	40MB	NONE	NEW DRU	SEAGATE ST31055W
3		2	2	9999	40MB	NONE	NEW DRU	SEAGATE ST31055W
4		2	4	9999	40MB	NONE	NEW DRU	SEAGATE ST31055W
5			NONE					
6			NONE					
7			NONE					

b. After all physical drives have been selected for the logical drive, press Esc key to continue to the next option.

After member physical drives are selected, a list of selections is displayed.

Maximum Drive Capacity :	9999MB
Assign Spare Drives	
Logical Drive Assignments	

5. Optionally set Maximum Physical Drive Capacity and assign spares.

a. Optionally select “Maximum Drive Capacity” from the above menu, and press Return.

Note – Changing the maximum drive capacity will reduce the size of the logical drive and leave some disk space unused.

Maximum Available Drive Capacity(MB):	9999
Maximum Drive Capacity(MB)	: 9999

As a rule, a logical drive should be composed of physical drives with the same capacity. A logical drive can only use the capacity of each drive up to the maximum capacity of the smallest drive.

b. Optionally add a local spare drive from the list of unused physical drives.

Note – A global spare cannot be created while creating a logical drive.

The spare chosen here is a local spare and will automatically replace any failed disk drive in this logical drive. The local spare will not be available for any other logical drive.

LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
P0	5794866F	NA	RAID1	3500	GOOD	S	2	0	0	
S1	7F5353A	NA	RAID1	3500	GOOD	S	2	0	0	
P2	54568151	NA	RAID1	3500	GOOD	S	2	0	0	
S3	2E7C1FD4	NA	RAID1	3500	GOOD	S	2	0	0	Scrunchy_4
4			NONE							
Maximum Drive Capacity : 34476MB Assign Spare Drives Disk Reserved Space: 256 MB Logical Drive Assignments										

LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
P0	5794866F	NA	RAID1	3500	GOOD	S	2	0	0	
S1	7F5353A	NA	RAID1	3500	GOOD	S	2	0	0	
P2	54568151	NA	RAID1	3500	GOOD	S	2	0	0	
S3	2E7C1FD4	NA	RAID1	3500	GOOD	S	2	0	0	Scrunchy_4
4			NONE							
	Slot	Ch1	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID		
		0	12	34732	160MB	NONE	FRM1_DRV	SEAGATE ST336605LSUN36G		
		0	13	34732	160MB	NONE	NEW DRV	SEAGATE ST336605LSUN36G		

Note – A logical drive created in a RAID level which has no data redundancy (NRAID or RAID 0) does not support spare drive rebuilding.

6. Optionally select “Logical Drive Assignments” to assign this logical drive to the secondary controller.

By default, all logical drives are automatically assigned to the primary controller.

Redundant Controller Logical Drive Assign to Secondary Controller ?
<input checked="" type="radio"/> Yes
<input type="radio"/> No

If you use two controllers for the redundant configuration, a logical drive can be assigned to either of the controllers to balance workload. Logical drive assignment can be changed any time later.

a. Press Esc key or No to exit from this window without changing the controller assignment.

- b. Select “Yes”, press Return to confirm, and then press Esc key to continue when all the preferences have been set.

A confirmation box is displayed on the screen.

- c. Verify all information in the box before choosing “Yes” to proceed.

```
Raid Level           : RAID 5
Online SCSI Drives  : 4
Maximum Drive Capacity : 9999 MB
Spare SCSI Drives   : 1
Logical Drive Assignment: Primary Controller

Create Logical Drive ?
  Yes                No
```

A message will indicate that the logical drive initialization has begun.

- d. Press Esc key to cancel the “Notification” prompt.
 - e. After the logical drive initialization is completed, use the Esc key to return to the Main Menu.
7. Select “view and edit Logical drives” to view details of the created logical drive.

3.5 Changing a Logical Drive Controller Assignment

By default, logical drives are automatically assigned to the primary controller. If you assign half the drives to the secondary controller, the maximum speed and performance will be somewhat improved due to the redistribution of the traffic.

To balance the workload between both controllers, you can distribute your logical drives between the primary controller (displayed as the Primary ID or PID) and the secondary controller (displayed as the Secondary ID or SID).

After a logical drive has been created, it can be assigned to the secondary controller. Then the host computer associated with the logical drive can be mapped to the secondary controller (see “Mapping Logical Drive Partitions to Host LUNs” on page 5-3).

1. To change a logical drive controller assignment, select “view and edit Logical drives” from the Main Menu, and press Return.

2. Select “logical drive assignments” and press Return.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
v	P0	6281CB39	NA	RAID5	200	GOOD	S	3	0	0	
v											
v											
v											
v											
v											
v											
S											
v											
	View scsi drives										
	Delete logical drive										
	Partition logical drive										
	logical drive Name										
	logical drive Assignments										
	Redundant Controller Logical Drive Assign to Secondary Controller ?										
	Yes <input type="checkbox"/> No <input checked="" type="checkbox"/>										
	6			NONE							
	7			NONE							

The reassignment is evident from the “view and edit Logical drives” screen.

A “P” in front of the LG number means that the logical drive is assigned to the primary controller. An “S” in front of the LG number means that the logical drive is assigned to a Secondary controller.

For example, “S1” indicates that logical drive 1 assigned to the secondary controller.

Note – The editable logical drive NAME name is used only in the RAID firmware administration and monitoring, and does not appear anywhere on the host. You can create a logical drive NAME after the logical drive is created: Select the logical drive in the above screen, and press Return. Then select “logical drive name,” “enter the desired name, and press Return.

3. Press Esc to return to the Main Menu.

4. Reset the controller to implement the change.

Select “system Functions” on the Main Menu. Then select “Reset Controller” and press Return.

You must reset the controller for the changes to take effect.

3.6 Partitioning a Logical Drive

You can divide a logical drive into several partitions, or use the entire logical drive as a single partition. You might configure up to 32 partitions for each logical drive.

For guidelines on setting up 128 LUNs, refer to “Mapping Logical Drive Partitions to Host LUNs” on page 5-3.



Caution – If you modify the size of a partition or logical drive, you lose all data on those drives.

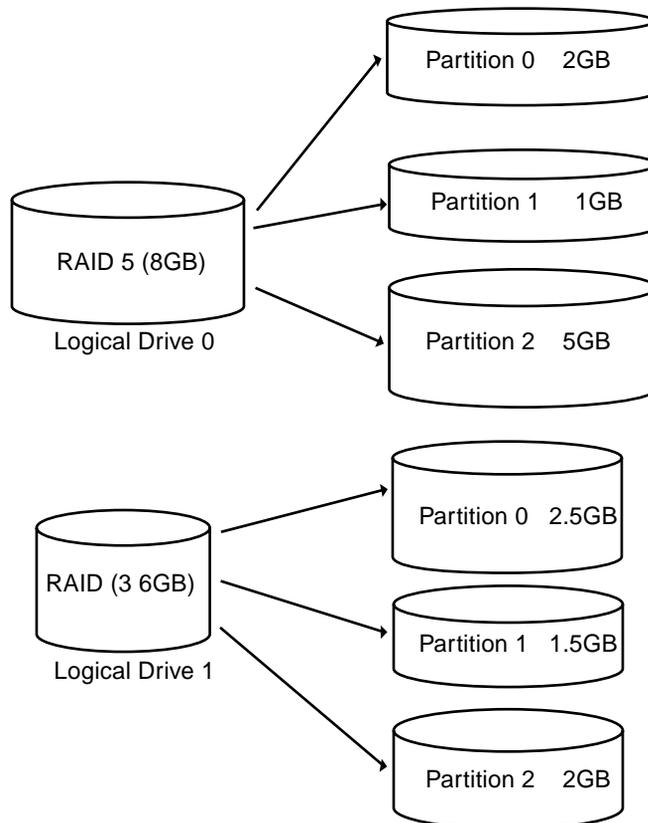


FIGURE 3-2 Partitions in Logical Configurations

To partition a logical drive, perform the following steps:

1. From the Main Menu, select “view and edit Logical drives.”
2. Select the logical drive you want to partition, then press Return

Q	LG	ID	LU	RAID	Size<MB>	Status	O	#LN	#SB	#FL	NAME
U	P0	B61E5AB	NA	RAID5	3000	GOOD	S	3	5	0	
U	1			NONE							
U	2			NONE							

3. Select “Partition logical drive” from the menu, then press Return.

Q	LG	ID	LU	RAID	Size<MB>	Status	O	#LN	#SB	#FL	NAME
U	P0	B61E5AB	NA	RAID5	3000	GOOD	S	3	5	0	
U	View scsi drives Delete logical drive Partition logical drive logical drive Name logical drive Assignments Expand logical drive add Scsi drives reGenerate parity cOPY and replace drive										
U	6			NONE							
U	7			NONE							

4. Select from the list of undefined partitions and press Return.

The following prompt then is displayed: “Partitioning the Logical Drive will make it no longer eligible for membership in a logical volume. Continue Partition Logical Drive?”

5. Select Yes.

A list of the partitions for this logical drive is displayed. If the logical drive has not yet been partitioned, all the logical drive capacity will list as “partition 0.”

6. Press Return and type the desired size for the selected partition, and then press Return to proceed.

Q	LG	ID	LU	RAID	Size<MB>	Partition	Offset<MB>	Size<MB>	NAME
U	P0	B61E5AB	NA	RAID5	3000	0	0	3000	
U	1			NONE					
U	2			NONE					
U	Partition Size <MB>: 3000								

A warning prompt is displayed: “This operation will result in the loss of all data on the partition! Partition Logical Drive?”

7. Select Yes.

8. Press Return to confirm.

The remaining capacity of the logical drive will be automatically allotted to the next partition. In the following figure, a partition size of 3000MB was entered; the remaining storage of 27000 MB is allocated to the partition below the partition created.

Q	LG	ID	LU	RAID	Size<MB>	Partition	Offset<MB>	Size<MB>	NAME
U	P0	B61E5A8	NA	RAID5	30000	0	0	3000	
U	1			NONE		1	3000	27000	
U	2			NONE		2			
U	3			NONE		3			

9. Repeat the preceding steps to partition the remaining capacity of your logical drive.

You can create up to 32 partitions per logical drive, with a total number of partitions not to exceed 128 partitions/LUNs per the RAID array.

Note – When you modify a partition or logical drive size, you must re-configure all host LUN mappings. All the host LUN mappings are removed with any change to partition capacity. See “Mapping Logical Drive Partitions to Host LUNs” on page 5-3.

Note – When a partition of logical drive/logical volume is *deleted*, the capacity of the deleted partition is added to the partition above the deleted partition.

3.7 Deleting a Logical Drive

You can keep the default logical drive configuration, or create new logical drives and different RAID levels with the following procedure. To re-configure the entire array with different logical drives and RAID levels, you must delete the existing logical drive(s) and re-configure the array prior to storing data on the array.

Note – We highly recommend that you keep the default configuration, which provides the maximum storage protection.

To delete a logical drive, perform the following steps.

1. **On the Main Menu, select “view and edit Logical drives.”**
2. **Highlight the line with the logical drive and press Return.**
3. **Select “Delete logical drive” and press Return.**



Caution – Deleting a logical drive will destroy data.

4. **Select Yes when prompted to confirm.**

3.8 Deleting a Partition of a Logical Drive

Q	LG	ID	LV	RAID	Size(MB)	Partition	Offset(MB)	Size(MB)	NAME
v	P0	4149A729	NA	RAID5	39	0	0	3999	
v	P1	76CD4DF6	NA	RAID0	119	1	3999	3999	
v	2			NONE					
v	3			NONE					
v	4			NONE		4	15999	3999	
v	5			NONE		5			
v	6			NONE		6			
v	7			NONE		7			

To delete a partition, perform the following steps.

1. Select “view and edit Logical drives” from the Main Menu and press Return.
2. Select the logical drive which has a partition you want to delete, then press Return.
3. Select “Partition logical drive.” Partitions of the logical drive will be displayed in tabulated form.
4. Select the partition you want to delete, then press Return. Return “0” on the partition size to delete this partition.

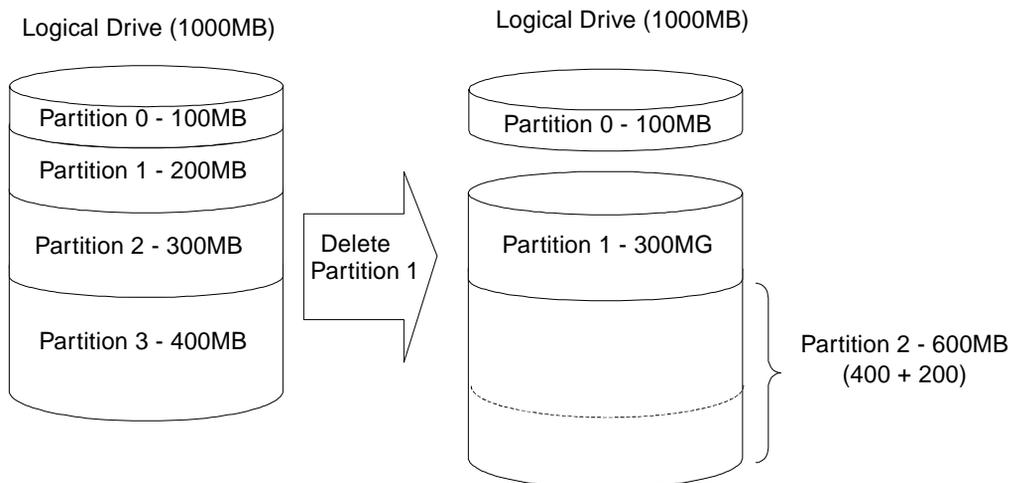


FIGURE 3-3 Example of Deleted Partitions

3.11 Performing a Logical Drive Parity Check

If no verifying method is applied to data writes, this function can be performed periodically to ensure bad sectors will not cause data loss in the event of drive failure. In a RAID unit, data is striped across multiple member drives and this function can regenerate parity and prevent bad blocks on drives.

LG	ID	LV	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
P0	4149A729	NA	RAID5	19998	GOOD	R	3	0	0	
P1	560E15F8	NA	RAID5	19998	GOOD	R	3	0	0	
View scsi drives										
Delete logical drive										
Regenerate Logical Drive Parity ?										
Yes No										
regenerate parity										
copy and replace drive										
7			NONE							

To check logical drive parity, perform the following steps.

1. Select “view and edit Logical drives” from the Main Menu and press Return.
2. Select the logical drive that you want to regenerate the parity for, and then press Return.
3. Select “reGenerate parity,” then press Return.
4. When prompted to “Regenerate Parity?” select Yes.

Note – If a regenerating process is stopped by a drive failure, the process cannot restart until logical drive rebuild has been completed.

3.12 Adding a SCSI Drive to a Logical Drive

RAID expansion allows users to expand a logical drive by adding new drives, or by copying the data from the original member drives to the new drives, and then replacing the original member drives without powering down the array.

Guidelines

- Expansion can only be performed on RAID 0, 3 and 5 logical drives. Expansion cannot be performed on an NRAID or RAID 1 logical drive.
- **When you add a drive to a logical drive, you add a partition** (the size of the new drive) to the logical drive. Therefore, if you have a single 200GB logical drive and add a 36GB drive, the total logical drive will be 236GB with two partitions (one 200GB and one 36GB).
- **All SCSI drives in a logical drive must be the same size**, namely, 36GB or 73GB drives.
- **The new partition must be mapped to a host LUN** in order for the HBA (host-bus adapter) to recognize its presence. If you want to add the new partition into an existing partition, operating environment support is necessary.
- Expanding logical drives by adding more SCSI hard disk drives cannot be canceled once begun. If a power failure occurs, the expansion will be paused and the controller will not automatically continue the expansion when the power comes back on. Restarting of the RAID expansion must be performed manually.
- If a member drive of the logical drive fails during RAID expansion, the expansion will be paused. The expansion will resume automatically after logical drive rebuild has been completed.

After new drives are added, the capacity of the original logical drive will be the same and the additional capacity is displayed as another partition (new partition).

When expansion is completed, data will be re-striped across the original and the newly added drives. After adding SCSI drives, the additional capacity is displayed as a new partition (see FIGURE 3-4).

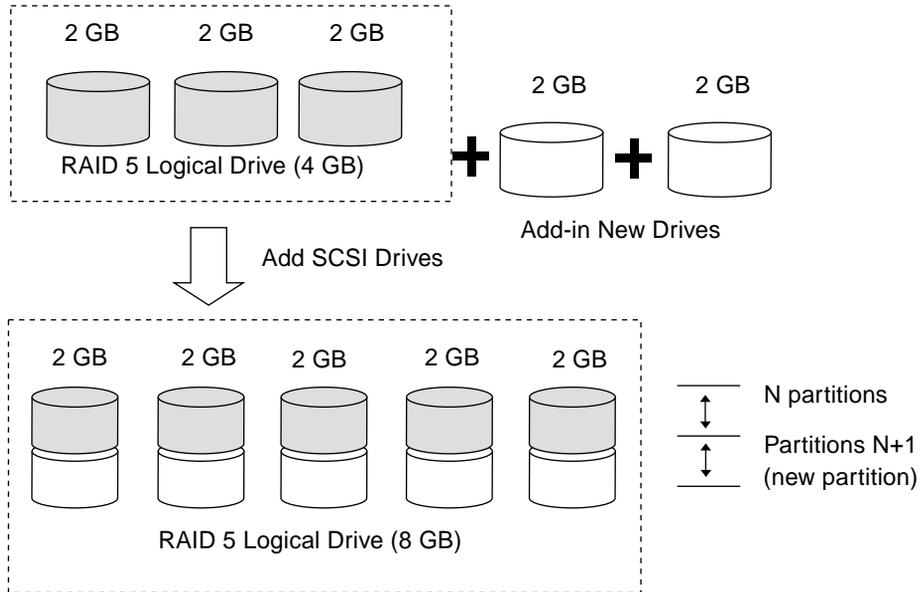


FIGURE 3-4 Expansion by Adding Drive

RAID levels supported: RAID 0, 3, and 5.

To add a SCSI drive to a logical drive, perform the following steps.

- 1. From the Main Menu, select “view and edit logical drive,” and highlight a logical drive.**

The drive selected for adding a SCSI drive should have a capacity no less than the original member drive. If possible, use drives of the same capacity because all drives in the array are treated as though they have the capacity of the smallest member drive in the logical array.

- Press ESC again to cancel the notification prompt. A status bar indicates the percentage of progress.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
v	P0	2E5B167A	NA	RAID5	9999	GOOD	R	3	0	0	
v	1	Adding									
v	2	[Progress Bar]									
v	3	11% Completed									
s	4		NONE								
v	5		NONE								

Upon completion, a confirming notification message is displayed:

“Add SCSI Drive to Logical Drive 0 Complete”

The capacity of the added drive is displayed as an unused partition.

< Main Menu >											
Quick installation	LUN	LV/LD	DRV	Partition	Size(MB)	RAID					
view and edit Logical drives	0	LD	0	0	9999	RAID5					
view and edit logical Volumes	1										
view and edit Host luns											
v	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
v	P0	34456224	NA	RAID5	14998	GOOD	S	4	0	0	
v	Partition	Offset(MB)	Size(MB)								
s	0	0	9999								
v	1	9999	4999								
view a	7										

The added capacity will be included in the logical drive automatically, meaning that you do not have to perform “Expand logical drive” later.

- However, if the logical drive has already been mapped with a host LUN, you have to map the added capacity to another host ID/LUN to make use of it.

In the above example of “view and edit Host luns,” the original capacity is 9999MB, its host LUN mapping remains unchanged and the added capacity is displayed as the second partition.

3.13 Copying and Replacing Drives with Drives of Larger Capacity

You can expand a logical drive by copying and replacing all member drives with drives of higher capacity. Refer to the diagram below. The capacity of member drives are copied and replaced one by one onto drives of larger capacity.

When all the member drives have been replaced, execute the “Expand logical drives” function to make use of the unused capacity.

Note – *The new created area will be a new partition.*

RAID levels supported: RAID 0, 3, and 5.

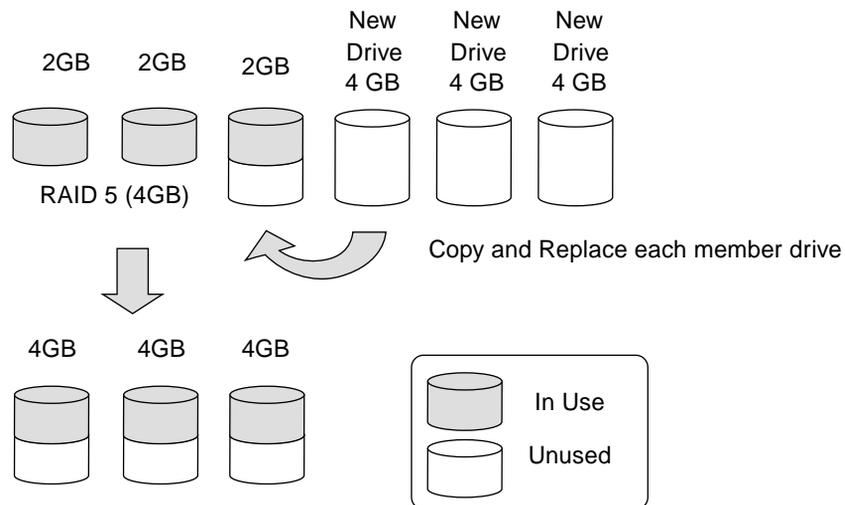


FIGURE 3-5 Expansion by Copy & Replace

To copy and replace drives with drives of larger capacity, perform the following steps.

1. From the Main Menu, select “view and edit Logical drives.”
2. Select a target logical drive, press Return and scroll down to choose “copy and replace drive.” Press Return to proceed.

5. Select Yes to confirm and proceed.

A notification message is displayed:

[21A1] LG:0 Logical Drive NOTICE: CHL:1 ID:3 Starting Clone

6. Press ESC to view the progress.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
v	P0	6D15A60	NA	RAID5	9999	GOOD	R	4	0	0	
v						Drive Copying					
v											
v											
v											
v											
s											
v						40% Completed					

7. Completion of the Copy and Replace process will be indicated by a notification message.

A notification message is displayed:

[21A2] LG:0 Logical Drive NOTICE: CHL:1 ID:3 Copy and Replace Completed

8. Follow the same method to copy and replace every member drive with drives of higher capacity as needed.

You may now perform “Expand logical drive” to make use of the capacity brought by the new drives and then map the additional capacity to a Host LUN.

Drive Copying Notification
[21A2] LG:0 Logical Drive NOTICE:CHL:1 ID:3 Copy and Replace Completed

3.14 Expanding a Logical Drive

The user can make use of unused capacity in a logical drive by expanding the logical drive.

An unused capacity is usually created by replacing the original members with drives of larger capacity; or, by adding new drive(s) to a logical drive. After expanding a logical drive, the additional capacity is displayed as another partition (a new partition). FIGURE 3-6 illustrates this idea.

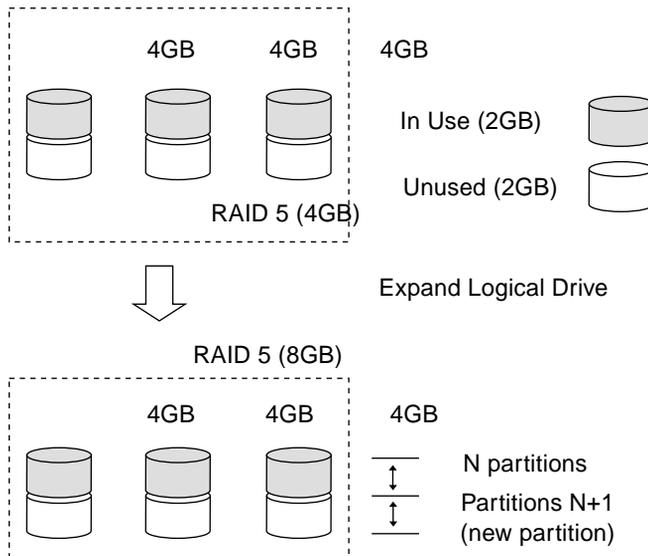


FIGURE 3-6 Logical Drive Expansion

Note – The new created area will be a new partition.

RAID levels supported: RAID 0, 1, 3, and 5.

The new partition must be mapped to a host LUN in order for the HBA (host-bus adapter) to recognize its presence. To add the new partition into an existing logical drive, operating environment support is needed.

In the following example, the logical drive is originally composed of three member drives and each member drive has the capacity of 1 Gigabyte. The “Copy and Replace” function has been performed on the logical drive and each member drive has been replaced by a new drive with the capacity of 2 Gigabytes. The next step is to perform “Expand logical drive” function to utilize the additional capacity brought by the new drives.

1. Select “view and edit Logical drives” from the Main Menu, select the logical drive with its members copied and replaced, and press Return.
2. Select “Expand logical drive” in the submenu and press Return to proceed. A confirming box is displayed.
3. Proceed by pressing Return or entering any value no larger than the “maximum drive expand capacity” and press Return.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
V	P0	0499A/C9	NA	RAID0	3000	GOOD	R	3	-	0	
V											
V											
V											
V											
S											
V											
V											
V											
a											
c											
6											
7				NONE							

View scsi drives
 Delete logical drive
 Partition logical drive
 logical drive Name
 logical drive Assignments
Expand logical drive

Maximum Available Drive Free Capacity: 1000MB
 Maximum Drive Expand Capacity(MB) : 1000_

4. Select Yes to confirm and proceed.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
V	P0	0499A/C9	NA	RAID0	3000	GOOD	R	3	-	0	
V											
V											
V											
V											
S											
V											
V											
a											
c											
6											
7				NONE							

View scsi drives
 Delete logical drive
 Partition logical drive
 logical drive Name
 logical drive Assignments
Expand logical drive

Expand Logical Drive ?
 Yes **No**

5. Upon completion, you will be prompted by the notification message:
[2188] Expansion of Logical Drive 0 Completed

6. Press ESC to return to the previous menu screen.

The total capacity of logical drive has been expanded to 6 Gigabytes.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
v	P0	0499A/C9	NA	RAID0	6000	GOOD	R	3	-	0	
v	1			NONE							
v	2			NONE							
v	3			NONE							
s	4			NONE							
v	5			NONE							
v	6			NONE							
v	7			NONE							

Viewing and Editing Logical Volumes

The creation of logical volumes is a legacy function, which is available but rarely used. This chapter describes the limitations and procedures for creating and using logical volumes.



Caution – This feature is not supported in the Configuration Service program. It is available but rarely used. It has been replaced by the use of partitions in logical drives.

Topics covered in this chapter include:

- “Understanding Logical Volumes (Multi-Level RAID)” on page 4-2
- “Creating a Logical Volume” on page 4-5
- “Expanding a Logical Volume” on page 4-7

4.1 Understanding Logical Volumes (Multi-Level RAID)

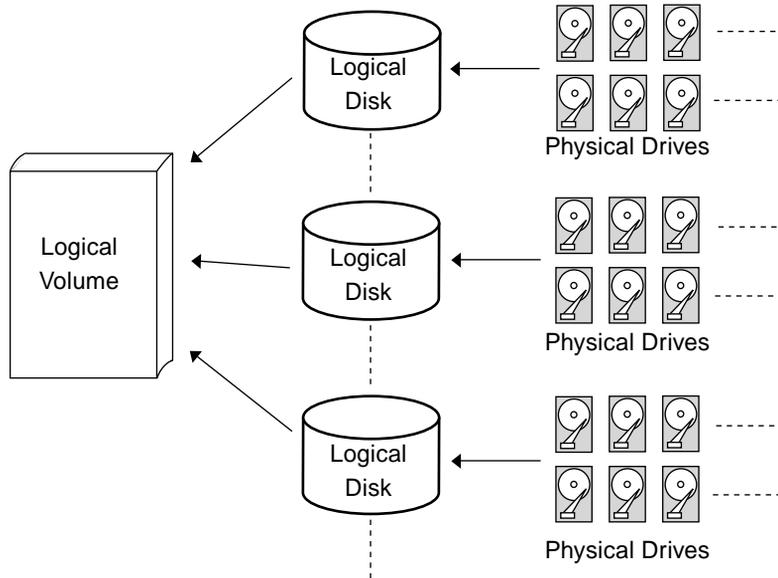


FIGURE 4-1 Logical Volume composed of Multiple Drives

A logical volume is a combination of RAID 0 (Striping) and other RAID levels. Data written to a logical volume is first broken into smaller data segments and striped across different logical drives in a logical volume. Each logical drive then distributes data segments to its member drives according to its mirroring, parity, or striping scheme.



Caution – The logical volume feature is not supported in the Configuration Service program. It is available but rarely used. It has been replaced by the use of partitions in logical drives.

A logical volume can be divided into a maximum of eight partitions. The preferred function is the creation of partitions into logical drives. You may configure up to 32 partitions for each logical drive, and up to 128 partitions total within an array. During normal operation, the host sees a non-partitioned logical volume or a partition of a partitioned logical volume as one single physical drive.

Note – Logical drives with many more partitions serve the same purpose.

■ Spare drives in a logical volume?

A local spare cannot be assigned to a logical volume. If a drive fails, it fails as a member of a logical drive; therefore, the controller allows local spare assignment to logical drives rather than logical volumes.

■ Logical volume limitations?

The logical volume cannot have any logical drive with a status of “fatal failed.” If there is a failed drive in any of the member logical drives (of a logical volume), the controller will start to rebuild that logical drive.

If any of the member logical drives fails fatally, the logical volume fails fatally and data will not be accessible.

To avoid a logical volume failure:

- Logical drives as members to a logical volume should be configured in RAID levels that provide redundancy, namely, RAID level 1 (0+1), 3, or 5.
 - Rebuild the logical drive as soon as possible whenever a drive failure occurs.
 - A logical drive should be composed of physical drives from different drive channels. Compose the logical drive with drives from different drive channels to avoid the fatal loss of data caused by bus failure.
- Partitioning the logical drive or partitioning the logical volume?

Once a logical drive has been divided into partitions, the logical drive can no longer be used as a member of a logical volume. The member logical drives of a logical volume should have one partition only with the entire capacity.

If you want to use a partitioned logical drive for a logical volume, delete the other partitions in this logical drive until there remains one partition only with the entire logical drive capacity.



Caution – Deleting the partition of the logical drive will also destroy all the data. Data should be backed up to somewhere else before making partition configuration.

When a logical drive is used as a member to a logical volume, this logical drive can no longer be partitioned in “View and Edit Logical Drives.” Instead, the Logical Volume can be partitioned in “View and Edit Logical Volume.”

The procedure for partitioning a logical volume is the same as that for partitioning a logical drive. After the logical volume has been partitioned, map each partition to a host ID/LUN to allow the host computer to utilize the partitions as individual drives.

- RAID expansion with logical volume?

The Logical Volume can also be expanded using the RAID expansion function. The concept of expanding a logical volume is similar to that of expanding a logical drive. To perform RAID expansion on a logical drive, replace each member physical drive with a drive of larger capacity or add a new drive, then perform logical drive expansion to utilize the newly-added capacity.

To perform RAID expansion on a logical volume, expand each member logical drive, then perform RAID expansion on the logical volume.

- Is there anything changed after using logical volume?

Without a logical volume: Logical drives can be assigned to the primary controller or to the secondary controller. The host I/Os directed to a logical drive will be served by the controller to which this logical drive is assigned. If a controller fails, the host I/Os originally assigned to the failed controller will be taken over by the existing controller. When the controller fails back (failed controller being replaced by a new one), logical drives will be returned to the replacement controller in the original configuration.

With a logical volume: Logical volumes can also be assigned to different controllers. The only difference is that the logical volumes will be considered as the base units for shifting the control during a controller failure.

- A logical volume with logical drives of different levels (Multi-level RAID arrays)?

A multi-level RAID with logical volume support offers the following available applications.

- **RAID (0+1):** This is a standard feature of RAID controllers. It has the benefits of RAID 1 (high availability) and RAID 0 (enhanced I/O performance through striping). Simply choose multiple drives for a RAID 1 logical drive, the RAID controller will implement RAID (0+1) automatically.
- **RAID (3+0):** A logical volume itself is a multi-level RAID implementation. A logical volume is composed of one or several logical drives with data “striping” (RAID 0). A logical volume with several RAID 3 member logical drives can be considered as a RAID (3+0), or RAID 53 as defined in “The RAID Book” (from The RAID Advisory Board).
- **RAID (5+0):** A logical volume with several RAID 5 member logical drives.
- **RAID (5+1):** Requires multiple RAID controllers. In a RAID (5+1) array, each layer-1 RAID controller handles one RAID 5 logical drive and a layer-2 RAID controller performs RAID 1 (mirroring) function to the virtual disks controlled by all of the layer-1 RAID controllers.

- **RAID (5+5):** Requires multiple RAID controllers. In the RAID (5+5) array, each layer-1 RAID controllers handles one to several RAID 5 logical drives and a layer-2 RAID controller performs RAID 5 to the virtual disks provided by all of the layer-1 RAID controllers.
- RAID 10 Logical volume with RAID 1 logical drives.
- RAID 30 Logical volume with RAID 3 logical drives.
- RAID 5 Logical volume with RAID 5 logical drives.

4.2 Creating a Logical Volume

A logical volume consists of one or several logical drives.

To create a logical volume, perform the following steps.

1. Select “view and edit logical Volumes” in the Main Menu.

The current logical volume configuration and status will be displayed on the screen.

q	LU	ID	Size(MB)	#LD	s
v	0				es
v	1				
v	2				parameters
v	3				vices
v	4				
v	5				
v	6				
v	7				

2. Select a logical volume number (0-7) that has not yet been defined, then press Return to proceed.

A prompt “Create Logical Volume?” is displayed.

LV	ID	Size(MB)	#LD
0			
1	Create Logical Volume ?		
2	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> No	
3			
4			
5			
6			
7			

3. Select “Yes” and press Return.

Q	LV	ID	Size(MB)	#LD								
v	0				s							
v	0				es							
v	1	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
v	2	*P0	69FC89B	NA	RAID5	39996	GOOD	R	4	1	0	
v	3	*P1	55A3373B	NA	RAID0	19998	GOOD	R	2	-	0	
v	4											
v	5											
v	6											
v	7											

- Select one or more logical drive(s) available on the list, and press Return to tag the logical drive(s) to be included in the volume.

An asterisk (*) is displayed on the selected drive.

Press Return to deselect a logical drive.

LV	ID	Size(MB)	#LD									
0				s								
0				es								
1	Logical Volume Assignments											
2	Redundant Controller Logical Volume Assign to Secondary Controller ?											
3	Yes						No					

Logical volumes can also be assigned to different controllers (primary or secondary). The default is primary. Press ESC if change is not preferred.

- As all the member logical drives are selected, press ESC to continue.

Logical Drive Count	: 1
Logical Volume Assignment	: Primary Controller
Create Logical Volume ?	
Yes	No

The logical volume creation confirm box is displayed. Select Yes to create the logical volume.

- Press Return and the information of the created logical volume is displayed.

4.3 Expanding a Logical Volume

To expand a logical volume, expand logical drive(s) in the logical volume and then perform “Expand logical volume.”

Q	LV	ID	Size(MB)	#LD	
v	#0	623F7A40	3792	1	s
v					es
v	1	View logical drive			
v		Delete logical volume			
v	2	Partition logical volume			ters
v		logical volume Assignments			
s	3	expand logical volume			
v	4				
v	5				
	6				
	7				

When prompted by “Expand Logical Volume?” select Yes to confirm, and the process will be completed immediately.

4.4 Viewing a Logical Volume Status Table

To check and configure logical drives, from the Main Menu select “view and edit Logical drives,” and press Return. The following screen displays the status of all logical drives.

Caution – This feature is not supported in the Configuration Service program. It is available but rarely used. It has been replaced by the use of partitions in logical drives.

Q	LV	ID	Size(MB)	#LD	
v	P0	466C5C8D	60000	1	s
v	1				es
v	2				
v	3				parameters
v	4				vices
v	5				
v	6				
v	7				

TABLE 4-1 Parameters Displayed in the Logical Volume Status Window

Parameters	Description
LV	Logical volume number P = primary controller S = secondary controller
ID	Logical volume ID number (controller-generated)
Size(MB)	Capacity of the logical volume in megabytes
#LD	The number of logical drive(s) in this logical volume

Viewing and Editing Host LUNs

The “view and edit Host luns” command allows you to map logical groups or logical volumes to the desired host channels. Each LG or LV may be mapped more than once to achieve a redundant data path (additional software is required).

Topics covered in this chapter include:

- “Planning for 128 LUNs” on page 5-1
- “Mapping Logical Drive Partitions to Host LUNs” on page 5-3
- “Deleting a Host LUN Mapping” on page 5-7

5.1 Planning for 128 LUNs

If you want to create 128 LUNs which is the maximum number of storage partitions which can be mapped for a RAID array, you will set up one of the following configurations:

- Create four host IDs. Create four logical drives. Partition each logical drive into 32 partitions (4 times 32 = 128). Map the 128 partitions to the four host IDs. This is the most commonly used configuration.

or

- Create six host IDs (this requires three host drives), perform one of the following steps, and then map the 128 partitions to the 6 host IDs.
 - Create 4 logical drives of 32 partitions each.
 - Create 5 logical drives with total number of partitions equaling 128 (four logical drives with 25 partitions each and one with 28 partitions).
 - Create 6 logical drives (six logical drives with 21 partitions each and one with 23 partitions).

For details on how to add host IDs, refer to “Creating Additional Host IDs” on page 7-9.

Note – For an overview of how partitions, LUNs, and host IDs work, refer to “Mapping Logical Drive Partitions to Host LUNs” on page 5-3.

To set up 128 LUNs, the following steps are required.

1. Create a minimum of four host IDs.

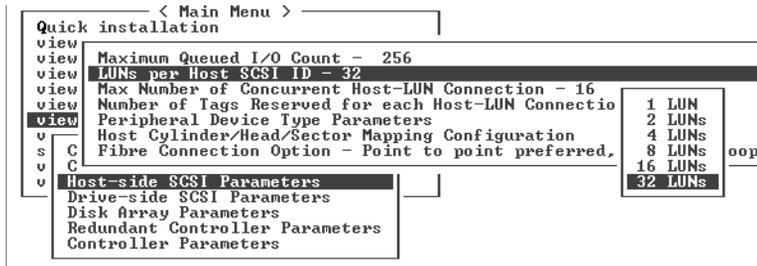
By default, you have two host IDs: Channel 1 ID 0 (primary controller) and Channel 3 ID 1 (secondary controller). You can have a total of two IDs per channel, one for the Primary Controller and one for the secondary controller.

For the detailed procedure, refer to “Creating Additional Host IDs” on page 7-9.

2. Confirm that the allowed number of LUNs per host id is 32.

Go to “view and edit Configuration parameters,” and select “hostside scsi parameters.”

If the “LUNs per Host SCSI ID” is not 32, highlight the line, press Return and select the number 32.



3. Create at least four logical drives.

For the detailed procedure, refer to “Creating Logical Drive(s)” on page 3-5.

4. For each logical drive, create a number of partitions per logical drive until you reach a total of 128 partitions, then map those partitions to the host IDs.

For the detailed procedures, refer to “Partitioning a Logical Drive” on page 3-13 and “Mapping Logical Drive Partitions to Host LUNs” on page 5-3.

5.2 Mapping Logical Drive Partitions to Host LUNs

The next step is to map each storage partition as one system drive (host ID/LUN). The host SCSI adapter will recognize the system drives after re-initializing the host bus.

A SCSI channel (SCSI bus) can connect up to 15 devices (excluding the controller itself) when the Wide function is enabled (16-bit SCSI). Each device has one unique ID.

The figure below illustrates the idea of mapping a system drive to host ID/LUN combination.

- The SCSI ID is like a cabinet, and the drawers are the LUNs (LUN is short for logical unit number).
- Each cabinet (SCSI ID) can have up to 32 drawers (LUNs).
- Data can be stored into one of the LUNs of the SCSI ID. Most SCSI host adapters treat a LUN like another SCSI device.
- The maximum number of LUNs which can be created for a RAID array is 128. To create a total of 128 LUNs, refer to “Planning for 128 LUNs” on page 5-1

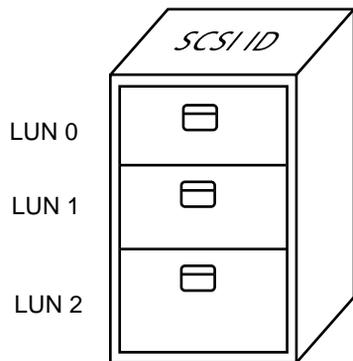


FIGURE 5-1 Filing Cabinet Represents the SCSI ID and File Drawers Represent the LUNs

SCSI ID/LUNs

Each SCSI ID/LUN looks like a storage device to the host computer.

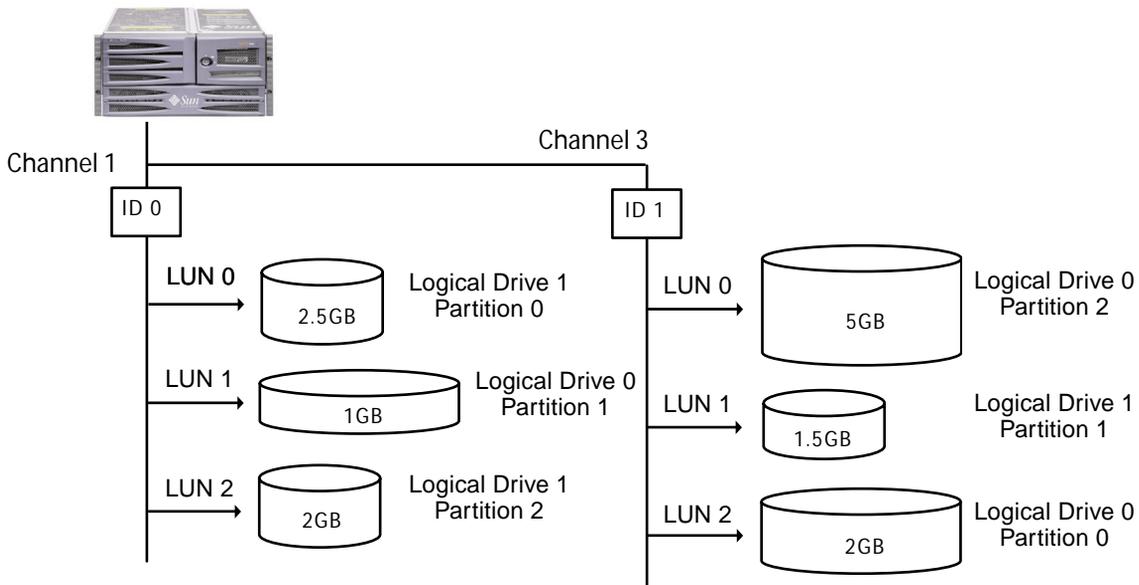
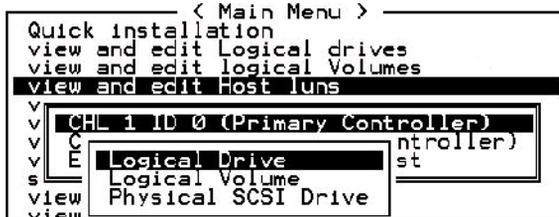


FIGURE 5-2 Mapping Partitions to Host ID/LUNs

To map a logical drive partition to a LUN, perform the following steps.

1. On the Main Menu, select “view and edit Host luns.”
2. Select a specific host-channel ID and press Return. Select a logical drive if prompted.



3. Select a LUN number, and press Return. Select a logical drive, then highlight a partition, and press Return.

LUN	LV/LD	DRV	Partition	Size(MB)	RAID
0					
1					

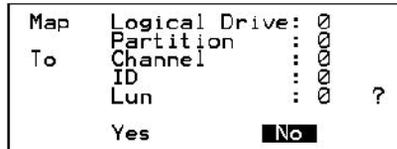
LG	ID	LV	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
P0	62A-2975	NA	RAID5	9998	GOOD	S	3	0	0	

Partition	Offset(MB)	Size(MB)
0	0	597

4. Select “Map Host LUN.”



5. Confirm the mapping scheme by selecting “Yes.”



The same partition might be mapped to multiple LUNs on multiple host channels. This feature is necessary for clustered environments and redundant path environments.

6. Press Esc key to return to the Main Menu.

7. Repeat the above procedure for each partition until all partitions are mapped to a LUN.

8. Select “system Functions” on the Main Menu, and select “Reset Controller” to implement the new configuration settings.

9. To verify unique mapping of each LUN (unique LUN number, unique DRV number, or unique Partition number), select the “view and edit Host luns” command and press Return.

10. Reboot the host(s) to complete the settings, and confirm that LUNs are connected to and visible by the hosts. For additional Solaris operating environment information, refer to Chapter 6 in the *Sun StorEdge 3310 SCSI Array Installation, Operation and Service Manual*.

5.2.1 Example Host LUN Mappings

The following example screens illustrate four channel IDs with 8 LUNs per channel:

< Main Menu >		LUN	LV/LD	DRV	Partition	Size(MB)	RAID
Quick installation		0	LD	0	0	400	RAID5
view and edit Logical drives		1	LD	0	1	400	RAID5
view and edit logical Volumes		2	LD	0	2	400	RAID5
view and edit Host Luns		3	LD	0	3	400	RAID5
v	CHL 1 ID 0 (Primary Controll	4	LD	0	4	400	RAID5
v	CHL 1 ID 1 (Secondary Contro	5	LD	0	5	400	RAID5
v	CHL 3 ID 2 (Secondary Contro	6	LD	0	6	400	RAID5
v	CHL 3 ID 3 (Primary Controll	7	LD	0	7	400	RAID5
v	Edit Host-ID/WWN Name List						

< Main Menu >		LUN	LV/LD	DRV	Partition	Size(MB)	RAID
Quick installation		0	LD	1	0	300	RAID5
view and edit Logical drives		1	LD	1	1	300	RAID5
view and edit logical Volumes		2	LD	1	2	300	RAID5
view and edit Host Luns		3	LD	1	3	300	RAID5
v	CHL 1 ID 0 (Primary Controll	4	LD	1	4	300	RAID5
v	CHL 1 ID 1 (Secondary Contro	5	LD	1	5	300	RAID5
v	CHL 3 ID 2 (Secondary Contro	6	LD	1	6	300	RAID5
v	CHL 3 ID 3 (Primary Controll	7	LD	1	7	300	RAID5
v	Edit Host-ID/WWN Name List						

< Main Menu >		LUN	LV/LD	DRV	Partition	Size(MB)	RAID
Quick installation		0	LD	3	0	350	RAID5
view and edit Logical drives		1	LD	3	1	350	RAID5
view and edit logical Volumes		2	LD	3	2	350	RAID5
view and edit Host Luns		3	LD	3	3	350	RAID5
v	CHL 1 ID 0 (Primary Controll	4	LD	3	4	350	RAID5
v	CHL 1 ID 1 (Secondary Contro	5	LD	3	5	350	RAID5
v	CHL 3 ID 2 (Secondary Contro	6	LD	3	6	350	RAID5
v	CHL 3 ID 3 (Primary Controll	7	LD	3	7	350	RAID5
v	Edit Host-ID/WWN Name List						

< Main Menu > Quick installation view and edit logical drives view and edit logical Volumes view and edit Host luns		LUN	LV/LD	DRV	Partition	Size(MB)	RAID
v	CHL 1 ID 0 (Primary Controll	0	LD	2	0	400	RAID5
v	CHL 1 ID 1 (Secondary Contro	1	LD	2	1	400	RAID5
v	CHL 3 ID 2 (Secondary Contro	2	LD	2	2	400	RAID5
s	CHL 3 ID 3 (Primary Controll	3	LD	2	3	400	RAID5
v	Edit Host-ID/WWN Name List	4	LD	2	4	400	RAID5
v		5	LD	2	5	400	RAID5
		6	LD	2	6	400	RAID5
		7	LD	2	7	400	RAID5

5.3 Deleting a Host LUN Mapping

To delete a host LUN mapping, perform the following steps:

1. On the Main Menu, select “view and edit Host luns.”
2. Highlight the channel and ID which is mapped to the host LUN, and press Return.
3. Highlight the specific host LUN and press Return.
4. Select Yes to delete the host LUN.

This option deletes the logical drive or logical volume mapping to the host channel. This option will not delete the data contained within the logical drive.

Note – All host LUN mappings are removed when there is any partition change.

Viewing and Editing SCSI Drives

This chapter describes how to view and edit physical drive parameters. Topics covered include:

- “SCSI Drive Status Table” on page 6-2
- “Assigning a Local Spare Drive” on page 6-4
- “Creating a Global Spare” on page 6-5
- “Viewing Drive Information” on page 6-6
- “Viewing the Connected Drives” on page 6-7
- “Scanning Drives” on page 6-8
- “Deleting a Spare Drive” on page 6-9
- “Setting Slot Numbers” on page 6-10
 - “Assigning a Slot Number to an Empty Slot” on page 6-10
 - “Adding or Deleting Drive Entries” on page 6-11
 - “Deleting a Slot Number” on page 6-11
 - “Removing an Empty Drive Entry” on page 6-12
 - “Identifying a Drive” on page 6-13
- “Cloning Failing Drive” on page 6-16
 - “Replacing After Clone” on page 6-16
 - “Perpetual Clone” on page 6-18
- “SCSI Drive Utilities (Do Not Use)” on page 6-21
 - “SCSI Drive Low-level Format” on page 6-21
 - “SCSI Drive Read/Write Test” on page 6-22

To view and edit physical drive parameters, select “view and edit scsi Drives” on the Main Menu, and press Return. The SCSI drive status table is displayed. Use this command to view the SCSI drives associated with the selected logical drive. The drive’s channel, ID, status, and model number can be viewed in the table displayed on screen.

To modify or view additional information, highlight a line in the SCSI drive table and press Return to view the available commands.

Note – The menu will vary according to the drive status.

6.1 SCSI Drive Status Table

To check and configure physical SCSI drives, on the Main Menu, select “view and edit scsi Drives,” then press Return. The resulting screen displays the status of all SCSI drives.

If there is a drive installed but not listed, the drive may be defective or not installed correctly.

When power is on, the controller scans all hard drives that are connected through the drive channels. If a hard drive was connected after the controller completes initialization, select any SCSI drive on the SCSI table, press Return, and use the “Scan scsi drive” function to let the controller recognize the newly added hard drive and configure it as a member of a logical drive.

Quick view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
view		0	0	70007	160MB	0	ON-LINE	SEAGATE ST373405LC
view		0	1	70007	160MB	0	ON-LINE	SEAGATE ST373405LC

TABLE 6-1 Parameters Displayed in the Drive Status Window

Parameters	Description
Slot	Slot number of the SCSI drive
Chl	SCSI channel of the connected drive
ID	SCSI ID of the drive
Size (MB)	Drive capacity in megabytes
Speed	xxMB Maximum synchronous transfer rate of this drive. Async The drive is using asynchronous mode.
LG_DRV	x The SCSI drive is a drive member of logical drive x. If Status shows “STAND-BY,” the SCSI drive is a local spare drive of logical drive x.
Status	<p>GLOBAL The SCSI drive is a global spare drive.</p> <p>INITING The drive is initializing.</p> <p>ON-LINE The drive is in good condition.</p> <p>REBUILD The drive is rebuilding.</p>

TABLE 6-1 Parameters Displayed in the Drive Status Window (*Continued*)

Parameters	Description
STAND-BY	Local spare drive or global spare drive. The local spare drive's LG_DRV column shows the logical drive number. The global spare drive's LG_DRV column shows "Global."
NEW DRV	The new drive has not been configured to any logical drive or as a spare drive.
USED DRV	The drive was previously configured as a logical drive or as a spare drive but is not currently configured.
BAD	Failed drive.
ABSENT	Drive slot is not occupied.
MISSING	Drive once existed, but is now missing.
SB-MISS	Spare drive missing.
Vendor and Product ID	Vendor and product model information of the drive.

To handle BAD drives, refer to "Controller Failure Symptoms, Logical Drive Rebuild, and Replacement" on page 8-5. If two drives show BAD and MISSING status, see "Determining Fatal Drive Failure" on page 8-10.

6.2 Assigning a Local Spare Drive

Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	4	9999	40MB	NONE	NEW DRV		
1	5	9999	40MB	NONE	NEW DRV		
1	6	9999	40MB	NONE	NEW DRV		
1	8	9999	40MB	NONE	NEW DRV		

A local spare drive is a standby drive assigned to serve one specified logical drive. When a member drive of this specified logical drive fails, the local spare drive becomes a member drive and automatically starts to rebuild.

Note – A logical drive composed in a non-redundancy RAID level (NRAID or RAID 0) does not support spare drive rebuild.

To assign a local spare drive, perform the following steps:

1. Select “view and edit scsi Drives.”
2. Highlight the drive which you want to designate as a spare and press Return.
3. Select “add Local spare drive.”
4. When the prompt “Add Local Spare?” is displayed, select Yes and assign the logical drive which will have the local spare.

6.3 Creating a Global Spare

To create a global spare, namely a spare drive which will automatically be used to replace any failed drive within the array, perform the following steps:

1. Select “view and edit scsi Drives.”
2. Highlight the drive which you want to designate as a spare, press Return, and select “add Global spare drive.”
3. When the prompt “Add Global Spare?” is displayed, select Yes.

Quick view view view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
	6	0	5	34732	160MB	2	ON-LINE	SEAGATE ST336605LSUN36G
view	1	View drive information					-LINE	SEAGATE ST336605LSUN36G
view	1	add Local spare drive					-LINE	SEAGATE ST336605LSUN36G
view	1	add Global spare drive					-LINE	SEAGATE ST336605LSUN36G
view		S a i s					T DRV	SEAGATE ST336605LSUN36G
sys		Add Global Spare Drive ?					T DRV	SEAGATE ST336605LSUN36G
view		Yes No					T DRV	SEAGATE ST336605LSUN36G
		disk Reserved space - unformatted					T DRV	SEAGATE ST336605LSUN36G
		0	13	34732	160MB	NONE	NEW DRV	SEAGATE ST336605LSUN36G
		0	14				SAF-TE	DotHill ERMLVD SCRUNCH

6.5 Viewing the Connected Drives

Prior to configuring disk drives into a logical drive, it is necessary to understand the status of physical drives in your enclosure.

To view the list of available SCSI drives, perform the following steps.

1. Use arrow keys to scroll down to “view and edit scsi Drives” and press Return.

Quic view view view view view	Slot	Ch1	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
		0	0	70007	160MB	0	ON-LINE	SEAGATE ST373405LC
		0	1	70007	160MB	0	ON-LINE	SEAGATE ST373405LC

2. Use arrow keys to scroll the table. Check to see if there are any drives installed but not listed here.

If there is a drive installed but not listed, the drive may be defective or not installed correctly, contact your RAID supplier.

3. If a hard drive was connected after the controller completes initialization, select a drive from the table, press

Return, and use the “Scan scsi drive” function to enable recognition of the newly added hard drive and to configure the new drive as a member of a logical drive.



Caution – Scanning an existing drive will remove its assignment to any logical drive. All data on that drive will be lost.

6.6 Scanning Drives

To scan a new SCSI drive installed after the controller completed initialization and to enable its use, perform the following steps.

1. On the Main Menu, select “view and edit scsi Drives” and press Return.
2. Select a drive from the SCSI drive table and press Return.



Caution – Scanning an existing drive will remove its assignment to any logical drive. All data on that drive will be lost.

view									
view	0	1	34732	160MB	0	ON-LINE	SEAGATE ST336605LSUN36G		
view	0	2	34732	160MB	0	ON-LINE	SEAGATE ST336605LSUN36G		
view									
syst	0	3	34732	160MB	1	ON-LINE	SEAGATE ST336605LSUN36G		
view									
view									
			View drive information			1	ON-LINE	SEAGATE ST336605LSUN36G	
			Scan scsi drive						
			set slot Number			1	ON-LINE	SEAGATE ST336605LSUN36G	
			add drive Entry						
			Identify scsi drive			E	USED DRV	SEAGATE ST336605LSUN36G	
			clone Failing drive						
			disk Reserved space - 256 mb			2	ON-LINE	SEAGATE ST336605LSUN36G	

FIGURE 6-2 “Scan scsi drive” Option

3. Select the “Scan scsi drive” function, and then press Return.

view										
view	0	1	34732	160MB	0	ON-LINE	SEAGATE ST336605LSUN36G			
view	0	2	34732	160MB	0	ON-LINE	SEAGATE ST336605LSUN36G			
view										
syst	0	3	34732	160MB	1	ON-LINE	SEAGATE ST336605LSUN36G			
view										
view										
			SCSI Channel 0			32	160MB	1	ON-LINE	SEAGATE ST336605LSUN36G
			SCSI Channel 2			32	160MB	1	ON-LINE	SEAGATE ST336605LSUN36G

The menu options will vary according to the drive status.

4. Select the drive channel and SCSI ID of the drive you want to scan, and press Return.

6.7 Deleting a Spare Drive

To delete a spare drive, perform the following steps.

1. Move the cursor to a local spare drive or global spare drive, then press Return.
2. Select “Delete global/local spare drive,” then press Return again.
3. Select Yes to confirm.

Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
	2	0	9999	40MB	0	ON-LINE	IBM DDRS-34560D
	2	1	9999	40MB	0	ON-LINE	IBM DDRS-34560D
	2	2	9999	40MB	0	ON-LINE	IBM DDRS-34560D
	2	3	9999	40MB	0	ON-LINE	IBM DDRS-34560D
	2	4	9999	40MB	0	ON-LINE	IBM DDRS-34560D
	2	5	9999	40MB	0	ON-LINE	IBM DDRS-34560D
	2	6	9999	40MB	0	ON-LINE	IBM DDRS-34560D
	2	7	9999	40MB	0	ON-LINE	IBM DDRS-34560D
	2	8	9999	40MB	GLOBAL	STAND-BY	IBM DDRS-34560D

View drive information
Delete global/local spare drive
Scan scsi drive
set slot Number
add drive Entry
Identify scsi driv

Delete Spare Drive ?
Yes No

The spare drive you deleted or any drive you replaced from a logical unit will be indicated as a “used drive.”

6.8 Setting Slot Numbers

This function is used to optionally add a slot number identifier in the Slot column of the SCSI drive table. This function has no effect on controller operation.

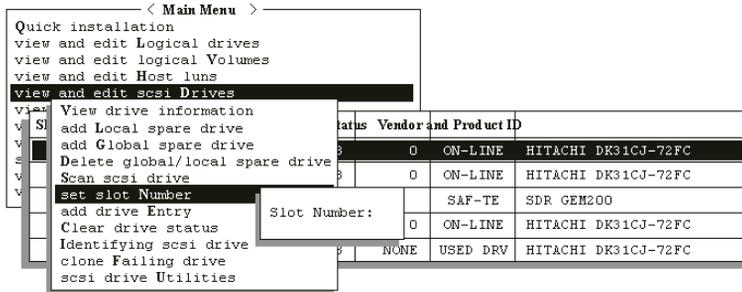


FIGURE 6-3 Set Slot Number

To set or edit a slot number, perform the following steps.

1. On the Main Menu, select “view and edit scsi Drives” and press Return.
A list of the connected SCSI drives is displayed.
2. Select a drive from the SCSI drive table and press Return.
3. Select the “set slot Number” function, and press Return.
An entry box will be displayed.
4. Type in a value (0– 15) that represents the slot number of the drive, and press Return.

This value does not need to be the device’s predetermined SCSI ID number. The slot number is displayed in the Slot column of the drive information list.

6.8.1 Assigning a Slot Number to an Empty Slot

When there is an empty slot (or sled) which does not contain a drive, the corresponding SCSI channel/ID will not appear in the drive information list.

You can assign a slot number to the empty slot and add a drive entry in order to use it later when a drive is installed.

6.8.2 Deleting a Slot Number

To delete the slot number of a SCSI drive, perform the following steps:

1. On the Main Menu, select “view and edit scsi Drives” and press Return.
2. Select the desired SCSI drive and press Return.
3. Select “Set Slot Number,” select “0” and press Return.

6.9 Adding or Deleting Drive Entries

This function is used to add an additional record to the SCSI drive table.

Use the command “Clear drive status” if you want to later remove a drive designation from the table.

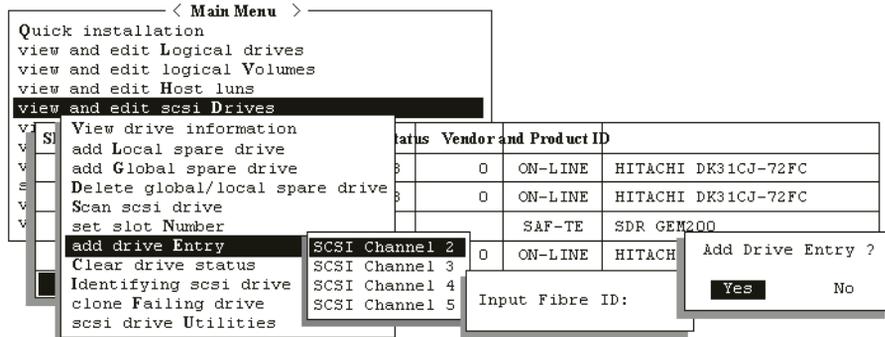


FIGURE 6-4 Add Drive Entry

1. On the Main Menu, select “view and edit scsi Drives” and press Return.
2. Select an insertion spot within the SCSI drive table and press Return.
3. Select the “add drive Entry” function, and press Return.
4. A channel list will be displayed. Select a channel.

5. Enter the desired ID number.

For installed SCSI drives, a table will be shown indicating the available IDs.

6. Then press the Return key and select “Yes.”

7. A confirmation box will then be displayed. Select “Yes” and press Return.

6.9.1 Removing an Empty Drive Entry

To remove an empty drive entry, delete its slot number (specify the value 0), then remove the drive entry, by performing the following steps.

1. On the Main Menu, select “view and edit scsi Drives” and press Return.

2. Select the desired SCSI drive and press Return.

3. Select “set slot Number,” select “0” and press Return.

4. Now select “Clear drive status” and press Return.

The empty drive entry will now disappear from the drive information list.

5. Then delete the empty drive entry (refer to “Adding or Deleting Drive Entries” on page 6-11).

Note – You will not be able to remove an empty drive entry if it has been assigned a slot number. Delete the slot number before removing the empty drive entry.

6.10 Identifying a Drive

If there is a failed drive in the RAID 5 logical drive, replace the failed drive with a new drive to keep the logical drive working. To identify a failed drive, refer to “Identifying a Drive” on page 6-13.



Caution – If, when trying to remove a failed drive, you mistakenly remove the wrong drive, you will no longer be able to access the logical drive because you have incorrectly failed another drive.

To diagnose a failed drive, identify a single drive, or test all drive activity LEDs, perform the following steps.

Note – The following procedure *only works* if there is no I/O activity.

1. On the Main Menu, select “view and edit scsi Drives” and press Return. Select the drive you want to identify, and then press Return.
2. Select the “Identifying scsi drive” function. Select “flash All drives” to flash the activity LEDs of all of the drives in the drive channel, and press Return.

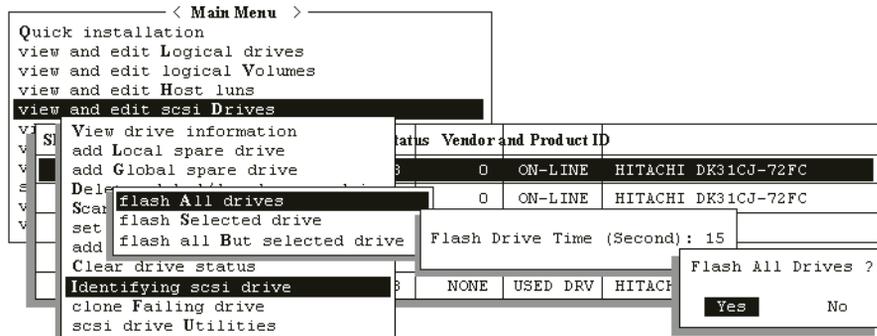


FIGURE 6-5 Flash All Drives Function

The option to change the Flash Drive Time is displayed.

3. Confirm your choice by pressing Return and selecting “Yes.”

Alternatively, to flash the read/write LED of only a selected drive, choose “flash Selected drive” or “flash all But selected drive” and perform the same procedure.

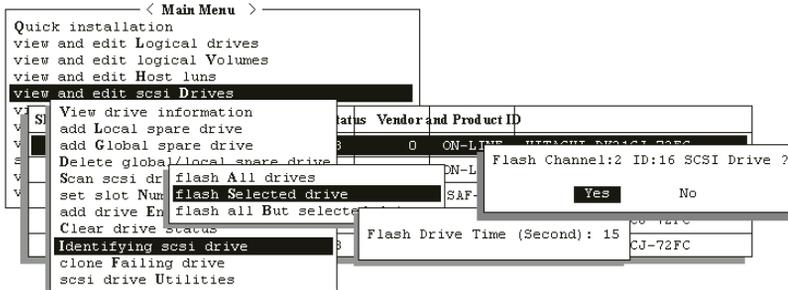


FIGURE 6-6 Flash Selected Drive Function

6.10.1 Flashing Selected SCSI Drive

The read/write LED of the drive you select lights steadily for a configurable period of time from 1 to 999 seconds.

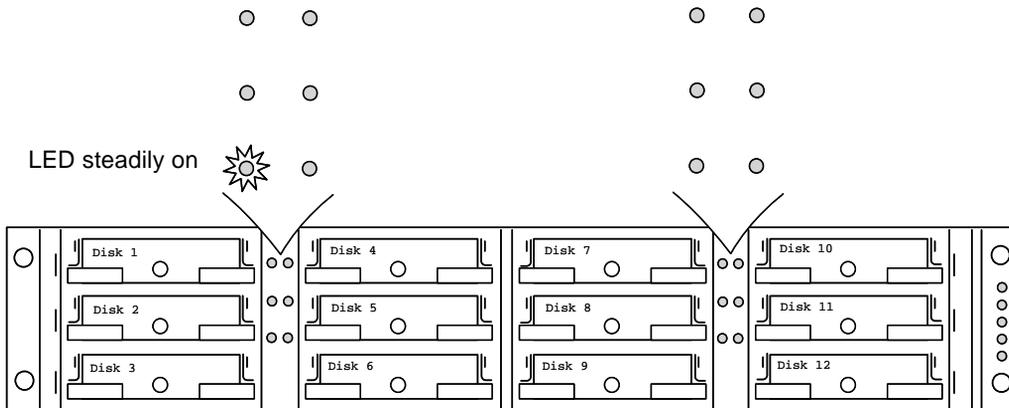


FIGURE 6-7 Flashing the Drive LED of a Selected Drive

6.10.2 Flashing All SCSI Drives

The “Flash All SCSI Drives” function lights LEDs of all good drives except the defective one.

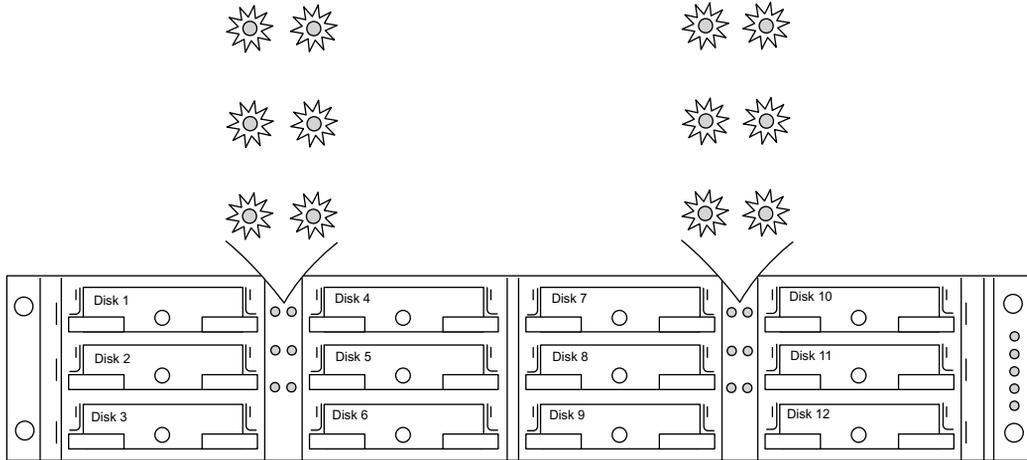


FIGURE 6-8 Flashing All Drive LEDs to Detect a Defective Non-Flashing Drive

6.10.3 Flashing All But Selected Drives

With this command, except for the selected drive, the read/write LEDs of all connected drives will light for a configurable period of time from 1 to 999 seconds.

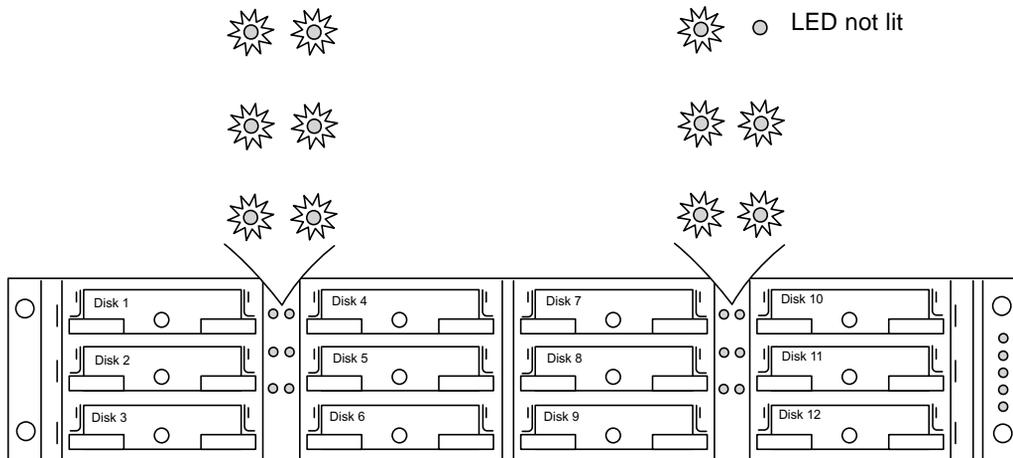


FIGURE 6-9 Flashing All Drive LEDs Except a Selected Drive LED

6.11 Cloning Failing Drive

To assist fault prevention, a system administrator can manually perform Clone Failing Drive to a drive which is about to fail. System administrators can decide when to replace a drive showing symptoms of defects by a healthy drive. A system administrator may also replace any drive at will even when a source drive is healthy.

The Clone Failing Drive is performed under the following conditions:

- Replacing drives about to fail (notified by controller).
- Manually replacing and cloning drive data on any drive to a new drive.

There are two options for cloning a failing drive:

- Replace after Clone
- Perpetual Clone

6.11.1 Replacing After Clone

Data on the source drive (the drive with predicted error or any selected member drive), will be cloned to a standby spare and the spare will then become the new source drive. The status of the original source drive is redefined as a “used drive.” System administrators may replace the used drive with a new one, and then configure the new drive as a spare drive.

Note – If there is no standby drive (local or global spare drive), you need to add a new drive and configure it as a standby drive. If there is no standby drive, the “clone failing drive” option will not be displayed.

To replace after clone, perform the following steps.

1. Select “view and edit scsi Drives” and press Return. Select the member drive that you wish to clone and press Return.
2. Select the “clone failing drive” function. This option is displayed only if there is a standby drive available.

3. Select “Replace After Clone.”

The controller automatically starts the cloning process using the existing stand-by (local or global spare drive) to clone the source drive (the target member drive with predicted error).

Quick view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID	
view	2	0	319	20MB	0	ON-LINE			
view	View drive information						0	ON-LINE	
view	Scan scsi drive						0	ON-LINE	
view	set slot Number						0	ON-LINE	
view	add drive Entry						0	STAND-BY	
view	Identify scsi drive						NONE	NEW DRV	
view	Clone Failing drive								
	Replace After Clone								
	Clone and Replace Drive ?							EW DRV	
	Yes No							EW DRV	
	2						EW DRV		

4. The cloning process begins with a notification message: [21A1] LG:0 Logical Drive NOTICE:CHL:1 ID:3 Starting Clone. Press ESC to proceed.

Quick view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
view	2	0	319	20MB	0	ON-LINE		
view	Drive Cloning							
view	<div style="border: 1px solid black; padding: 5px; text-align: center;"> <div style="background-color: black; width: 28%; height: 10px; margin: 0 auto;"></div> <p>28% Completed</p> </div>							
view	2	4	319	20MB	NONE	NEW DRV		
view	2	5	319	20MB	NONE	NEW DRV		

The cloning process is indicated by a status bar.

5. Select the drive indicated as “CLONING” by pressing Return.

(To quit the status bar, press ESC to return to the table of the connected drives).

Note – If there is no standby drive (local/global spare drive), you need to add a new drive and configure it as a standby drive.

The cloning process begins with a notification message:

[21A1] LG:0 Logical Drive NOTICE:CHL:1 ID:3 Starting Clone

2. Press ESC to view current progress on a status bar.

Quic view view view view view view view view view view view	Slot	Ch1	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
		2	0	319	20MB	0	ON-LINE	
							Drive Copying	
							35% Completed	
		2	4	319	20MB	NONE	NEW DRV	
		2	5	319	20MB	NONE	NEW DRV	

3. To quit viewing the status bar, press ESC to return to the previous menu screen.
4. Select the drive indicated as “CLONING” by pressing Return.
5. Select “clone Failing drive” again to view the progress.

You may identify the source drive and choose to “View clone progress” or “Abort clone” if you have selected the wrong drive.

Quic view view view view view view view view view view view	Slot	Ch1	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
		2	0	319	20MB	0	ON-LINE	
		2	1	319	20MB	0	ON-LINE	
		2	2	319	20MB	0	ON-LINE	
		2	3	319	20MB	0	clone	
							EW DRV	
							EW DRV	
							EW DRV	
							EW DRV	
							NONE	NEW DRV

V
S
s
a
I

Source Drive: Channel 2 ID 0
View clone progress
Abort clone

clone Failing drive

The cloning progress is completed when the following notification message is displayed: *[21A2] LG:0 Logical Drive NOTICE:CHL:1 ID:3 Clone Completed*

6.12 SCSI Drive Utilities (Do Not Use)

These utilities are used for specific troubleshooting methods and should only be used by qualified technicians.

To access these utilities, select “view and edit scsi Drives” on the Main Menu, select the drive that the utility is to be performed on, then press Return.

Select “scsi drive Utilities,” then press Return. Choose “SCSI Drive Low-level Format” or “Read/Write Test.”

Quick view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
							LINE	IBM DDRS-34560D
							ENT	IBM DDRS-34560D
							LINE	IBM DDRS-34560D
							LINE	IBM DDRS-34560D
							LINE	IBM DDRS-34560D
							LINE	IBM DDRS-34560D
							LINE	IBM DDRS-34560D
	7	2	6	9999	40MB	NONE	NEW DRV	IBM DDRS-34560D
	8	2	8	9999	40MB	1	ON-LINE	IBM DDRS-34560D

SCSI Drive Low-level Format

Quick view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
							LINE	IBM DDRS-34560D
							ENT	IBM DDRS-34560D
								-34560D
								-34560D
								-34560D
								-34560D
								-34560D
	7						NEW DRV	IBM DDRS-34560D
	8	2	8	9999	40MB	1	ON-LINE	IBM DDRS-34560D



Caution – All data on the disk drive will be destroyed when you use this command.

Viewing and Editing SCSI Channels

This chapter describes how to view and edit SCSI channels. Topics covered include:

- “SCSI Channel Status Table” on page 7-2
- “Configuring SCSI Channels as Host or Drive” on page 7-5
- “Permanent SCSI Drive Channel IDs” on page 7-7
- “Creating Additional Host IDs” on page 7-9
- “Deleting a Host Channel SCSI ID” on page 7-10
- “Drive Channel SCSI IDs (Do Not Change)” on page 7-11
- “Setting a SCSI Channel Termination (Do Not Change)” on page 7-12
- “Setting a Transfer Clock Speed” on page 7-13
- “Setting a Transfer Width” on page 7-14
- “Viewing and Editing SCSI Target / Drive Channel” on page 7-15
 - “Slot Number” on page 7-16
 - “Maximum Synchronous Transfer Clock” on page 7-16
 - “Maximum Transfer Width” on page 7-16
 - “Parity Check” on page 7-17
 - “Disconnecting Support” on page 7-17
 - “Maximum Tag Count” on page 7-18

7.1 SCSI Channel Status Table

To check and configure SCSI channels, from the Main Menu select “view and edit Scsi channels,” and press Return. The following screen displays the status of all SCSI channels for this controller.



Caution – Do not change the PID and SID values of drive channels.

```

      < Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels

```

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Drive	7	6	80.0MHz	Wide	L	0n	80.0MHz	Wide
1	Host	0		80.0MHz	Wide	L	0n	80.0MHz	Wide
2	Drive	7	6	80.0MHz	Wide	L	0n	80.0MHz	Wide
3	Host	NA	2	80.0MHz	Wide	L	0n	80.0MHz	Wide
6(D)	RCCOM								

Note – A mapped host channel sometimes shows the current sync clock as “Async/Narrow” and correctly identify the change in speed. The host adapter driver is designed to downgrade the negotiation rate on certain errors (predominantly parity errors). There is little or no performance change.

Highlight a SCSI channel and press Return to view the additional commands available for that channel.

TABLE 7-1 Parameters Displayed in the SCSI Channel Window

Parameters	Description
Chl	SCSI channel's ID.
Mode	Channel mode.
	RCCom Redundant controller communication channel
	Host The channel is functioning as a host channel.
	Drive The channel is functioning as a drive channel.
PID	Primary controller's SCSI ID mapping:
	* Multiple SCSI IDs were applied (host channel mode only).
	x The SCSI ID for host LUNs mapped to this channel in Host Channel mode. SCSI ID for the primary controller in drive channel mode.
	NA No SCSI ID applied.
SID	Secondary controller's SCSI ID mapping:
	* Multiple SCSI IDs (Host Channel mode only).
	x The SCSI ID for host LUNs mapped to this channel in host channel mode. SCSI ID for the secondary controller in drive channel mode.
	NA No SCSI ID applied
DefSynClk	Default SCSI bus synchronous clock:
	xx.xMHz Maximum synchronous transfer rate set to xx.x.
	Async Channel is set for asynchronous transfers.
DefWid	Default SCSI bus width:
	Wide Channel is set to allow wide (16-bit) transfers.
	Narrow Channel is set to allow narrow (8-bit) transfers.
S	Signal:
	S Single-ended
	L LVD
	F Fibre
Term	Terminator status:
	On Termination is enabled.
	Off Termination is disabled.

TABLE 7-1 Parameters Displayed in the SCSI Channel Window (*Continued*)

Parameters	Description
	<p>NA For a redundant controller communications channel (RCCOM).</p>
CurSynClk	<p>Current SCSI bus synchronous clock:</p> <p>xx.xMHz The current speed at which the channel is communicating.</p> <p>Async. The channel is communicating asynchronously or not device is detected.</p> <p>(empty) The default SCSI bus synchronous clock has changed. Reset the controller for changes to take effect.</p>
CurWid	<p>Current SCSI bus width:</p> <p>Wide The channel is currently servicing wide 16-bit transfers.</p> <p>Narrow The channel is currently servicing wide 8-bit transfers.</p> <p>(empty) The default SCSI bus width has changed. Reset the controller for the changes to take effect.</p>

7.1.1 SCSI Drive Channel Commands

In the “view and edit Scsi channels” window, highlight a SCSI drive channel and press Return to view the additional commands available for drive channels.

The screenshot shows a terminal window with a menu at the top and a table below. The menu options are: Quick installation, view and edit logical drives, view and edit logical Volumes, view and edit Host luns, view and edit scsi Drives, and view and edit Scsi channels (highlighted). The table below has columns: Ch, efWid, S, Term, CurSynClk, and CurWid. The rows represent SCSI channels 0 through 6.

Ch	efWid	S	Term	CurSynClk	CurWid
0	Wide	L	Off	80.0MHz	Wide
1	Wide	L	Off	20.0MHz	Wide
2	Wide	L	Off	Async	Narrow
3	Wide	L	Off	80.0MHz	Wide
6	Serial	F	NA	1 GHz	Serial

7.1.2 SCSI Host Channel Commands

In the “view and edit Scsi channels” window, highlight a SCSI host channel and press Return to view the additional commands available for host channels.

	Clk	DefWid	S	Term	CurSynClk	CurWid
0	Hz	Wide	L	Off	80.0MHz	Wide
1	Host	0	1	80.0MHz	Wide	L Off 20.0MHz Wide
2	Drive	6	7	80.0MHz	Wide	L Off Async Narrow
3	Host	3	2	80.0MHz	Wide	L Off 80.0MHz Wide
6(C)	RCCOM	NA	NA	AUTO	Serial	F NA 1 GHz Serial

7.2 Configuring SCSI Channels as Host or Drive

All Sun StorEdge RAID arrays are preconfigured when they arrive from the factory. Default channel settings and rules are specified as follows:

- Channel 0 (CH 0) MUST be a drive channel.
- Channels 1, 2, and 3 (CH 1, CH 2, CH 3) may be drive or host channels.
- Default channel settings are:
 - CH 0 and CH 2 = Drive channels
 - CH 1 and CH 3 = Host channels

The most common reason to change a host channel to a drive channel is when you attach an Expansion Unit to a RAID array and only need one host channel.

To change the use of a SCSI channel, reconfigure the channel according to the following procedure:

1. Select “view and edit Scsi channels” from the Main Menu.

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
  
```

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Drive	7	6	80.0MHz	Wide	L	On	80.0MHz	Wide
1	Host	0		80.0MHz	Wide	L	On	80.0MHz	Wide
2	Drive	7	6	80.0MHz	Wide	L	On	80.0MHz	Wide
3	Host	NA	2	80.0MHz	Wide	L	On	80.0MHz	Wide
6(D)	RCCOM								

The communications path for the controllers is displayed as “RCCOM” (Redundant Controller Communications).

2. Highlight the channel which you want to modify and press Return.

Use the arrow key to select Yes if you want to change the host or drive assignment.

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
  
```

Ch	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Drive	7	6	80.0MHz	Wide	L	On	80.0MHz	Wide
1	Host	0		80.0MHz	Wide	L	On	80.0MHz	Wide
2	Drive	7	6	80.0MHz	Wide	L	On	80.0MHz	Wide
3	Host	NA	2	80.0MHz	Wide	L	On	80.0MHz	Wide
6(D)	RCCOM								

Change Mode to Host Channel ?									
Yes <input type="checkbox"/> No <input checked="" type="checkbox"/>									

```

View and edit scsi target
parity check - Enabled
view chip inFormation
  
```



Caution – The channels of redundant controllers must be the same. For example, if the primary controller uses channel 2 to connect to a group of drives, the secondary controller must also use channel 2 to connect to the same group of drives. Changes to the primary controller are automatically be made to the secondary controller.

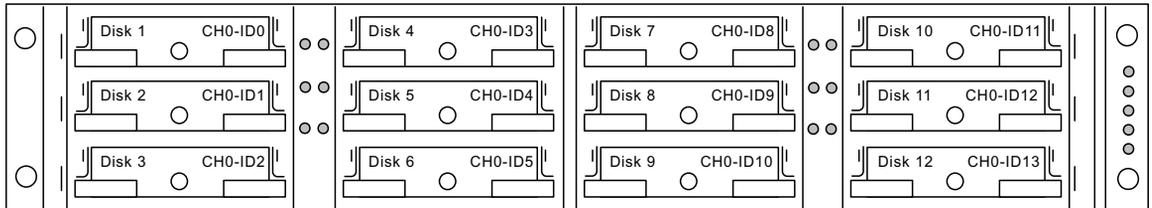
7.3 Permanent SCSI Drive Channel IDs

Each array must be configured as a single bus configuration or a dual bus configuration, based on where the SCSI bus cable is attached on the I/O module. For bus configuration details, refer to the chapter on Connecting Ports, in the *Sun StorEdge Installation, Operation, and Service Manual*.

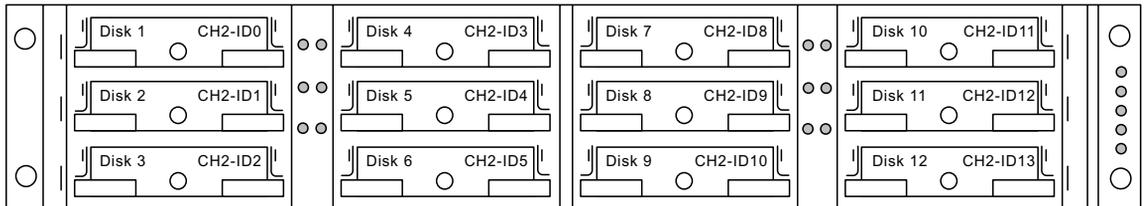
The drive-bus configuration determines how drives and drive IDs are assigned to drive channels on the controller.

- A **single bus configuration** assigns all 12 disk drive IDs in a controller to one channel (typically CH 0 for the RAID array and CH 2 for an Expansion Unit).

RAID Array - Single Bus Configuration - Default IDs

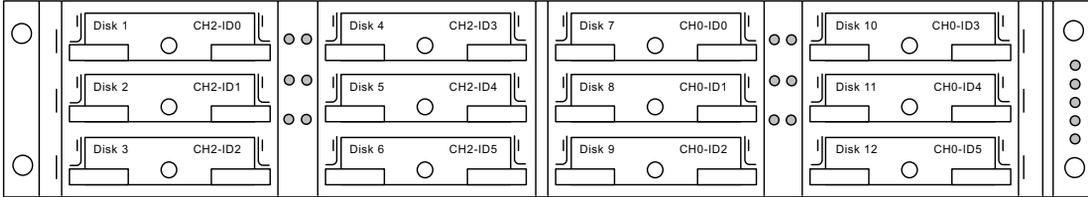


Expansion Unit - Single Bus Configuration - Default IDs

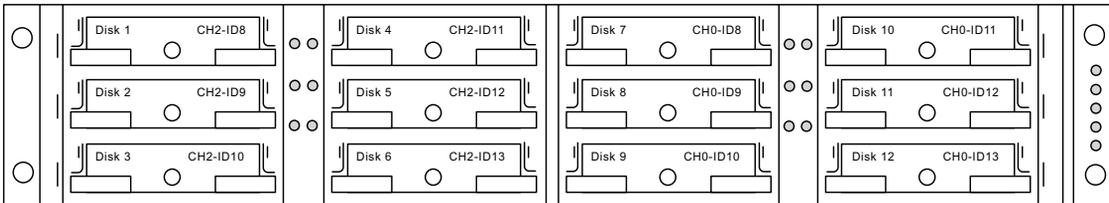


- A **dual bus configuration** assigns 6 disk drive IDs to CH 0 and 6 disk drive IDs to CH 2 in the RAID array, and then typically adds an additional 6 disk drive IDs to both CH 0 and CH2 when it is connected to an Expansion Unit.

RAID Array - Dual Bus Configuration - Default IDs



Expansion Unit - Dual Bus Configuration - Default IDs



For JBOD cabling and drive ID information, refer to the appendix on JBODs in the *Sun StorEdge Installation, Operation, and Service Manual*. JBODs are arrays without controllers which are connected directly to host servers.

7.4 Creating Additional Host IDs

All RAID arrays are preconfigured when they arrive from the factory.

Default host channel IDs are:

- Channel 1 ID 0 (primary controller)
- Channel 3 ID 1 (secondary controller)

Each host channel might have two editable ID numbers:

- Primary controller ID
- Secondary controller ID

Each ID number must be a unique number within the host channel. You can:

- Edit each host ID number to change the SCSI target number of each controller host channel that is seen by the host.
- Add additional host ID numbers (by adding a second host ID to channels 1 and 3, and additional host IDs if you make Channel 2 into a host channel).

Note – To map 128 partitions into 128 LUNs, you must add additional host IDs. A minimum of four host IDs are required; a maximum of six host IDs are possible. For details on mapping 128 LUNs, refer to “Mapping Logical Drive Partitions to Host LUNs” on page 5-3.

To select a unique ID number for a host channel, perform the following steps.

1. Select “view and edit Scsi channels.”

Press Return.

2. Highlight the host channel on which you want to edit the Primary/Secondary ID, and press Return.

Press Return to access the list of IDs and use the arrow keys to select an ID number (0 through 15), then press Return again.

3. Select “view and edit scsi Id” and press Return. Select “Add Channel SCSI ID.”

4. Select the controller on which you want to add a host ID.

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels

```

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	Cur
0	Host	0	1	40.0MHz	Wide	L	On	Async	Nar
1	ID 0 Primary Controller ID 1 Secondary Controller				Wide	L	On	Async	Nar
2					Wide	S	On	20.0MHz	Wi
3	Add Channel SCSI ID Delete Channel SCSI ID				Wide	L	On	Async	Nar

ID 2
ID 3
ID 4
ID 5
ID 6
ID 7
ID 8
ID 9
ID 10
ID 11
ID 12
ID 13
ID 14
ID 15

5. Select an ID number for that controller.

Note – To create a total of 128 LUNs, you must have a minimum of four host IDs (two each for Channels 1 and 3) and might have a maximum of six host IDs (two each for Channels 1 and 2, and 3). Each host ID can have up to 32 partitions, which are then mapped to LUNs to create a total not to exceed 128.

6. From the Main Menu, select “system Functions,” then “Reset controller.”

The configuration change takes effect only after the controller is reset.

7.5 Deleting a Host Channel SCSI ID

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels

```

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	0	1	40.0MHz	Wide	L	On	Async	Narrow
1	ID 0 Delete Secondary Controller SCSI ID 1 ? ID 1							ync	Narrow
2					No			0MHz	Wide
3	Delete Channel SCSI ID				Wide	L	On	Async	Narrow

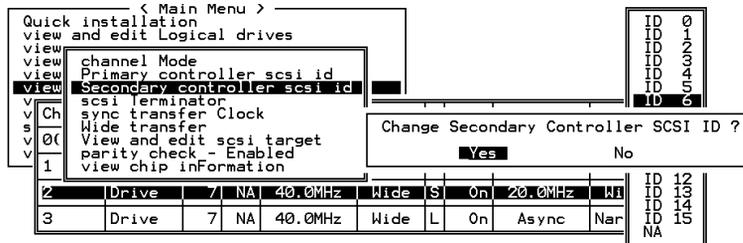
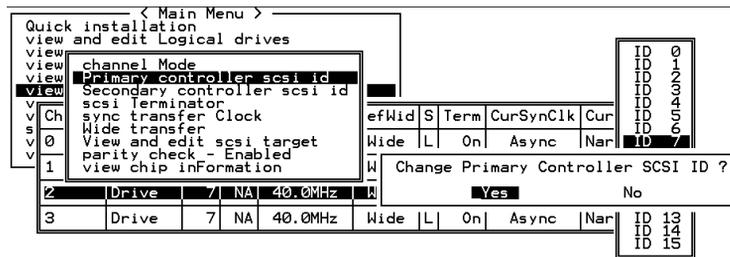
To delete a host channel SCSI ID, perform the following steps.

1. On the Main Menu, select “view and edit Scsi channels” and press Return.
2. Press Return on a host channel line with the SCSI ID you want to delete.
3. Select “Delete Channel SCSI ID.”

The dialog box “Delete Primary/Secondary Controller SCSI ID?” is displayed.

4. Select **Yes**, then press **Return** to confirm.
5. On the **Main Menu**, select “**system Functions**”, then “**Reset controller.**”
The configuration change will take effect only after the controller is reset.

7.6 Drive Channel SCSI IDs (Do Not Change)



On the **Main Menu**, highlight a drive channel and press **Return**.

DO NOT USE the following commands that appear:

- “Primary controller scsi id”
- “Secondary controller scsi id”

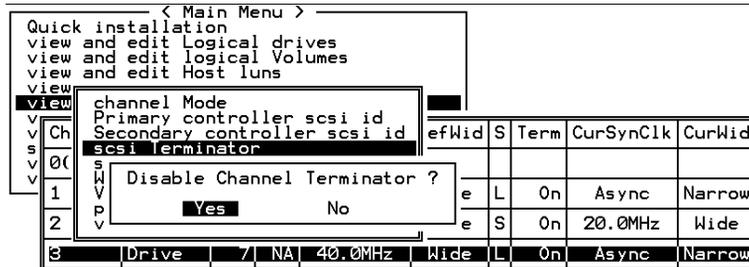
These commands are used to change the default drive SCSI IDs.



Caution – To edit a drive channel SCSI ID may create conflicts with the controller communication channels and create confusion in tracking the status of drive IDs in either single bus or dual bus configurations.

For default drive channel IDs, refer to “Permanent SCSI Drive Channel IDs” on page 7-7.

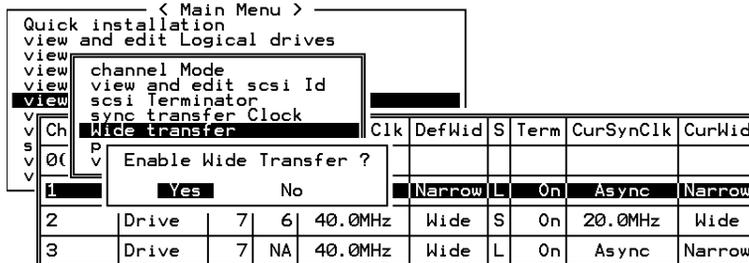
7.7 Setting a SCSI Channel Termination (Do Not Change)



Typically the default setting is not changed.

1. Select the channel you want the terminator enabled or disabled, then press Return.
2. Select "scsi Terminator", then press Return.
A dialog box is displayed.
3. Select Yes, then press Return.

7.9 Setting a Transfer Width



Typically the default setting for transfer speed is not changed for the host or drive channel.

To view the options for transfer speed, select “view and edit Scsi channels,” highlight the drive or host channel, press Return, and select “Wide transfer” (if the channel default is “Narrow”) or select “narrow transfer” (if the channel default is “Wide”).

Note – Every time you change the SCSI Transfer Width, you must reset the controller for the changes to take effect.

7.10 Viewing and Editing SCSI Target / Drive Channel

On the Main Menu, select “view and edit Scsi channels,” highlight the drive channel, and press Return. Select “View and edit scsi target,” then press Return.

< Main Menu >													
Quick installation													
view and edit Logical drives													
view													
view	channel Mode												
view	Primary controller scsi id												
view	Secondary controller scsi id												
view	scsi Terminator												
v	Ch								efWid	S	Term	CurSynClk	CurWid
v	s												
v	0	View and edit scsi target											
v	1	parity check - Enabled							arrow	L	0n	Async	Narrow
		view chip information											
	2	Drive	7	6	40.0MHz	Wide	S	0n	20.0MHz	Wide			
	3	Drive	7	NA	40.0MHz	Wide	L	0n	Async	Narrow			

A list of all the SCSI targets and their current settings is displayed. Press Return on a SCSI target to display an editable parameter menu on the screen.

Quick	Slot	Ch1	ID	SyncClk	XfrWid	ParityChk	Disconnect	TagCount	
view	2	0		10	Wide	Enabled	Enabled	Def(32)	
view									
view									
view									
v	Ch	Slot number					bled	Enabled	Def(32)
v	s	maximum sync. xfer Clock					bled	Enabled	Def(32)
v	0	maximum xfer Width					bled	Enabled	Def(32)
v	1	Parity check					bled	Enabled	Def(32)
		Disconnect support					bled	Enabled	Def(32)
		maximum Tag count					bled	Enabled	Def(32)
		Restore to default setting					bled	Enabled	Def(32)
	2	2	5	10	Wide	Enabled	Enabled	Def(32)	
	3	2	8	10	Wide	Enabled	Enabled	Def(32)	
		2	9	10	Wide	Enabled	Enabled	Def(32)	

Note – It is only recommended to alter the SCSI target settings when adjustments should be made to specific devices on a drive channel. You may change the SCSI parameters for specific drives when mixing different drives or connecting other SCSI device like a CD-ROM on a drive channel. Note that neither mixing drives nor connecting CD-ROM is recommended for the controller.

7.10.6 Maximum Tag Count

Quick view	Slot	Ch1	ID	SyncClk	XfrWid	ParityChk	Disconnect	TagCount
view	1	2	0	9	Wide	Enabled	Enabled	Def (32)
view	2	2	1	9	Wide	Enabled	Enabled	Def (32)
view	3	2	2	9	Wide	Enabled	Enabled	Def (32)
view	4	2	3	9	Wide	Enabled	Enabled	Def (32)
view	5	2	4	9	Wide	Enabled	Enabled	Def (32)
view	6	2	6	9	Wide	Enabled	Enabled	Def (32)
view	7	2	8	9	Wide	Enabled	Enabled	Def (128)

Ch	Slot number	bled	Enabled	Def	Default
0	maximum sync. xfer Clock	bled	Enabled	Def	Disable
1	maximum xfer Width	bled	Enabled	Def	Disable
2	Parity check	bled	Enabled	Def	Disable
3	Disconnect support	bled	Enabled	Def	Disable
4	maximum tag count	bled	Enabled	Def	Disable
5	Restore to default setting	bled	Enabled	Def	Disable

Set Maximum Tag Count ?	
Yes	No

1. Select “maximum Tag count,” then press Return.
A list of available tag count numbers is displayed.
2. Select a number, then press Return.
3. Select Yes in the dialog box that follows to confirm the setting.



Caution – Disabling the Maximum Tag Count will disable the internal cache of the SCSI drive.

Viewing and Editing Configuration Parameters

This chapter describes viewing & editing configuration parameters. Topics covered:

- “Optimization Modes (Caching Parameters)” on page 8-2
 - “Optimization for Random or Sequential I/O” on page 8-3
 - “Write-Back/Write-Through Cache Enable/Disable” on page 8-4
- “Controller Failure Symptoms, Logical Drive Rebuild, and Replacement” on page 8-5
 - “General Controller Failure Response” on page 8-5
 - “Automatic Logical Drive Rebuild” on page 8-6
 - “Manual Rebuild” on page 8-6
 - “Concurrent Rebuild in RAID (0+1)” on page 8-9
 - “Identifying a Failed Drive for Replacement” on page 8-9
 - “Restoring Your Configuration (NVRAM) From a File” on page 8-9
 - “Determining Fatal Drive Failure” on page 8-10
- “Controller Parameters” on page 8-12
 - “Controller Name” on page 8-12
 - “LCD Title Display - Controller Logo (Not Applicable)” on page 8-13
 - “Password Validation Timeout” on page 8-13
 - “Controller Unique Identifier” on page 8-15
 - “SDRAM ECC Function (Do Not Change)” on page 8-15
- “Drive-side SCSI Parameters” on page 8-16
 - “SCSI Motor Spin-Up (Do Not Change)” on page 8-17
 - “SCSI Reset at Power-Up (Do Not Change)” on page 8-18
 - “Disk Access Delay Time” on page 8-19
 - “SCSI I/O Timeout” on page 8-19
 - “Maximum Tag Count (Tag Command Queuing)” on page 8-21
 - “SAF-TE and S.E.S. Enclosure Monitoring” on page 8-22
 - “Periodic Drive Check Time” on page 8-22
 - “Auto-Detect Failure Drive Swap Check Time” on page 8-23
- “Disk Array Parameters” on page 8-24

- “Rebuild Priority” on page 8-24
- “Verification on Writes” on page 8-25
- “Host-side SCSI Parameters” on page 8-27
 - “Overview of SCSI Channel, SCSI ID and LUN” on page 8-27
 - “Maximum Concurrent Host-LUN Connections” on page 8-28
 - “Number of Tags Reserved for Each Host LUN Connection” on page 8-29
 - “Maximum Queued I/O Count” on page 8-29
 - “LUNs Per Host SCSI ID” on page 8-30
 - “Cylinder/Head/Sector Mapping” on page 8-31
- “Peripheral Device Type Parameters” on page 8-33
- “Setting an IP Address” on page 8-34

8.1 Optimization Modes (Caching Parameters)

Mass storage applications fall into two major categories: database applications and video/imaging applications. The controller supports two embedded optimization modes: Optimization for Random I/O and Optimization for Sequential I/O.

Limitations: There are limitations on the optimization modes. First, one optimization mode must be applied to all logical units in a RAID array. Second, once the optimization mode is selected and data written in logical units, the only way to change the optimization mode is to back up all data to another location, delete all logical configurations of drives, reconfigure the logical drive configuration with the new optimization mode, and reboot the array.

The limitation derives from the consideration with the redundant configuration of controllers. Data inconsistency might occur when a controller pre-configured with one optimization mode is used to replace a failed controller with different mode.

- **Database and Transaction-based Applications:** This kind of application usually includes an SQL server, Oracle server, Informix, or other database services. These applications keep each transaction within the minimal size, so that I/O transfers will not be clogged by one large transaction. Due to its transaction-based nature, these applications do not read or write data in a sequential order. Instead, access to data occurs randomly. The transaction size ranges from 2K to 4K. Transaction-based performance is usually measured in I/Os per second or IOPS.
- **Video Recording/Playback and Imaging Applications:** This kind of application usually belongs to video playback, video post-production editing, or other similar applications. These applications read or write large files to and from storage in a sequential order. The size of each I/O can be 128K, 256K, 512K, or up to 1MB. Performance is measured in MB/Sec.

When an array works with applications such as video or image oriented applications, the application reads/writes from the drive as large-block, sequential files instead of small-block and randomly-accessed files.

- Optimization for Random I/O (32K block size):

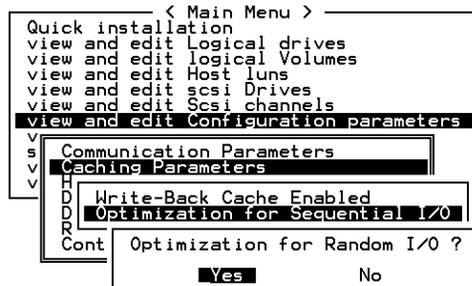
The logical drive, cache memory, and other controller parameters will be adjusted for the use of database/transaction-processing applications.

- Optimization for Sequential I/O (128K block size):

Optimization for Sequential I/O provides larger stripe size (block size, also known as Chunk size) than Optimization for Random I/O. Numerous controller's internal parameters will also be changed to optimize for sequential or random I/O. The change will take effect after the controller resets.

The logical drive, cache memory, and other controller internal parameters will be adjusted for the use of video/imaging applications.

8.1.1 Optimization for Random or Sequential I/O



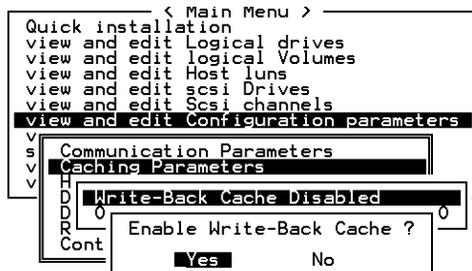
The default optimization mode is for “Sequential.” Optimization mode for sequential is automatically applied to any logical configuration of drives larger than 512GB.

To select the optimization mode for all drives, perform the following steps.

1. On the Main Menu, select “view and edit Configuration parameters,” then select “Caching Parameters.”
2. Select “Optimization for Random I/O” or “Optimization for Sequential I/O.”
3. Then press Return. The “Random” or “Sequential” dialog box is displayed, depending on the option you have selected.
4. Select Yes in the dialog box that follows to confirm the setting.

8.1.2 Write-Back/Write-Through Cache Enable/Disable

The write-back cache function significantly enhances controller performance. Write-through strategy is considered more secure if power failure should occur. Because a battery module is installed, power will be supplied to the data cached in memory and the cached writes can be completed when power is restored.



To change the caching parameter option, perform the following steps.

1. On the Main Menu, select “view and edit Configuration parameters” and press Return.
2. Select “Caching Parameters,” and press Return.
3. Select “Write-Back Cache,” then press Return. “Enabled” or “Disabled” will display the current setting with the Write-Back caching.
4. Select Yes in the dialog box that follows to confirm the setting.

8.2 Controller Failure Symptoms, Logical Drive Rebuild, and Replacement

Controller failure symptoms are as follows:

- The surviving controller sounds an audible alarm.
- The center LED (status symbol) flashes yellow on the failed controller.
- The surviving controller sends event messages announcing the controller failure of the other controller.

A warning message, “SCSI Bus Reset Issued” is displayed for each of the SCSI channels.

In addition, an alert message says “Redundant Controller Failure Detected.”

8.2.1 General Controller Failure Response

If one controller in the redundant controller configuration fails, the surviving controller temporarily takes over for the failed controller until it is replaced.

A failed controller is managed by the surviving controller which disables and disconnects from its counterpart while gaining access to all the signal paths. The surviving controller then manages the ensuing event notifications and takes over all processes. It is always the primary controller regardless of its original status, and any replacement controller afterward will assume the role of the secondary controller.

The failover and failback processes are completely transparent to the host. Controllers are hot-swappable, and replacing a failed unit takes only a few minutes.

To maintain your redundant controller configuration, replace the failed controller as soon as possible.

8.2.2 Automatic Logical Drive Rebuild

Rebuild with Spare

When a member drive in a logical drive fails, the controller will first examine whether there is a local spare drive assigned to this logical drive. If yes, it will automatically start to rebuild the data of the failed disk to it.

If there is no local spare available, the controller will search for a global spare. If there is a global spare, it will automatically use it to rebuild the logical drive.

Failed Drive Swap Detect

If neither a local spare drive nor a global spare drive is available, and the “Periodic Auto-Detect Failure Drive Swap Check Time” is “disabled,” the controller will not attempt to rebuild unless you apply a forced-manual rebuild.

To enable this feature, on the Main Menu, choose “view and edit Configuration parameters,” then select “Drive-side SCSI Parameters,” and select “Periodic Auto-Detect Failure Drive Swap Check Time.”

When the “Periodic Auto-Detect Failure Drive Swap Check Time” is “Enabled” (that is, a check time interval has been selected), the controller will detect whether or not the failed drive has been swapped (by checking the failed drive's channel/ID). Once the failed drive has been swapped, the rebuild will begin immediately.

If the failed drive is not swapped but a local spare is added to the logical drive, the rebuild will begin with the spare.

For a flowchart of automatic rebuild, see FIGURE 8-1.

8.2.3 Manual Rebuild

When a user applies forced-manual rebuild, the controller will first examine whether there is any local spare assigned to the logical drive. If yes, it will automatically start to rebuild.

If there is no local spare available, the controller will search for a global spare. If there is a global spare, the logical drive rebuild will begin. See FIGURE 8-2.

If neither local spare nor global spare is available, the controller will examine the SCSI channel and ID of the failed drive. After the failed drive has been replaced by a healthy one, the logical drive rebuild begins on the new drive. If there is no drive available for rebuilding, the controller will not attempt to rebuild until the user applies another forced-manual rebuild.

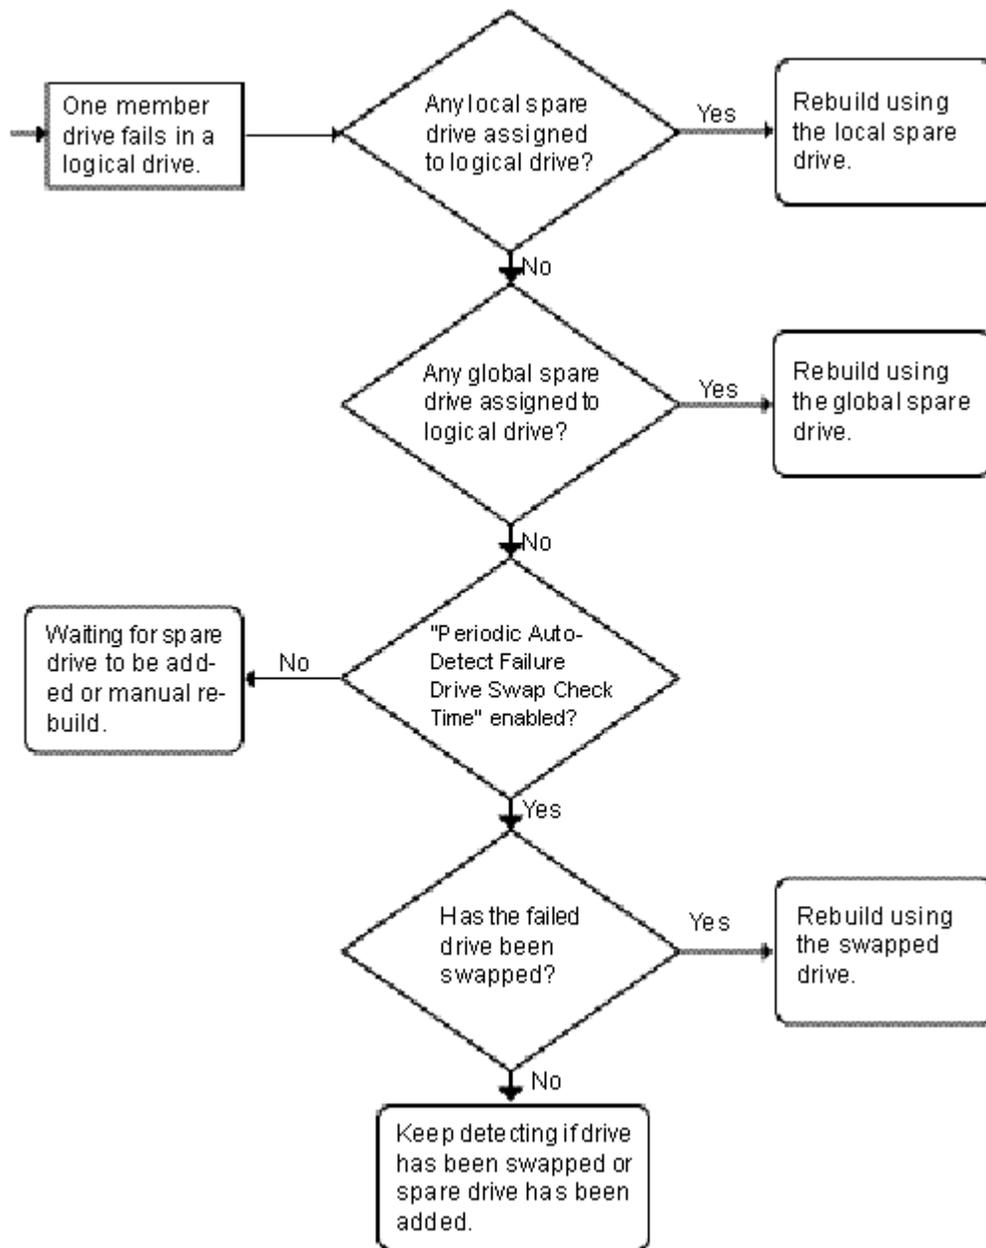


FIGURE 8-1 Automatic Rebuild

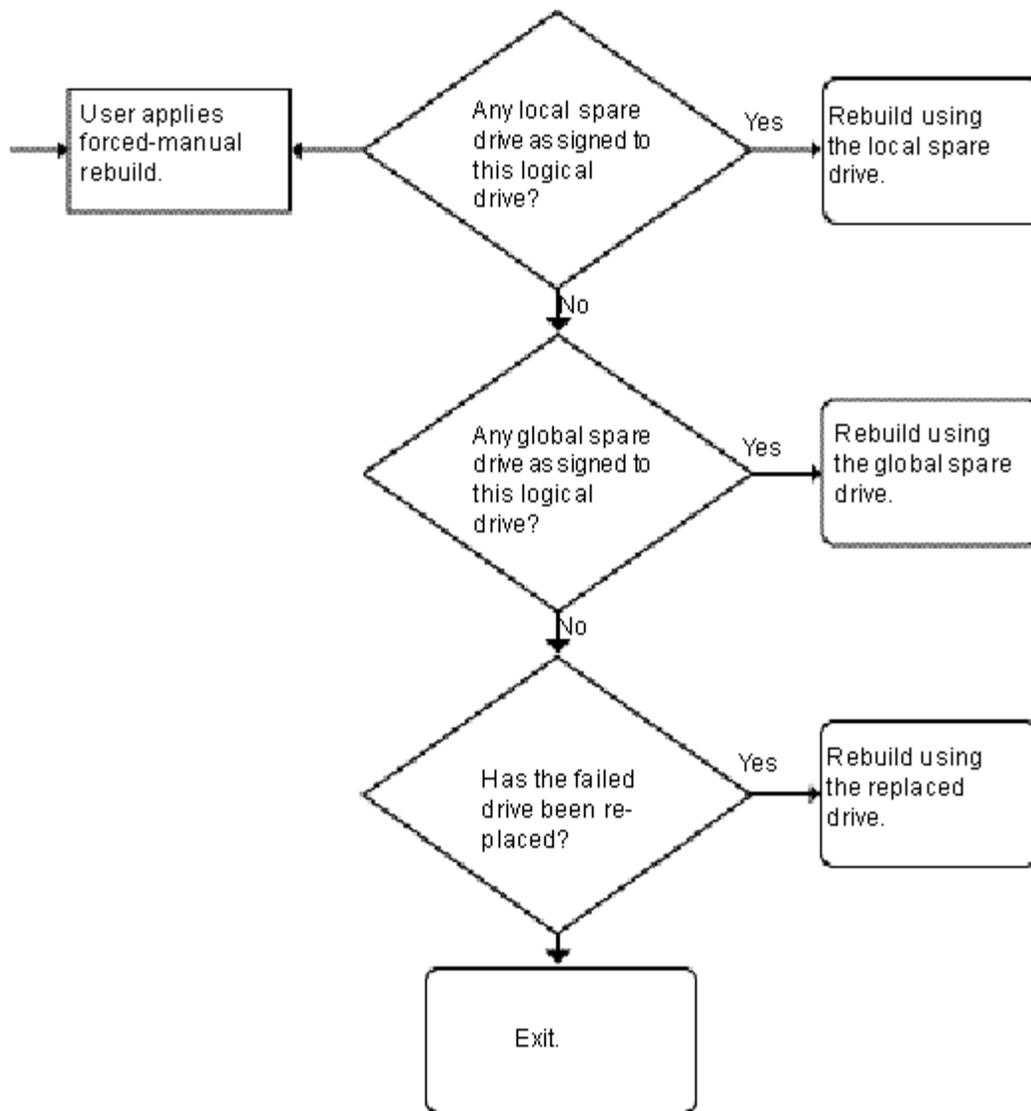


FIGURE 8-2 Manual Rebuild

8.2.4 Concurrent Rebuild in RAID (0+1)

RAID (0+1) allows multiple-drive failure and concurrent multiple-drive rebuild. Drives newly swapped must be scanned and set as local spares. These drives will be rebuilt at the same time (you do not need to repeat the rebuilding process for each drive.)

8.2.5 Identifying a Failed Drive for Replacement

If there is a failed drive in the RAID 5 logical drive, replace the failed drive with a new drive to keep the logical drive working. To identify a failed drive, refer to “Identifying a Drive” on page 6-13.



Caution – If, when trying to remove a failed drive, you mistakenly remove the wrong drive, you will no longer be able to access the logical drive because you have incorrectly failed another drive.

8.2.6 Restoring Your Configuration (NVRAM) From a File

If you have saved a configuration file and want to apply the same configuration to another array or reapply it to the array that had the configuration originally, you must be certain that the channels and SCSI IDs in the configuration file are correct for the array where you are restoring the configuration.

The NVRAM configuration file will restore all configuration settings (channel settings, host IDs, etc.) but does not rebuild logical drives.

To save a configuration file, refer to “Saving Configuration (NVRAM) to Disk” on page 10-6.



Caution – If the channels or SCSI IDs are not a correct match for the array, you will lose access to the mismatched channels or drives when you restore the configuration with the configuration file.

Note – In the Configuration Service program, you can save a configuration file that can restore all configurations and rebuild all logical drives. However, it will also erase all data when it rebuilds all logical drives, so operation is performed only when no data has been stored or all data has been transferred to another array.

To restore configuration settings from a saved NVRAM file, perform the following steps.

1. **From the Main Menu choose “system Functions.”**
2. **Select “Controller maintenance” and press Return.**
3. **Select “Restore NVRAM from disks,” and press Return.**
4. **Press Yes to confirm.**

A prompt will notify you that the controller NVRAM data has been successfully restored from disks.

8.2.7 Determining Fatal Drive Failure

With the redundant RAID array system, your system is protected with the RAID parity drive and by the default global spare (you may have more than one).

Note – A FATAL FAIL status occurs when there is one more drive failing than the number of spare drives available for the logical drive. If a logical drive has two global spares available, then three failed drives must occur for FATAL FAIL status.

In an extremely rare occurrence where two or more drives appear to fail at the same time, perform the following steps.

1. **Discontinue all input/output activity immediately.**
2. **To cancel the beeping alarm, use a paperclip to press the Reset button (below the LEDs on the front-right ear of the array).**
3. **Physically check that all the drives are firmly seated in the array and that none have been partially or completely removed.**
4. **Check again the firmware Main Menu and check the “view and edit Logical drives,” and look for:**

Status: FAILED DRV (one failed drive) or

Status: FATAL FAIL (two or more failed drives)

5. Highlight the logical drive, press Return, and select “view scsi drives.”

If two physical drives have a problem, one drive will have a BAD status and one drive will have a MISSING. The MISSING status is a reminder that one of the drives may be a “false” failure. The status does not tell you which drive might be a false failure.

6. Do one of the following:

- Choose “system Functions” from the Main Menu and press Return. Select “Reset controller,” and press Return.

or

- Power off the array. Wait five seconds, and power on the array.

7. Repeat steps 4 and 5 to check the logical and SCSI drive status.

After resetting the controller, if there is a false bad drive, the array will automatically start rebuilding the failed RAID set.

If the array does not automatically start rebuilding the RAID set, check the status under “view and edit Logical drives.”

- If the status is “FAILED DRV,” manually rebuild the RAID set (refer to “Manual Rebuild” on page 8-6).
- If the status is still “FATAL FAIL,” you have lost all data on the logical drive and must re-create the logical drive. Proceed with the following procedures:
 - Replace the failed drive with a new disk drive (refer to the *Sun StorEdge 3310 SCSI Array Installation, Operation and Service Manual Guide*)
 - “Deleting a Logical Drive” on page 3-16
 - “Creating Logical Drive(s)” on page 3-5

For additional troubleshooting tips, refer to the Sun StorEdge 3310 SCSI Array Release Notes located at:

www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3310

8.3 Controller Parameters

8.3.1 Controller Name

The Controller Name only is displayed in the firmware program and is used to identify different controllers.

Note – The controller’s name and password jointly share a 16-character alphanumeric field. If you set up a password, check that both the controller name and any password can fit within the 16-character field.

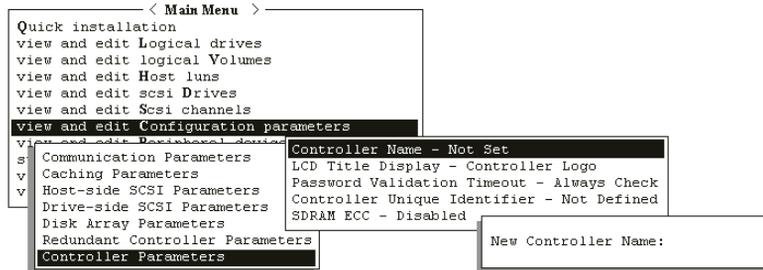
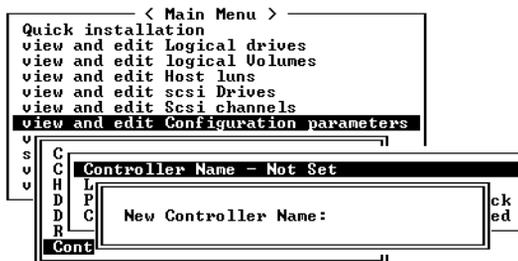


FIGURE 8-3 Controller Name

1. Select “view and edit Configuration parameters,” “Controller Parameters,” then press Return.
2. From the “Controller Parameters” menu, select “Controller Name”, and then press Return.



Depending on the controller’s current settings, you will be prompted to either enter a new name or modify the existing name for the designated controller.

3. Enter a name for the controller, and confirm by pressing the Enter key.

8.3.2 LCD Title Display - Controller Logo (Not Applicable)

This function is not applicable to this product.

8.3.3 Password Validation Timeout

This function sets a timeout when a password is required to be entered.

If a single password is set, the operator must enter this case-sensitive, alphanumeric password each time the controller is reset, causing an initial display of the Terminal Interface screen. In most cases, the default value “Always Check” should be left unchanged.

Although this function allows you to set the timeout setting, it does not provide a means of counting “retries”. In other words the user may continue to retry entering a password until the preset timeout expires, unless the default “Always Check” value is selected. The other options available are “Disable” or setting a value for 1, 2, or 5 minutes.

Leaving this setting at “Always Check” means that there is no defined timeout, and the operator has unlimited opportunities to enter the correct password, but each try is validated before access to the firmware’s functions is permitted. If this function is disabled, any entry will provide immediate access to the Main Menu topics, despite whether or not a password has been established.

Note – This firmware only allows one password to be entered. Subsequently, there are not unique timeout options for each potential operator.

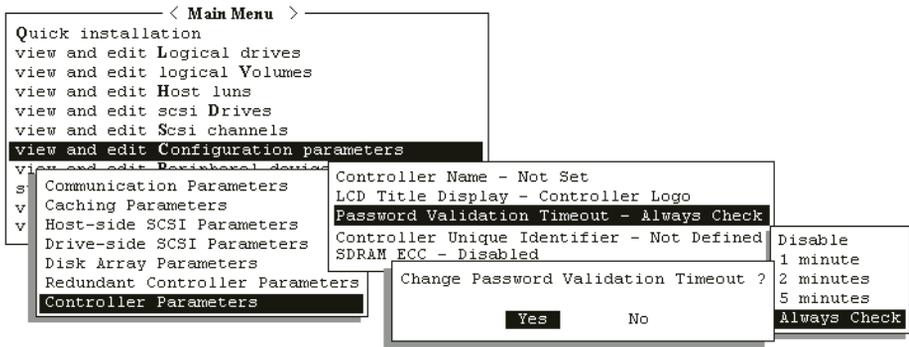
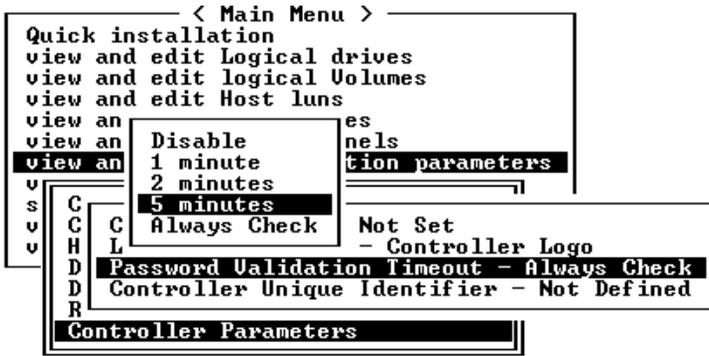


FIGURE 8-4 Password Validation Timeout

To set Password Validation Timeout, follow these steps.

1. From the Main Menu, select “view and edit Configuration parameters,” then select “Controller Parameters,” and press Return.
2. Select “Password Validation Timeout”, and then press Return.



3. Select a validation timeout from the list displayed and press Return.
 Enable a validation timeout from one minute to Always Check. The Always Check timeout will disable any configuration change without entering the correct password.
4. A prompt will then be displayed. Select “Yes” to confirm.

8.3.4 Controller Unique Identifier

The Controller Unique Identifier is automatically set by the SAF-TE device. The Controller Unique Identifier is used to create Ethernet Addresses and World Wide Names.



Caution – However, *if the array is powered off during the controller replacement or if you replaced a controller in a single-controller configuration, you must set the Controller Unique Identifier to the correct value or the array could become inaccessible.*

To set the "Controller Unique Identifier" parameter to the correct value, perform the following steps.

1. **On the firmware main menu, select “view and edit configuration parameters,” then select “Controller Parameters,” and press Return.**
2. **From the Controller Parameters menu, select “Controller Unique Identifier <hex>” and press Return.**
3. **Type in the value 0 (to automatically read the chassis serial number from the midplane) or type the hex value for the original serial number of the chassis (used when the midplane has been replaced).**

The value 0 is immediately replaced with the hex value of the chassis serial number.

A non-zero value should only be specified if the chassis has been replaced but the original chassis serial number must be retained; this feature is especially important in a Sun Cluster environment to maintain the same disk device names in a cluster.

4. **To implement the revised parameter value, select “system Functions” on the Main Menu, select “Reset Controller” and press Return.**

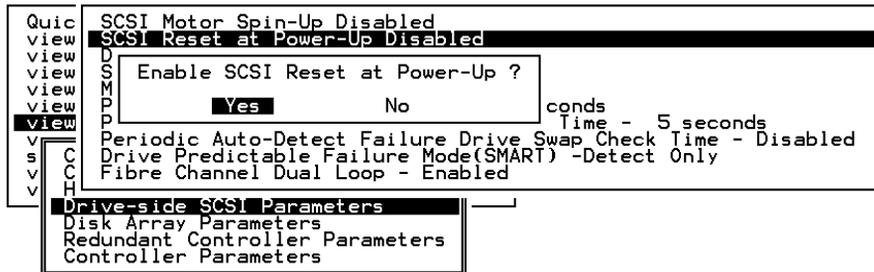
8.3.5 SDRAM ECC Function (Do Not Change)

The default setting is always set to Enabled.

8.4.2 SCSI Reset at Power-Up (Do Not Change)

By default, when the controller is powered up, it will send a SCSI bus reset command to the SCSI bus. When disabled, it will not send a SCSI bus reset command on the next power-up.

When connecting dual host computers to the same SCSI bus, the SCSI bus reset will interrupt all the read/write requests being performed. This may cause some operating environments or host computers to act abnormally. Disable the SCSI Reset at Power-up to avoid this situation.

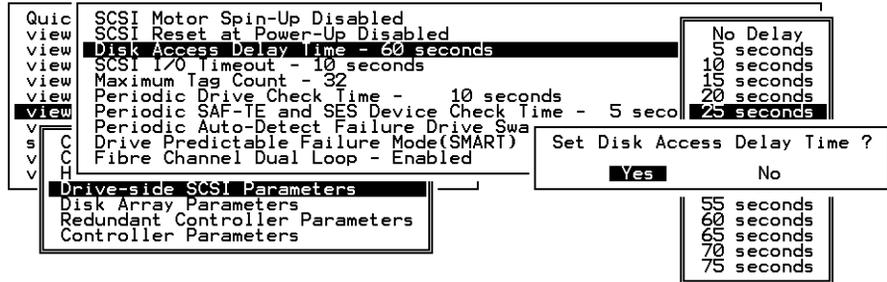


1. Select “view and edit Configuration parameters” on the Main Menu.
2. Select “Drive-side SCSI Parameters,” then press Return.
The Drive-side SCSI parameters menu is displayed.
3. Select “SCSI Reset at Power-Up”, then press Return.
4. Select Yes in the dialog box that follows to confirm the setting.
5. Power off all hard drives and controller, and power them on again.

All the hard drives will not spin-up at this time. The controller will then spin-up the hard drives one by one at a four-second interval.

8.4.3 Disk Access Delay Time

This function sets the delay time before the controller tries to access the hard drives after power-on. The default is 15 seconds. The range is No Delay to 75 seconds.



1. Select “view and edit Configuration parameters” on the Main Menu.
2. Select “Drive-side SCSI Parameters,” then press Return. The Drive-side SCSI parameters menu is displayed.
3. Select “Disk Access Delay Time,” then press Return. A list of selections is displayed.
4. Select the desired delay time, and press Return. Select Yes to confirm the setting.

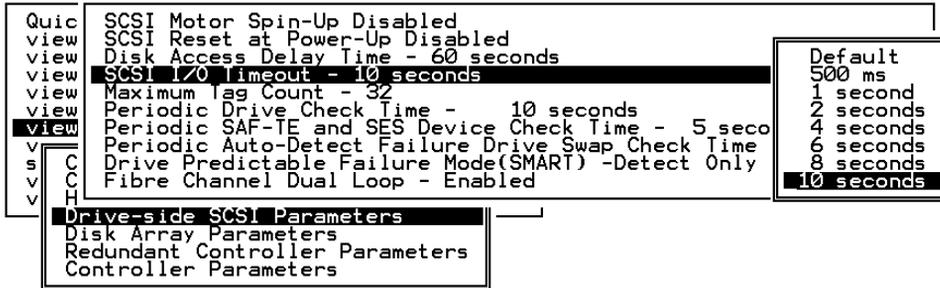
8.4.4 SCSI I/O Timeout

The “SCSI I/O Timeout” is the time interval for the controller to wait for a drive to respond. If the controller attempts to read data from or write data to a drive but the drive does not respond within the SCSI I/O timeout value, the drive will be considered a failed drive.

The default setting for “SCSI I/O Timeout” is 7 seconds. It is recommended not to change this setting. Setting the timeout to a lower value will cause the controller to judge a drive as failed while a drive is still retrying or while a drive is unable to arbitrate the SCSI bus. Setting the timeout to a greater value will cause the controller to keep waiting for a drive, and it may sometimes cause a host timeout.

When the drive detects a media error while reading from the drive platter, it will retry the previous reading or recalibrate the head. When the drive encounters a bad block on the media, it reassigns the bad block to another spare block. However, all of this takes time. The time to perform these operations can vary between different brands and models of drives.

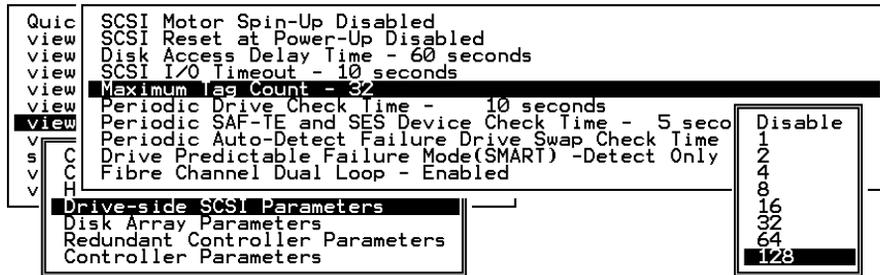
During SCSI bus arbitration, a device with higher priority can utilize the bus first. A device with lower priority will sometimes receive a SCSI I/O timeout when devices of higher priority devices keep utilizing the bus.



1. Select “view and edit Configuration parameters” on the Main Menu.
2. Select “Drive-side SCSI Parameters,” then press Return. The Drive-side SCSI parameters menu is displayed.
3. Select “SCSI I/O Timeout –Default (7 seconds),” then press Return.

A list of selections is displayed. Move the cursor bar on a selection, then press Return. Select Yes in the dialog box that follows to confirm the setting.

8.4.5 Maximum Tag Count (Tag Command Queuing)



This is the maximum number of tags that can be sent to each drive at the same time. A drive has a built-in cache that is used to sort all of the I/O requests (“tags”) that are sent to the drive, allowing the drive to finish the requests faster.

The cache size and maximum number of tags varies between different brands and models of drive. Using the default setting of “32” is highly recommended.

Note – Changing the maximum tag count to “Disable” will cause the Write-Back cache in the hard drive to be ignored (i.e., not used).

The controller supports tag command queuing with an adjustable tag count from 1 to 128. The default setting is “Enabled” with a maximum tag count of 32.

To change the default setting, perform the following steps.

1. Select “view and edit Configuration parameters” on the Main Menu.
2. Select “Drive-side SCSI Parameters,” then press Return. The Drive-side SCSI parameters menu is displayed.
3. Select “Maximum Tag Count”, then press Return. A list of available tag count numbers is displayed.
4. Select a number, then press Return. Select Yes in the dialog box that follows to confirm the setting.
5. For the changes to take effect, select “system Functions,” select “Reset Controller,” and press Return.

8.4.6 SAF-TE and S.E.S. Enclosure Monitoring

```

Quick view SCSI Motor Spin-Up Disabled
view view SCSI Reset at Power-Up Disabled
view view Disk Access Delay Time - 60 seconds
view view SCSI I/O Timeout - 10 seconds
view view Maximum Tag Count - 32
view view Periodic Drive Check Time - 10 seconds
view view Periodic SAF-TE and SES Device Check Time - 5 seco
view view Periodic Auto-Detect Failure Drive Swap Check Time
view view Drive Predictable Failure Mode(SMART) -Detect Only
view view Fibre Channel Dual Loop - Enabled

Drive-side SCSI Parameters
Disk Array Parameters
Redundant Controller Parameters
Controller Parameters
  
```

Disabled
50 ms
100 ms
200 ms
500 ms
1 second
2 seconds
5 seconds
10 seconds
20 seconds
30 seconds
60 seconds

If there are remote devices within your RAID enclosure monitored via SAF-TE/S.E.S., use this function to decide at what interval the controller will check the status of these devices.

Select “Periodic SAF-TE and SES Device Check Time,” and press Return. Move the cursor to the desired interval; then press Return. Select Yes in the dialog box to confirm the setting.

8.4.7 Periodic Drive Check Time

The Periodic Drive Check Time is an interval for the controller to check the drives on the SCSI bus at controller startup (a list of all detected drives can be seen under “view and edit scsi Drives”). The default value is Disabled. Disabled means that if a drive is removed from the bus, the controller will not know that the drive is removed until a host tries to access that drive.

Changing the check time to any other value allows the controller to check at the selected interval all of the drives that are listed under “view and edit scsi Drives.” If any drive is then removed, the controller will know even if a host does not access that drive.

```

Quick view SCSI Motor Spin-Up Disabled
view view SCSI Reset at Power-Up Disabled
view view Disk Access Delay Time - 60 seconds
view view SCSI I/O Timeout - 10 seconds
view view Maximum Tag Count - 32
view view Periodic Drive Check Time - 10 seconds
view view Periodic SAF-TE and SES Device Check Time - 5 seco
view view Periodic Auto-Detect Failure Drive Swap Check Time
view view Drive Predictable Failure Mode(SMART) -Detect Only
view view Fibre Channel Dual Loop - Enabled

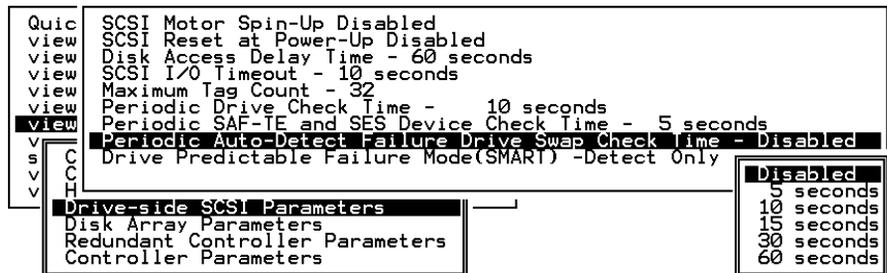
Drive-side SCSI Parameters
Disk Array Parameters
Redundant Controller Parameters
Controller Parameters
  
```

Disable
1/2 seconds
1 second
2 seconds
5 seconds
10 seconds
30 seconds

8.4.8 Auto-Detect Failure Drive Swap Check Time

The Drive-Swap Check Time is the interval at which the controller checks to see whether a failed drive has been swapped. When a logical drive's member drive fails, the controller will detect the failed drive (at the selected time interval). Once the failed drive has been swapped with a drive that has adequate capacity to rebuild the logical drive, the rebuild will begin automatically.

The default setting is Disabled, meaning that the controller will not auto-detect the swap of a failed drive. When the Periodic Drive Check Time is set to Disabled, the controller is not able to detect any drive removal that occurs after the controller has been powered on. The controller only detects drive removal when a host attempts to access the data on the drive.



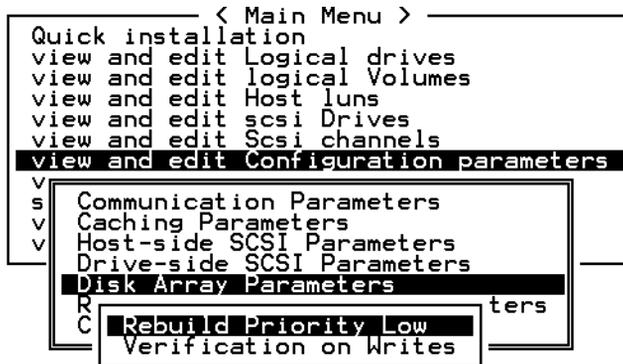
To enable this feature, perform the following steps.

1. Select “Periodic Auto-Detect Failure Drive Swap Check Time,” and press Return.
2. Select the desired interval and press Return. Select Yes in the dialog box that follows to confirm the setting.

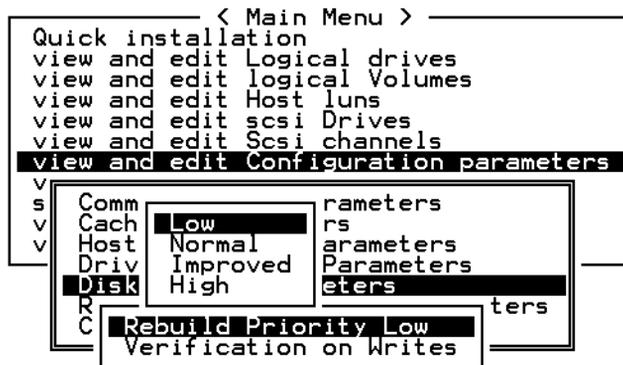
By choosing a time value to enable the Periodic Drive Check Time, the controller polls all connected drives in the controller's drive channels at the assigned interval. Drive removal is detected even if a host does not attempt to access data on the drive.

8.5 Disk Array Parameters

Select “View and edit Configuration parameters” on the Main Menu and press Return. Then select “Disk Array Parameters.”



8.5.1 Rebuild Priority



The RAID controller provides a background rebuilding ability. This means the controller is able to serve other I/O requests while rebuilding the Logical Drives. The time required to rebuild a drive set will largely depend on the total capacity of the Logical Drive being rebuilt. Additionally, the rebuilding process is totally transparent to the host computer or the operating environment.

1. Select “view and edit Configuration parameters,” then “Disk Array Parameters,” then press Return again.

The Disk Array Parameters menu is displayed.

2. Select “Rebuild Priority,” then press Return.

A list of the priority selections (Low, Normal, Improved, or High) is displayed.

The background rebuild process has four priority options:

- Low (the default which uses the controller’s minimum resources to rebuild)
- Normal (to speed up the rebuilding process)
- Improved (to speed up the rebuilding process)
- High (to use the controller’s maximum resources to complete the rebuilding process in the shortest possible time)

3. Select the desired setting, then press Return.

8.5.2 Verification on Writes

Normally, errors may occur when a hard drive writes data. In order to avoid the write error, the controller can force the hard drives to verify the written data. There are three selectable methods:

- Verification on LD Initialization Writes
Performs Verify-after-Write while initializing the logical drive.
- Verification on LD Rebuild Writes
Performs Verify-after-Write during the rebuilding process.
- Verification on LD Normal Drive Writes
Performs Verify-after-Write during normal I/O requests.

Each method can be enabled or disabled individually. Hard drives will perform Verify-after-Write according to the selected method.

Note – The “verification on Normal Drive Writes” method will affect the write performance during normal use.

To select the type of verification wanted, perform the following steps:

1. Select “view and edit Configuration parameters,” then “Disk Array Parameters,” then press Return again.

The Disk Array Parameters menu is displayed.

2. Press Return on the “Verification on Writes” in the “Disk Array Parameters” menu.

The items for selection is displayed on screen.

```

      < Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
v
s
v
v
Communication Parameters
Caching Parameters
Host
Driv Verification on LD Initialization Writes Disabled
Disk Verification on LD Rebuild Writes Disabled
R Verification on Normal Drive Writes Disabled
C
R
C Verification on Writes
  
```

3. Select the desired item, then press Return.

```

      < Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
v
s
v
v
Communication Parameters
Caching Parameters
Host
Driv Verification on LD Initialization Writes Disabled
Disk Verification on LD Rebuild Writes Disabled
R Verification on Normal Drive Writes Disabled
C
R
C Verification on Writes
V
Enable Initialize RAID with Verify Data ?
Ver: Yes No
  
```

4. Select Yes in the confirm box to enable or disable the function.
Follow the same procedure to enable or disable each method.

8.6 Host-side SCSI Parameters

Select the “view and edit Configuration parameters” command, select “Host-side SCSI Parameters,” and press Return to access the following parameters:

- Maximum Queued I/O Count
- LUNs per Host SCSI ID
- Max Number of Concurrent Host-LUN Connection - Def (4)
- Number of Tags Reserved for each Host-LUN Connection
- Peripheral Device Type Parameters
- Host Cylinder/Head/Sector Mapping Configuration

8.6.1 Overview of SCSI Channel, SCSI ID and LUN

A SCSI channel (SCSI bus) can connect up to 15 devices (not including the SCSI controller itself) when the Wide function is enabled (16-bit SCSI).

It can connect up to 7 devices (not including the controller itself) when the Wide function is disabled (8-bit SCSI).

Each device has one unique SCSI ID. Two devices owning the same SCSI ID are not allowed.

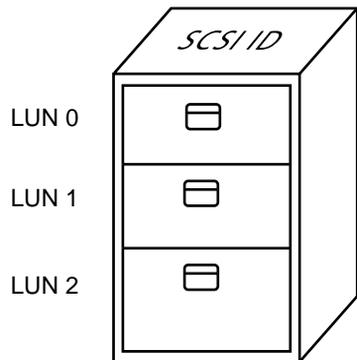


FIGURE 8-5 Filing Cabinet Represents The SCSI ID and File Drawers Represent the Luns.

From a SCSI's point of view, a SCSI ID is like a cabinet, and the drawers are the LUNs (Logical Unit Numbers). Each SCSI ID can have up to 32 LUNs. Data can be stored into one of the LUNs of the SCSI ID. Most SCSI host adapters treat a LUN like another SCSI device.

8.6.2 Maximum Concurrent Host-LUN Connections

This function is used to set the maximum number of concurrent host-LUN connections. The default setting is 4 LUNs, with a predefined range of 1 to 64.

Note – This function should only be changed if you have more than four logical drives or partitions. Increasing this number may increase your performance.

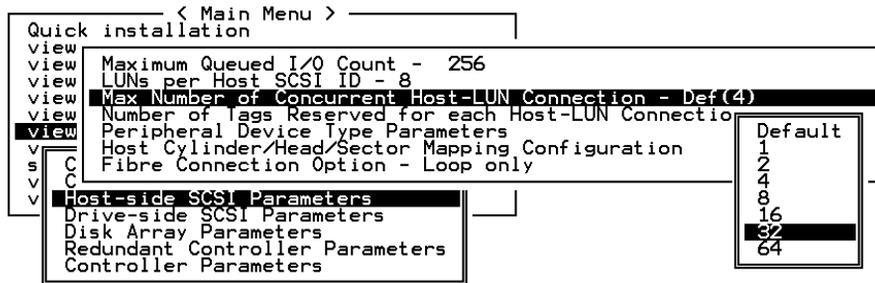
Maximum concurrent host LUN connection (nexus in SCSI) is the arrangement of the controller internal resources for use with a number of the current host nexus.

For example, you can have four hosts (A, B, C, and D) and four host IDs/LUNs (ID 0, 1, 2 and 3) in a configuration where:

- host A accesses ID 0 (one nexus)
- host B accesses ID 1 (one nexus)
- host C accesses ID 2 (one nexus)
- host D accesses ID 3 (one nexus)

These connections are all queued in the cache and are called four nexus.

If there are I/Os in the cache with 4 different nexus, and another host I/O comes with a nexus different than the four in the cache (for example, host A accesses ID 3), the controller will return busy. This occurs with the concurrent active nexus; if the cache is cleared, it will accept four different nexus again. Many I/Os can be accessed via the same nexus.



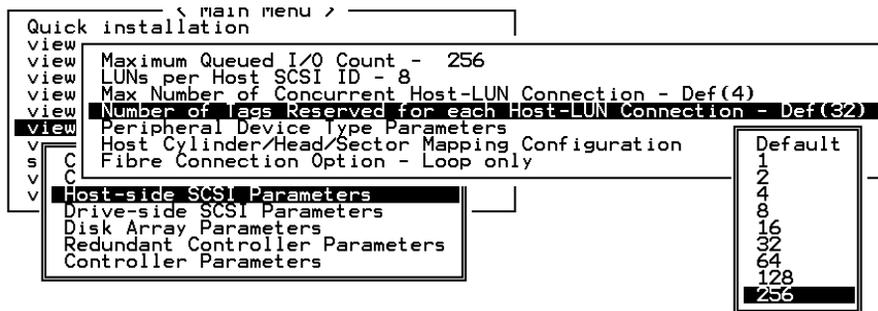
To change the default number of nexus for a host (the default setting is 4), perform the following steps:

1. From the Main Menu, select “view and edit Configuration parameters,” “Host-side SCSI Parameters,” then press **Return**.
2. Select “Max Number of Concurrent Host-LUN Connection,” then press **Return**.
A list of available selections is displayed. Select an item, then press **Return**.
3. Select **Yes** in the dialog box that follows to confirm the setting.

8.6.3 Number of Tags Reserved for Each Host LUN Connection

This function is used to modify the tag command queuing on the Host-LUN connection. The default setting is 32 tags, with a predefined range of 1 to 256. Unless necessary, the default factory setting should not be changed.

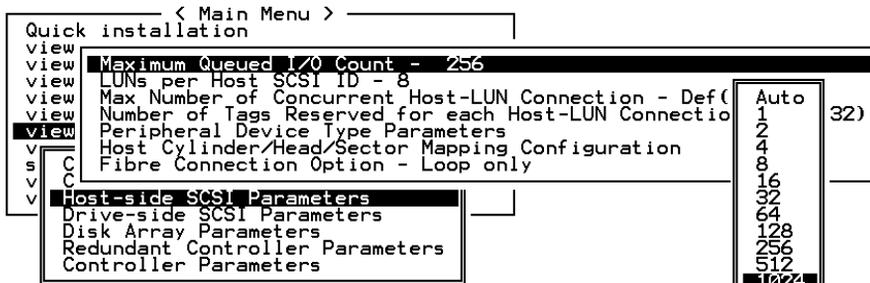
Each “nexus” has “32” (the default setting) tags reserved. This setting ensures that the controller accepts at least 32 tags per nexus. The controller will be able to accept more than that as long as the controller internal resources allow it; if the controller does not have enough internal resources, at least 32 tags can be accepted per nexus.



1. From the Main Menu, select “view and edit Configuration parameters,” “Host-side SCSI Parameters,” then press Return.
2. Select “Number of Tags Reserved for each Host-LUN Connection,” then press Return. A list of available selections is displayed.
3. Select an item, then press Return. Select Yes in the dialog box that follows to confirm the setting.

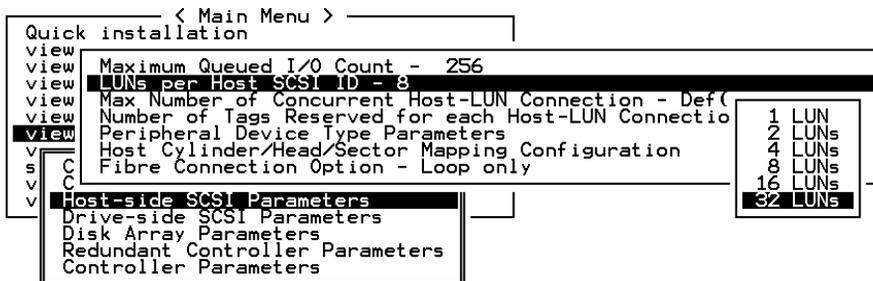
8.6.4 Maximum Queued I/O Count

This function allows you to configure the maximum size of the I/O queue the controller can accept from the host computer by byte size. The predefined range is from 1 to 1024 bytes, or you may choose the “Auto” (automatically configured) mode. The default value is 256 bytes.



1. From the Main Menu, select “view and edit Configuration parameters,” “Host-side SCSI Parameters,” then press Return.
2. Select “Maximum Queued I/O Count,” then press Return. A list of available selections is displayed.
3. Select an item, then press Return. Select Yes in the dialog box that follows to confirm the setting.

8.6.5 LUNs Per Host SCSI ID



This function is used to change the number of LUNs per host SCSI ID. The default setting is 32 LUNs, with a predefined range of 1 to 32 LUNs per logical drive available. A maximum of 128 LUNs per array is allowed.

1. From the Main Menu, select “view and edit Configuration parameters,” “Host-side SCSI Parameters,” then press Return.
2. Select “LUNs per Host SCSI ID,” then press Return. A list of selections is displayed.
3. Select an item, then press Return. Select Yes in the dialog box that follows to confirm the setting.

8.6.6 Cylinder/Head/Sector Mapping

In SCSI arrays, the drive capacity is decided by the number of blocks. Some operating environments read the capacity of the array based on the cylinder/head/sector count of the drives.

For Sun Solaris, the cylinder cannot exceed 65535, so you can choose cylinder < 65535. The controller automatically adjusts the head/sector count, and then the operating environment can read the correct drive capacity. For a Sun Solaris configuration, the recommended values are listed in the following table.

TABLE 8-1 Cylinder/Head Sector Mapping under Sun Solaris

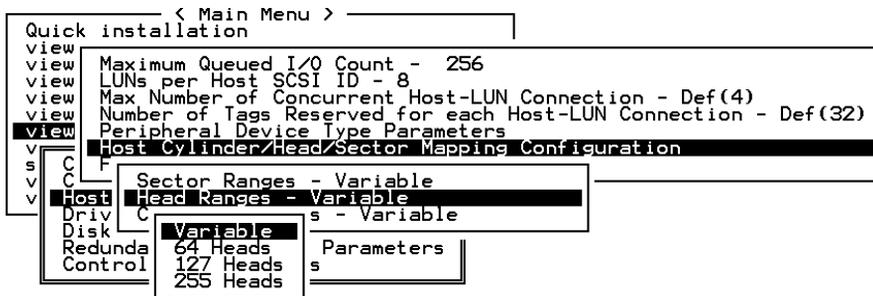
Capacity	Cylinder	Head	Sector
< 64 GB	?	64	32
64 - 128 GB	?	64	64
128 - 256 GB	?	127	64
256 - 512 GB	?	127	127
512 GB - 1 TB	?	255	127

Currently, Solaris does not support drive capacity larger than 1 terabyte.

To configure Sector Ranges, Head Ranges, and Cylinder Ranges, perform the following steps.

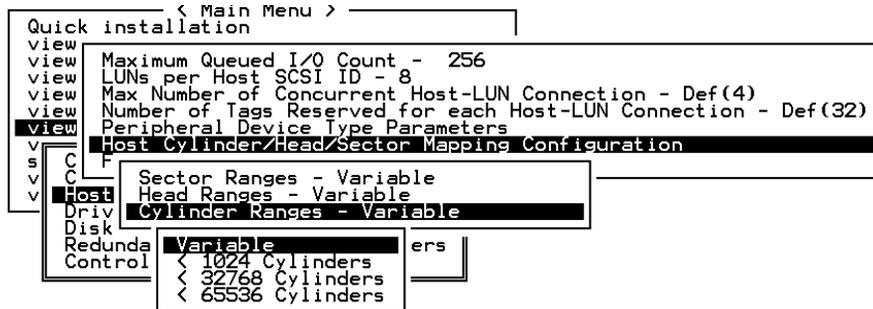
1. Select “view and edit Configuration parameters” on the Main Menu.
2. Select “Host-Side SCSI Parameters,” then select “Host Cylinder/Head/Sector Mapping Configuration.”

3. Select Sector Ranges and choose the desired value.



4. Select Head Ranges and the desired value.

5. Select Cylinder Ranges and the desired value.



8.7 Peripheral Device Type Parameters

The functions included within this section are provided for arrays without a preset logical RAID unit, connected to a host. Subsequently, the majority of these functions do not apply to the Sun StorEdge product line.



Caution – Modification of these settings will likely cause problems with your storage array. All parameters are typically preset.

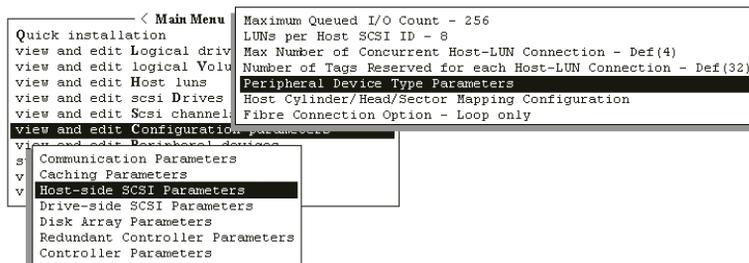


FIGURE 8-6 Peripheral Device Type Parameters

Peripheral Device Qualifier: The default setting for this parameter is “Connected.”

Device Supports Removable Media: The default setting for this parameter is “Disabled.”

LUN Applicability: The default setting is “All Undefined LUNs.”

Related to the other LUN functions, LUN Applicability is primarily used when no logical drive has been created and mapped to a host LUN, and the RAID controller is the only device connecting to the host SCSI card. For such instances, the operating environment will typically not load the driver of the host SCSI adapter. If the driver is not loaded, the in-band SCSI utility cannot communicate to the RAID controller.

If “LUN-0's only” is selected, only LUN-0 of the host ID is displayed as a device with the user-defined peripheral device type. If “All Undefined LUNs” is selected, each LUN in that host ID is displayed as a device with the user-defined peripheral device type.

8.8 Setting an IP Address

The controller Ethernet port offers out-of-band management through two programs:

- The Configuration Service program. Refer to the *Sun StorEdge 3310 SCSI Configuration Service User Guide* for details.
- The firmware program when you Telnet to the controller IP address to enter the firmware program.

To access the array using the Ethernet port, you must set up an IP address for the controller.

To set the IP address, netmask and gateway values of the RAID controller, perform the following steps.

- 1. Access the array through the COM port on the controller module of the array.**
- 2. On the Main Menu, select “view and edit Configuration parameters.”**
- 3. Select “Communication Parameters,” then select “Internet Protocol (TCP/IP).”**
- 4. Press Return on the chip hardware address, and then select “Set IP Address.”**
- 5. Enter the desired IP address, NetMask, and Gateway values.**
- 6. Reset the controller for the configuration to take effect:**

Select “system Functions” from the Main Menu, then select “Reset controller” and press Return.

Viewing and Editing Peripheral Devices

This chapter describes viewing and editing parameters for peripheral devices. Topics covered include:

- “Viewing Peripheral Device Controller Status” on page 9-2
- “Viewing Peripheral Device SAF-TE Status” on page 9-2
- “Setting Peripheral Device Entry” on page 9-5
- “Setting the UPS Power Fail Signal” on page 9-7
- “Viewing the Controller Voltage and Temperature Status Window” on page 9-8

```
      < Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
system Functions
v View Peripheral Device Status
v Set Peripheral Device Entry
  Define Peripheral Device Active Signal
  Adjust LCD Contrast
  Controller Peripheral Device Configuration
```

FIGURE 9-1 View and Edit Peripheral Devices Menu

9.1 Viewing Peripheral Device Controller Status

To view the status of each controller, perform the following steps.

1. Select “view and edit Peripheral devices” on the Main Menu.
2. Select the “View Peripheral Device Status” option. A table will display the status.

The screenshot shows a terminal window with the following content:

```
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
```

The "view and edit Peripheral devices" option is highlighted. Below it, the "View Peripheral Device Status" option is selected, leading to a table:

ITEM	STATUS	LOCATION
Redundant Controller	Enabled	Primary
SAF-TE Device	Operational	Channel 0 ID 14

FIGURE 9-2 View Peripheral Device Status

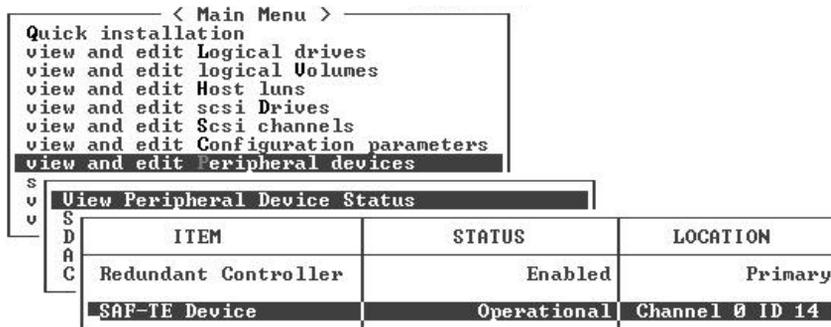
9.2 Viewing Peripheral Device SAF-TE Status

To check the status of SAF-TE components (temperature sensors, cooling fans, the beeper speaker, power supplies, and slot status), perform the following steps.

The SAF-TE controller is located on the SCSI I/O module.

1. From the Main Menu select “view and edit Peripheral devices,” and press Return.

2. Select “View Peripheral Device Status,” and press Return.



3. Select “SAF-TE Device,” and press Return to view the status of temperature sensors, power supplies, beeper speaker, and cooling fans.

The temperature sensor displays the current temperature of each sensor in degrees Fahrenheit.

The drive slot status indicates that a slot is filled by displaying a SCSI ID number:

- Single bus configuration: 0 through 13 ID numbers if all 12 drives are filled. SCSI IDs 6 and 7 are reserved for host communication. If there is an empty slot, the message “No Device Inserted” is displayed. See FIGURE 9-3.
- Dual-bus configuration (not supported): the message “No Device Inserted” for the six drives on one channel and six ID numbers for the second channel. See FIGURE 9-4.

To check that you have all slots filled in a dual bus configuration, see “SCSI Drive Status Table” on page 6-2 and check the column labeled “Ch1 ID.”

Product ID	StorEdge 2310 A	Drive Slot 1	SCSI ID 1
Revision Level	0.62	Drive Slot 2	SCSI ID 2
Unique ID	3030303132323338	Drive Slot 3	SCSI ID 3
Cooling Fan 0	Operational	Drive Slot 4	SCSI ID 4
Cooling Fan 1	Operational	Drive Slot 5	SCSI ID 5
Power Supply 0	Operational and On	Drive Slot 6	SCSI ID 8
Power Supply 1	Operational and On	Drive Slot 7	SCSI ID 9
Temp Sensor 0	78	Drive Slot 8	SCSI ID 10
Temp Sensor 1	78	Drive Slot 9	SCSI ID 11
Temp Sensor 2	80	Drive Slot 10	SCSI ID 12
Temp Sensor 3	86	Drive Slot 11	SCSI ID 13
Temp Sensor 4	91		
Temp Sensor 5	82		
Temp Sensor 6	82		
Temp Alert	Normal		
Speaker Status	Off or No Speaker		
Drive Slot 0	SCSI ID 0		

FIGURE 9-3 Example of SAF-TE Device Status Window in a Single Bus Configuration

In a dual bus configuration example, the below SAF-TE window displays “No Device Inserted” for six drives which are actually inserted into slots. The SAF-TE protocol does not support a dual-bus configuration and only recognizes one bus (half the drives) if you have a dual-bus configuration.

Product ID	StorEdge 3310 A	Drive Slot 1	No Device Inserted
Revision Level	A000	Drive Slot 2	No Device Inserted
Unique ID	3132333435362020	Drive Slot 3	No Device Inserted
Cooling Fan 0	Operational	Drive Slot 4	No Device Inserted
Cooling Fan 1	Operational	Drive Slot 5	No Device Inserted
Power Supply 0	Operational and On	Drive Slot 6	SCSI ID 0
Power Supply 1	Operational and On	Drive Slot 7	SCSI ID 1
Temp Sensor 0	89	Drive Slot 8	SCSI ID 2
Temp Sensor 1	86	Drive Slot 9	SCSI ID 3
Temp Sensor 2	82	Drive Slot 10	SCSI ID 4
Temp Sensor 3	77	Drive Slot 11	SCSI ID 5
Temp Sensor 4	82		
Temp Sensor 5	84		
Temp Sensor 6	82		
Temp Alert	Normal		
Speaker Status	Off or No Speaker		
Drive Slot 0	No Device Inserted		

FIGURE 9-4 Example of SAF-TE Device Status Window in a Dual Bus Configuration

9.3 Setting Peripheral Device Entry

The menu functions within the Set Peripheral Device Entry include the following:

- Redundant Controller
- UPS Status

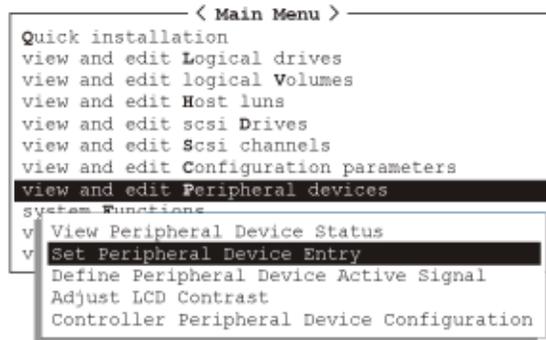


FIGURE 9-5 Set Peripheral Device Entry

9.3.1 Redundant Controller Mode (Do Not Change)

The redundant controller mode is automatically set to Enabled. Do not change this setting.

For more information about redundant controller operation, refer to “Controller Defaults and Limitations” on page 1-13.

9.3.2 Enable UPS Status

This function is used to enable the Uninterruptible Power Supply (UPS) status if a UPS unit is installed for power redundancy and backup. The default value for this function is Disabled.

1. Select the “UPS Status” option and press Return.

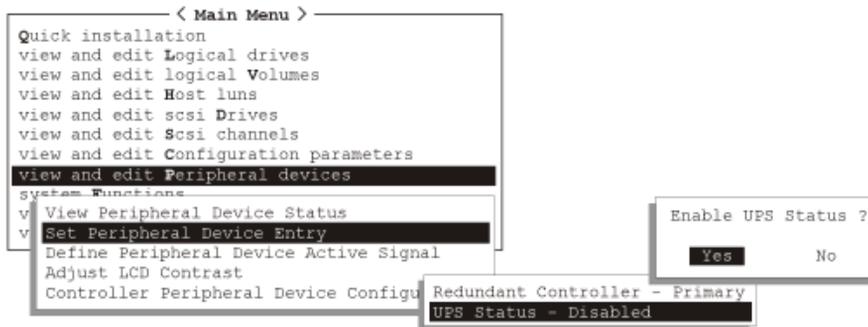


FIGURE 9-6 UPS Status

2. A confirmation prompt will be displayed. Select “Yes,” and press Return to confirm.

9.4 Setting the UPS Power Fail Signal

The “UPS Power Fail Signal” function is to prioritize the alert notification level if power to your UPS device should fail for any reason. The default High priority should not be changed.

1. On the Main Menu, select “view and edit Peripheral devices.”
2. Select the “Define Peripheral Device Active Signal” command.

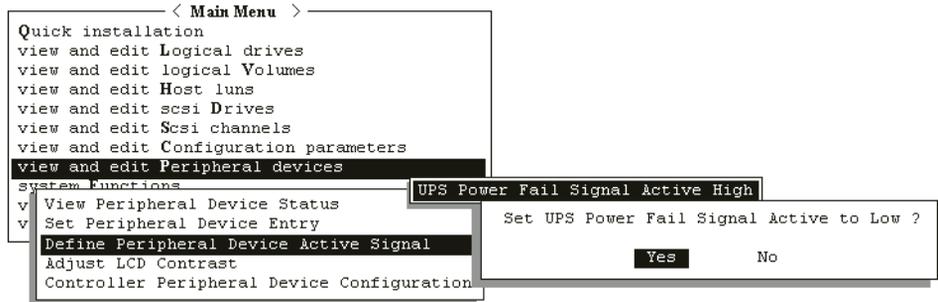


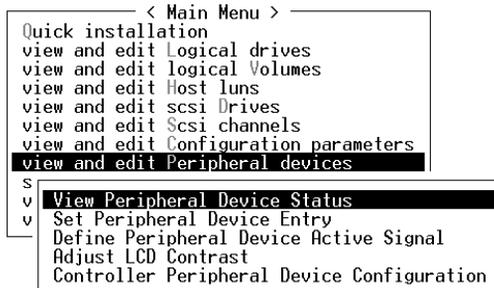
FIGURE 9-7 UPS Power Fail Signal

3. If you want to change the setting, select the “UPS Power Fail Signal” function and press Return.
4. A prompt will be displayed. Select “Yes” to change the setting.

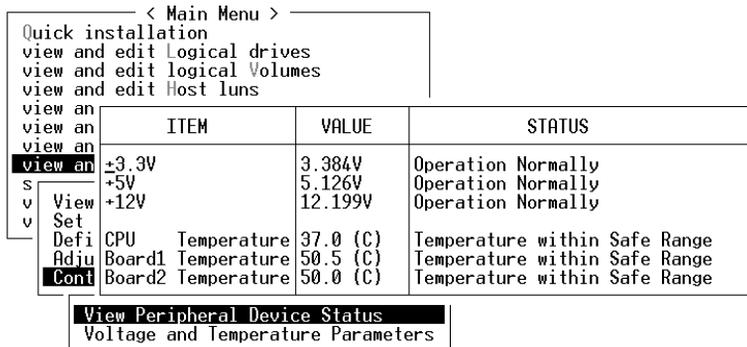
9.5 Viewing the Controller Voltage and Temperature Status Window

To check the status of controller voltage and temperature, perform the following steps.

1. From the Main Menu, select “view and edit Peripheral devices,” and press Return.



2. Select “Controller Peripheral Device Configuration,” and press Return.
3. Select “View Peripheral Device Status,” and press Return to view voltage and temperature status of the RAID unit.



The components checked for voltage and temperature will be displayed on screen and will be defined as normal or out-of-order.

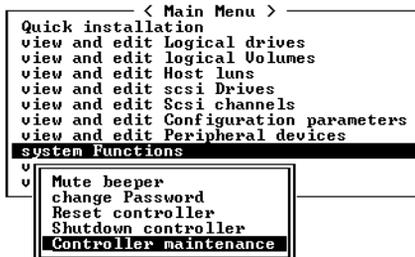
System Functions, Info, and Event Logs

This chapter contains a discussion of system functions, configuration information, and describes how to view event logs. Topics covered include:

- “System Functions” on page 10-2
 - “Muting the Beeper” on page 10-2
 - “Setting a New Password” on page 10-3
 - “Changing Password” on page 10-3
 - “Disabling the Password” on page 10-4
 - “Resetting Controller” on page 10-4
 - “Shutting Down the Controller (Restricted Use)” on page 10-5
 - “Restoring Configuration (NVRAM) from a File” on page 10-5
- “Saving Configuration (NVRAM) to Disk” on page 10-6
- “Viewing Event Logs on Screen” on page 10-7

10.1 System Functions

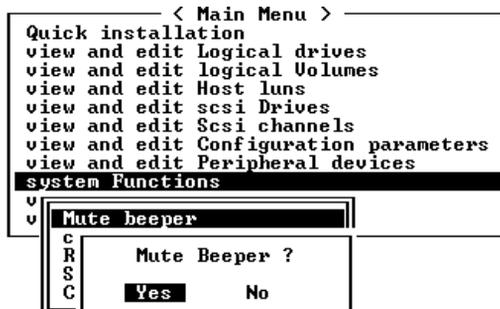
Select “system Functions” in the Main Menu, then press Return. The system Functions menu is displayed. Select an item, then press **Return**.



10.1.1 Muting the Beeper

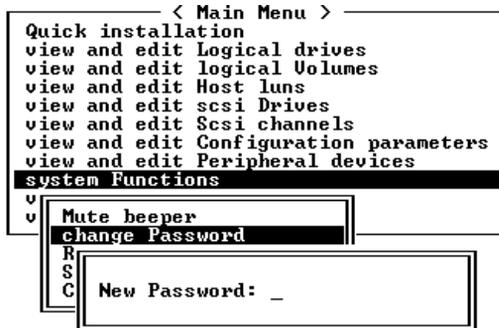
To change the beeper setting, perform the following steps.

1. Select “system Functions” in the Main Menu, then press Return.
2. When the controller’s beeper has been activated, choose “Mute beeper,” then press Return.



3. Select “Yes” and press Return in the next dialog box to turn the beeper off temporarily for the current event. The beeper will still be activated on the next event.

10.1.2 Setting a New Password



To change the password, perform the following steps.

1. Select “system Functions” in the Main Menu, then press Return, and select “change Password.”
2. Enter the desired password in the column, then press Return. The next dialog box will display “Re-Enter Password”. Enter the password again to confirm and press Return.

The new password will now become the controller’s password. Providing the correct password may be necessary when entering the Main Menu from the Initial screen.

10.1.3 Changing Password

1. Select “system Functions” in the Main Menu, then press Return.

Use the controller’s password to protect the array from unauthorized entry. Once the controller’s password has been set, the user can only configure and monitor the RAID controller by providing the correct password.

Rules

- The controller is able to verify the password when entering the Main Menu from the initial screen or making configuration change. If the controller will be left unattended, the “Password Validation Timeout” can be set to “Always Check.” Setting validation timeout to “always check” protects the controller configuration from any unauthorized change.
- The controller password and controller name share a 16-character space. The maximum characters for the controller password is 15. When the controller name occupies 15 characters, there is only one character left for the controller password, and vice versa.

2. To set or change the controller password, move the cursor bar to “Change Password,” then press Return.

If a password has previously been set, the controller will ask for the old password first. If the password has not been set, the controller will ask for the new password. The password cannot be replaced unless a correct old password is provided.

3. Enter the old password, then press Return.

If the password is incorrect, it will not allow you to change the password. Instead, it will display the message *Password incorrect!* then return to the previous menu.

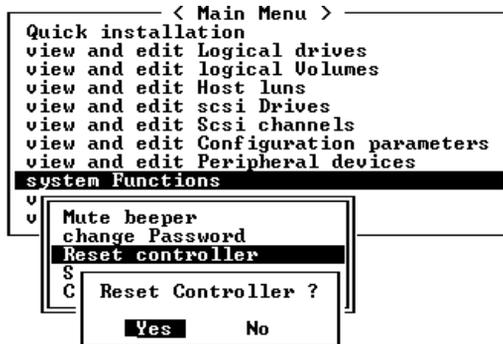
If the password is correct, or there is no preset password, it will ask for the new password.

10.1.4 Disabling the Password

To disable or delete the password, press Return only in the password column that is used for entering a new password. The existing password will be deleted. No password checking will occur when entering the Main Menu from the Initial screen.

10.1.5 Resetting Controller

1. Select “system Functions” in the Main Menu, then press Return.
2. To reset the controller without powering off the array, select “Reset Controller,” then press Return.



3. Select Yes in the dialog box that follows, then press Return.

The controller will now reset as well as power-off or re-power-on.

10.1.6 Shutting Down the Controller (Restricted Use)

Do NOT use this command unless directed by Technical Support. The “Shutdown Controller” command erases the cache content.

10.2 Restoring Configuration (NVRAM) from a File

If you have saved a configuration file and want to apply the same configuration to another array or re-apply it to the array which had the configuration originally, you must be certain that the channels and SCSI IDs in the configuration file are correct for the array where you are restoring the configuration.

The NVRAM configuration file will restore all configuration settings (channel settings, host IDs, etc.) but does not rebuild logical drives. To save a configuration file, refer to “Saving Configuration (NVRAM) to Disk” on page 10-6.



Caution – If the channels or SCIS IDs are not a correct match for the array, you will lose access to the mismatched channels or drives when you restore the configuration with the configuration file.

Note – In the Configuration Service program, you can save a configuration file which can restore all configurations and rebuild all logical drives. However, it will also erase all data when it rebuilds all logical drives so operation is only performed when no data has been stored or all data has been transferred to another array.

To restore configuration settings from a saved NVRAM file, perform the following steps.

1. From the Main Menu, choose “system functions.”
2. Select “Controller maintenance” and press Return.
3. Select “Restore NVRAM from disks,” and press Return.
4. Press Yes to confirm.

A prompt will notify you that the controller NVRAM data has been successfully restored from disks.

10.3 Saving Configuration (NVRAM) to Disk

This function is used to save controller-dependent configuration information and is highly recommended whenever a configuration change is made.

The logical configuration information will be stored within the logical drive.

Note – A logical drive must exist for the controller to write NVRAM content onto it.

1. From the Main Menu, select “system Functions.”
2. Use arrow keys to scroll down and select “Controller Maintenance,” “save nvram to disks,” then press Return.



3. Select Yes to confirm.

A prompt will inform you that NVRAM information has been successfully saved.

To restore the configuration, refer to “Restoring Configuration (NVRAM) from a File” on page 10-5.

10.4 Viewing Event Logs on Screen

A controller event log records an event or alarm which occurs after the system is powered on

Note – The Event Monitoring Units in each RAID unit and each Expansion Unit send messages to the controller log which report problems and status of the fans, temperature, and voltage.



Caution – Powering off or resetting the controller will automatically delete all recorded event logs.

To view the event logs on screen, select “view and edit Event logs” on the Main Menu and press Return.

```
      < Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
system Functions
view system Information
view and edit Event logs
```

The controller can store up to 1000 event logs. An event log can record a configuration or operation event as well as an error message or alarm event.

TABLE 10-1 Example Event Logs

[0181]	Controller Initialization Completed
[2181]	LG:0 Logical Drive NOTICE: Starting Initialization
[2182]	Initialization of Logical Drive 0 Completed
[2181]	LG:1 Logical Drive NOTICE: Starting Initialization
[2182]	Initialization of Logical Drive 2 Completed

To clear the saved event logs, scroll the cursor down to the last event and press Return. A prompt is displayed “Clear Above xx Event Logs?” Choose Yes to clear the recorded event logs.

Firmware Specifications

This appendix contains the firmware specifications in feature summary tables.

TABLE A-1 Basic RAID Management Features

Feature	Description
RAID Levels	0, 1(0+1), 3, 5, 10, 30, and 50. Enhanced RAID Levels supported (with logical volume implementation)
Maximum Number of logical drives	8
RAID level dependency to each logical drive	Independent. Logical drive configured in different RAID levels can co-exist in an array
Maximum drive number of each logical drive	31 (RAID 3 or 5); 45 (RAID 0 or NRAID); 44 (RAID 1)
Logical drive identification	Unique, Controller Randomly generated logical drive ID; Logical drive name user-configurable
Maximum partitions of each logical drive	32
Maximum number of logical drives in a logical volume	8
Maximum number of logical volumes	8
Maximum number of LUN per Host ID	Up to 32, user configurable
Concurrent I/O	Supported
Tag Command Queuing	Supported
Dedicated Spare Drive	Supported, defined as the spare drive specifically assigned to a logical drive
Global Spare Drive	Supported, the spare drive serving all logical drives

TABLE A-1 Basic RAID Management Features

Co-existing Dedicated and Global Spare Drives	Supported
Auto-rebuild onto spare drive	Supported
Auto-scan of replacement drive upon manually initiated rebuild	Supported
Auto-rebuild onto failed drive replacement	Supported. With no spare drive assigned, the controller will auto-scan the failed drive and starts to rebuild automatically once the failed drive has been replaced.
Background firmware download	Firmware can be downloaded during active I/Os. Administrator may find appropriate time to reset controller later.
Auto recovery from logical drive failure	Supported. When user accidentally removes the wrong drive to cause the 2 nd drive failure of a one-drive-failed RAID5 / RAID3 logical drive, switch off the controller, put the drive back and power on the controller. The logical drive will be restored to one-drive-failed status.

TABLE A-2 Caching Operation Features

Feature	Description
Write-back Cache	Supported.
Write-through Cache	Supported.
Supported Memory type	SDRAM memory for enhanced performance Fast Page Memory with Parity for enhanced data security
Read-ahead Operation	Intelligent Dynamic read-ahead operation for sequential data accessing
Multi-Threaded Operation	Yes
Scatter / Gather	Supported
I/O sorting	Supported. Optimized I/O sorting for enhanced performance
Variable Stripe Size	RAID5: Optimization for Random I/O (32k), optimization for Sequential I/O (128k), user selectable. RAID3: Optimization for Random I/O (4k), optimization for Sequential I/O (16k), user selectable.

TABLE A-3 Advanced Features

Feature	Description
Drive Low-level format	Supported.
Drive Identification	Supported. Force the drive to light on the activity indicator for user to recognize the correct drive.
Drive Information Listing	Supported. Drive vendor name, model number, firmware revision, capacity (blocks), serial number, narrow/wide and current sync. speed.
Drive Read/Write testing	Supported.
Configuration on Disk	Supported. The logical drive information is recorded on drive media.
Save/ Restore NVRAM to / from Disks	Supported. Save all the settings stored in the controller NVRAM to the logical drive members
Save / Restore NVRAM to / from file	Supported. Save all the settings stored in the controller NVRAM to a file (via GUI or TextRAID manager) on user's computer.
Host LUN Geometry User Configurable Default Geometry:	Capacity <64GB: Head=63, Sector=32, Cylinder=? (depends on capacity) 64GB<capacity<128GB:Head=64, Sector=64, Cylinder=? (depends on capacity) 128GB<capacity<256GB: Head=127, Sector=64, Cylinder=? ? (depends on capacity) 256GB<capacity<512GB: Head=127, Sector=127, Cylinder=? 512GB<capacity<1TB: Head=255, Sector=64, Cylinder=? (depends on capacity) 1TB<capacity: Head=225, Sector=225, Cylinder=? (depends on capacity)
User Configurable Geometry range:	Sector: 32,64,127,255 or Variable Head: 64,127,255 or Variable Cylinder: <1024, <32784,<65536 or Variable
Drive Motor Spin-up	Supported. The controller will send spin-up (start unit) command to each drive at 4 sec. intervals.
Drive-side Tag Command Queue	Supported. User adjustable up to 128 for each drive
Host-side Maximum Queued I/O count	User adjustable up to 1024

TABLE A-3 Advanced Features

Maximum concurrent Host LUN connection	User adjustable up to 64
Number of Tags Reserved for each Host-LUN connection	User adjustable up to 256
Drive I/O timeout	User adjustable

TABLE A-4 RAID Expansion Features

Feature	Description
On-line RAID Expansion	Supported.
Mode-1 RAID Expansion-add Drive	Supported. Multiple drives can be added concurrently.
Mode-2 RAID Expansion – Copy and Replace drives	Supported. Replace members with drives of larger capacity.
Expand Capacity with no extra drive bays required	Supported in Mode 2 RAID expansion. Provide “Copy and Replace Drive” function to replace drives with drives of greater capacity. No need to add another enclosure for the extra drives.
Operating environment support for RAID Expansion	No operating environment driver required. No software has to be installed for this purpose.

TABLE A-5 Redundant Controller Features

Feature	Description
Active-active redundant controller	Supported
Synchronized cache for both controllers	Supported. Through a dedicated synchronizing channel, synchronized cache over Fibre loops is supported.
Write-back cache enabled in redundant controller mode	Yes; with synchronized cache connection between controllers.
Automatic Failover	Yes for all Power PC controllers (user's interaction necessary)
Automatic Failback	Yes for all Power PC controllers (user's interaction necessary)

TABLE A-5 Redundant Controller Features

Controller Hot-Swap	<p>No need to shut down the failed controller before replacing the failed controller. (Customer's design-in hot-swap mechanism necessary)</p> <p>Support on-line hot-swap of the failed controller. There is no need to reset or shutdown the failed controller. One controller can be pulled out during active I/Os to simulate the destructive controller failure. (Customer's design-in hot-swap mechanism necessary)</p>
Redundant Controller Communication channel	Sentinel RAID: SCSI; RCC Reset signals built-in EonRAID: Fibre channel(s); RCC cable necessary
No Single-point-of-failure	Supported.
Automatic engagement of replacement controller	Supported in PowerPC series
Dynamic cache memory allocation	Yes. Cache memory is dynamically allocated, not fixed.
Environment management	Supported. SAF-TE, S.E.S., ISEMS (I ² C interface); and on-board controller voltage/temp monitor are all supported in both single and redundant controller mode. In the event of controller failure, serves can be taken over by the existing controller.
Cache battery backup	Supported. Battery backup solutions for cache memory are supported in both single controller and redundant modes.
Load sharing	Supported. Workload can be flexibly divided between different controllers by assigning logical configurations of drives (LDs/LVs) to different controllers.
User-configurable channel mode	Supported. Channel modes configurable (SCSI or Fibre) as HOST or DRIVE in both single controller and redundant controller mode.
Redundant Controller rolling firmware upgrade	Firmware upgrade can be downloaded to the primary controller and then be adopted by both controllers, without interrupting host I/O.
Redundant Controller firmware synchronization	In the event of controller failure, a replacement controller running a different version of firmware can restore a redundant array with a failed controller. Different firmware versions can be auto-synchronized later.

TABLE A-6 Data Safety Features

Feature	Description
Regenerate Parity of logical drives	Supported. Can be performed periodically by the user to ensure that bad sectors do not cause data loss in the event of drive failure.
Bad block auto-reassignment	Supported. Automatic reassignment of bad block
Battery backup for cache memory	Supported. The battery backup solutions provide long-lasting battery support to the cache memory when power failure occurs. The unwritten data in the cache memory can be committed to drive media when power is restored.
Verification on Normal Writes	Supported. Performs read-after-write during normal write processes to ensure data is properly written to drives.
Verification on Rebuild Writes	Supported. Performs read-after-write during rebuild write to ensure data is properly written to drives.
Verification on LD initialization writes	Supported. Performs read-after-write during logical drive initialization to ensure data is properly written to drives.
Drive S.M.A.R.T. support	Not used. Default: Disabled.
Clone Failing Drive	Users may choose to clone data from a failing drive to a backup drive manually

TABLE A-7 User Interface Features

RS-232C Terminal	Supports terminal modes: ANSI, VT-100, ANSI Color. Provides menu-driven user-friendly text-based interface.
Buzzer alarm	Warns user when any failure or critical event occurs.

Parameter Summary Tables

This appendix explains how to fine-tune your array by making minor configuration changes to parameters. It also lists parameter defaults that should not be changed. Topics covered include:

- “Optimizing Controller Parameters” on page B-1
- “Keeping Specific Parameter Defaults” on page B-3

B.1 Optimizing Controller Parameters

Although the factory defaults provide optimized controller operation, you may want to fine-tune your array through minor modifications to the parameters listed below.

Some parameters can only be changed at the initial array configuration; other parameters may be changed at any time. Please note the listing of parameters whose defaults should not be changed unless directed by Technical Support.

TABLE B-1 Controller Parameter Settings

User-Defined Parameter	Default	Alternate Settings
Fault Management:		
③ Clone Failing Drive	Manual function	Replace After Clone Perpetual Clone
① Rebuild Priority	Low (higher priority requires more array resource)	Low, Normal, Improved, High
① Verification on Write	Disabled	On LD Initialization On LD Rebuild On Normal Drive Writes
① Event Notification	Reports to user interface and on board alarm	Over Dial-out Modem Over SNMP Trap Over Java-Based Software
Optimization Mode:		
① Write-back Cache	Enabled	Disabled
① Optimization for Random/Sequential	Sequential (cannot be changed after the creation of a logical drive)	Sequential or Random Refer to “Optimization for Random or Sequential I/O” on page 8-3 for a description.
SCSI Parameters:		
① Maximum Tag Count	32	1-128 or Disabled
① Maximum Queued I/O Count	256	Auto, or 32 to 1024
② LUN's per SCSI ID	8	Up to 32
① Periodic Drive Check Time	Disabled	1/2 to 30 seconds
① SAF-TE/SES Check Time	5	Disabled to 60 seconds
① Periodic Drive Swap Check Time	Disabled	5 to 60 seconds
① Number of Host-LUN Connection	4	1 to 64
① Tag per Host-LUN Connection	32	1 to 256

Legend:

- ① Parameters that are configured at the initial array configuration.
- ② Parameters that can be changed late.
- ③ Non-critical.

B.2 Keeping Specific Parameter Defaults

TABLE B-2 Parameter Defaults to Keep

User-Defined Parameter	Default to Keep (unless directed by Technical Support)
Fault Management:	
S.M.A.R.T.	Disabled on drives
SDRAM ECC	Enabled
SCSI Parameters:	
Data Transfer Rate	80 MHz
Wide Transfer	Enabled
Parity Check	Enabled
Spin-Up Parameters:	
Motor Spin-Up	Disabled
Reset at Power-UP	Enabled
Initial Disk Access Delay	15 (none to 75 seconds)

Event Messages

This appendix lists the following event messages:

- “Controller Events” on page C-2
- “SCSI Drive Events” on page C-3
- “Logical Drive Events” on page C-5
- “General Target Alerts” on page C-6

Events are categorized as follows:

Category	Description/Response
Alert	Errors that need to attend to immediately; may require reconnecting cables, replacing a component, or rebuilding a drive.
Warning	Errors which may indicate a temporary condition, a possible component problem or the need to adjust controller parameters. Press ESC key to clear message.
Notification	Informational message sent from controller firmware; press ESC key to clear message.

C.1 Controller Events

The controller records all array events during power on; it records up to 1,000 events.



Caution – To power off or to reset the controller will cause an automatic deletion of all the recorded event logs.

Alerts

- [0104] Controller ALERT: DRAM Parity Error Detected
- [0105] Controller <primary/secondary> SDRAM ECC <multi-bits/single-bit> Error Detected
- [0110] CHL:_ FATAL ERROR (␣)
- [0111] Controller ALERT: Redundant Controller Failure Detected
- [0114] Controller ALERT: Power Supply Unstable or NVRAM Failed

Warnings

- [0107] Memory Not Sufficient to Fully Support Current Config.

Notifications

- [0111] Controller NOTICE: Redundant Controller Firmware Updated
- [0181] Controller Initialization Completed
- [0187] Memory is Now Sufficient to Fully Support Current Config.
- [0189] NVRAM Factory Defaults Restored
- [0189] NVRAM Restore from Disk is Completed
- [0189] NVRAM Restore from File is Completed

C.2 SCSI Drive Events

Warnings

- [1101] CHL:_ ID:_ SCSI Target ALERT: Unexpected Select Timeout
- [1102] CHL:_ ID:_ SCSI Target ALERT: Gross Phase/Signal Error Detected
- [1103] CHL:_ ID:_ SCSI Target ALERT: Unexpected Disconnect Encountered
- [1104] CHL:_ ID:_ SCSI Drive ALERT: Negotiation Error Detected
- [1105] CHL:_ ID:_ SCSI Target ALERT: Timeout Waiting for I/O to Complete
- [1106] CHL:_ ID:_ SCSI Target ALERT: SCSI Parity/CRC Error Detected
- [1107] CHL:_ ID:_ SCSI Drive ALERT: Data Overrun/Underrun Detected
- [1108] CHL:_ ID:_ SCSI Target ALERT: Invalid Status/Sense Data Received ()
- [110f] CHL:_ LIP() Detected
- [110f] CHL:_ SCSI Drive Channel Notification: SCSI Bus Reset Issued
- [110f] CHL:_ SCSI Drive Channel ALERT: SCSI Bus Reset Issued
- [1111] CHL:_ ID:_ SCSI Target ALERT: Unexpected Drive Not Ready
- [1112] CHL:_ ID:_ SCSI Drive ALERT: Drive HW Error ()\
- [1113] CHL:_ ID:_ SCSI Drive ALERT: Bad Block Encountered - _ ()
- [1114] CHL:_ ID:_ SCSI Target ALERT: Unit Attention Received
- [1115] CHL:_ ID:_ SCSI Drive ALERT: Unexpected Sense Received ()
- [1116] CHL:_ ID:_ SCSI Drive ALERT: Block Reassignment Failed - _ ()
- [1117] CHL:_ ID:_ SCSI Drive ALERT: Block Successfully Reassigned - _ ()
- [1118] CHL:_ ID:_ SCSI Drive ALERT: Aborted Command ()
- [1142] SMART-CH:_ ID:_ Predictable Failure Detected (TEST)
- [1142] SMART-CH:_ ID:_ Predictable Failure Detected
- [1142] SMART-CH:_ ID:_ Predictable Failure Detected-Starting Clone
- [1142] SMART-CH:_ ID:_ Predictable Failure Detected-Clone Failed\

Notifications

[11c1] CHL:_ ID:_ SCSI Drive NOTICE: Scan scsi drive Successful

C.3 SCSI Channel Events

Alerts

[113f] CHL:_ ALERT: Redundant Loop Connection Error Detected on ID:_

[113f] CHL:_ SCSI Drive Channel ALERT: SCSI Channel Failure

[113f] CHL:_ ALERT: Fibre Channel Loop Failure Detected

[113f] CHL:_ ALERT: Redundant Loop for Chl:_ Failure Detected

[113f] CHL:_ ALERT: Redundant Path for Chl:_ ID:_ Expected but Not Found

[113f] CHL:_ ID:_ ALERT: Redundant Path for Chl:_ ID:_ Failure Detected

Notifications

[113f] CHL:_ NOTICE: Fibre Channel Loop Connection Restored

[113f] CHL:_ ID:_ NOTICE: Redundant Path for Chl:_ ID:_ Restored

C.4 Logical Drive Events

Alerts

- [2101] LG: <NA/Logical Drive Index> Logical Drive ALERT: CHL:_ ID:_ SCSI Drive Failure
- [2103] LG:_ Logical Drive ALERT: Rebuild Failed
- [2106] LG:_ Logical Drive ALERT: Add SCSI Drive Operation Failed
- [2102] LG:_ Logical Drive ALERT: Initialization Failed
- [2104] LG:_ Logical Drive ALERT: Parity Regeneration Failed
- [2105] LG:_ Logical Drive ALERT: Expansion Failed
- [2111] LG:_ Logical Drive ALERT: CHL:_ ID:_ Clone Failed

Notifications

- [2181] LG:_ Logical Drive NOTICE: Starting Initialization
- [2182] Initialization of Logical Drive _ Completed
- [2183] LG:_ Logical Drive NOTICE: Starting Rebuild
- [2184] Rebuild of Logical Drive _ Completed
- [2185] LG:_ Logical Drive NOTICE: Starting Parity Regeneration
- [2186] Parity Regeneration of Logical Drive _ Completed
- [2187] LG:_ Logical Drive NOTICE: Starting Expansion
- [2188] Expansion of Logical Drive _ Completed
- [2189] LG:_ Logical Drive NOTICE: Starting Add SCSI Drive Operation
- [218a] Add SCSI Drive to Logical Drive _ Completed
- [218b] LG:_ Logical Drive NOTICE: Add SCSI Drive Operation Paused
- [218c] LG:_ Logical Drive NOTICE: Continue Add SCSI Drive Operation
- [21a1] LG:_ Logical Drive NOTICE: CHL:_ ID:_ Starting Clone”
- [21a2] LG:_ Logical Drive NOTICE: CHL:_ ID:_ Clone Completed”

C.5 General Target Alerts

SAF-TE Device

[3f21] SAF-TE Device () ALERT: Power Supply Failure Detected ()

[3f22] SAF-TE Device () ALERT: Cooling Fan Not Installed ()

[3f22] SAF-TE Device () ALERT: Cooling Fan Failure Detected ()

[3f23] SAF-TE Device () ALERT: Elevated Temperature Alert ()

[3f24] SAF-TE Device () ALERT: UPS Power Failure Detected ()

Controller on-board

[3f23] Peripheral Device ALERT: CPU Temperature <high/low threshold>
Temperature Detected (._C)

[3f23] Peripheral Device ALERT: Board1 Temperature <high/low threshold>
Temperature Detected (._C)

[3f23] Peripheral Device ALERT: Board2 Temperature <high/low threshold>
Temperature Detected (._C)

[3f22] Peripheral Device ALERT: Controller FAN _ Not Present or Failure
Detected

[3f22] Peripheral Device ALERT: Controller FAN _ <high/low threshold> Speed
Detected (_RPM)

[3f21] Peripheral Device ALERT: +3.3V <upper/lower threshold> Voltage
Detected ()

[3f21] Peripheral Device ALERT: +5V <upper/lower threshold> Voltage Detected
()

[3f21] Peripheral Device ALERT: +12V <upper/lower threshold> Voltage Detected
()

I²C Device

[3f23] Peripheral Device ALERT: Temperature Sensor _ Failure Detected

[3f23] Peripheral Device ALERT: Temperature Sensor _ Not Present

[3f23] Peripheral Device ALERT: <high/low threshold> Temperature _ Detected (_(F/C))

[3f22] Peripheral Device ALERT: FAN _ Failure Detected

[3f22] Peripheral Device ALERT: FAN _ Not Present

[3f22] Peripheral Device ALERT: <high/low threshold> FAN _ Speed Detected (_RPM)

[3f21] Peripheral Device ALERT: Power Supply _ Failure Detected

[3f21] Peripheral Device ALERT: Power Supply _ Not Present

[3f21] Peripheral Device ALERT: <high/low threshold> Power Supply _ Voltage Detected (_)

[3f24] Peripheral Device ALERT: UPS _ AC Power Failure Detected

[3f24] Peripheral Device ALERT: UPS _ Battery Failure Detected

SES Devices

[3f21] SES (C_ I_) Power Supply _: <Vendor descriptor strings/Device Not Supported>!

[3f21] SES (C_ I_) Power Supply _: <Vendor descriptor strings/Device Not Installed>!

[3f21] SES (C_ I_) Power Supply _: <Vendor descriptor strings/Device Unknown Status>!

[3f21] SES (C_ I_) Power Supply _: <Vendor descriptor strings/Device Not Available>!

[3f22] SES (C_ I_) Cooling element _: <Vendor descriptor strings/Device Not Supported>!

[3f22] SES (C_ I_) Cooling element _: <Vendor descriptor strings/Device Not installed>!

[3f22] SES (C_ I_) Cooling element _: <Vendor descriptor strings/Device Unknown Status>!

[3f22] SES (C_ I_) Cooling element _: <Vendor descriptor strings/Device Not Available>!

[3f23] SES (C_ I_) Temperature Sensor _: <Vendor descriptor strings/Device Not Supported>!

[3f23] SES (C_ I_) Temperature Sensor _: <Vendor descriptor strings/Device Not installed>!

[3f23] SES (C_ I_) Temperature Sensor _: <Vendor descriptor strings/Device Unknown Status>!

[3f23] SES (C_ I_) Temperature Sensor _: <Vendor descriptor strings/Device Not Available>!

[3f24] SES (C_ I_) UPS _: <Vendor descriptor strings/Device Not Supported>!

[3f24] SES (C_ I_) UPS _: <Vendor descriptor strings/Device Not installed>!

[3f24] SES (C_ I_) UPS _: <Vendor descriptor strings/Device Unknown Status>!

[3f24] SES (C_ I_) UPS _: <Vendor descriptor strings/Device Not Available>!

[3f21] SES (C_ I_) Voltage sensor _: <Vendor descriptor strings/Device Not Supported>!

[3f21] SES (C_ I_) Voltage sensor _: <Vendor descriptor strings/Device Not installed>!

[3f21] SES (C_ I_) Voltage sensor _: <Vendor descriptor strings/Device Unknown Status>!

[3f21] SES (C_ I_) Voltage sensor _: <Vendor descriptor strings/Device Not Available>!

[3f21] SES (C_ I_) Current sensor _: <Vendor descriptor strings/Device Not Supported>!

[3f21] SES (C_ I_) Current sensor _: <Vendor descriptor strings/Device Not installed>!

[3f21] SES (C_ I_) Current sensor _: <Vendor descriptor strings/Device Unknown Status>!

[3f21] SES (C_ I_) Current sensor _: <Vendor descriptor strings/Device Not Available>!

General Peripheral Device

[3f21] Peripheral Device ALERT: Power Supply Failure Detected

[3f22] Cooling Fan Not Installed

[3f22] Cooling Fan Failure Detected

[3f24] Elevated Temperature Alert

[3f24] UPS Power Failure Detected

Glossary

The glossary lists acronyms and defines RAID terms found through the documentation. It also includes definitions of the operational states for disk drives, logical drives, and redundant controllers.

Acronyms

ANSI	American National Standards Institute
CH	Channel
CISPR	International Special Committee on Radio Interference
EMU	Event Monitoring Unit
FC-AL	Fibre Channel-Arbitrated Loop. FC-AL is implemented as either a loop or a Fabric. A loop can contain up to 126 nodes, accessible through only one or two servers.
FRU	Field-Replaceable Unit
GB	Gigabyte. 1,000,000,000 (one billion) bytes.
GBIC	Gigabit Interface Converter. A hot-swappable input/output device that plugs into a Gigabit Ethernet port or Fibre Channel.
HBA	Host Bus Adapter
ID	Identifier number
IEC	International Electrotechnical Commission
JBOD/Expansion Unit	Just a Bunch of Disks (a Sun StorEdge unit with drives and no controllers)
LAN	Local Area Network
LD	Logical drive

- LUN** Logical unit number. The major and minor device numbers make up the logical unit numbering sequence for a particular device connected to a computer.
- LVD** A low-noise, low-power, and low-amplitude signaling technology that enables data communication between a supported server and storage devices. LVD signaling uses two wires to drive one signal over copper wire and requires a cable that is no longer than 25 meters (82 ft.).
- MB** Megabyte. 1,000,000 bytes or characters of data.
- NVRAM** Non-Volatile Random Access Memory. A memory unit equipped with a battery so that the data stays intact even after the main power had been switched off.
- PID** Primary controller identifier number
- RAID** Redundant Array of Independent Disks. A configuration in which multiple drives are combined into a single virtual drive to improve performance and reliability.
- SAN** Storage Area Networking. A high-speed, open-standard scalable network of storage devices and servers providing accelerated data access.
- SCSI** Small Computer Systems Interface. An industry standard for connecting disk and tape devices to a workstation.
- SES** SCSI Enclosure Services driver. An interface to SCSI Enclosure Services devices. These devices sense and monitor the physical conditions within an enclosure, as well as enable access to the status reporting and configuration features of the enclosure (such as indicator LEDs on the enclosure).
- SID** Primary controller identifier number
- S.M.A.R.T.** Self Monitoring Analysis and Reporting Technology. The industry standard reliability prediction indicator for both the IDE/ATA and SCSI hard disk drives. Hard disk drives with SMART offer early warning of some hard disk failures so critical data can be protected.
- SMTP** Simple Mail Transfer Protocol. A protocol for sending e-mail messages between servers and from mail clients to mail servers. The messages can then be retrieved with an e-mail client using either POP or IMAP.
- SNMP** Simple Network Management Protocol. A set of protocols for managing complex networks. SNMP works by sending messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.
- WWN** worldwide name. A number used to identify array logical drives in both the array software and in Solaris operating environment.

Terminology

active-active controllers	A pair of components, such as storage controllers in a failure-tolerant RAID array that share a task or set of tasks when both are functioning normally. When one component of the pair fails, the other takes the entire load. Dual active controllers (also called dual-active controllers) are connected to the same set of devices and provide a combination of higher I/O performance and greater failure tolerance than a single controller.
automatic rebuild	A process where data is automatically reconstructed after a drive failure and written to a standby (spare) drive. An automatic rebuild will also occur when a new drive is installed manually in place of a failed drive. If the rebuild process is interrupted by a reset, use the Rebuild command on the Array Administration menu to restart the rebuilding process.
background rate	The background rate is the percentage of available array controller CPU time assigned to array administration activities, such as rebuilding failed drives, checking parity, and initialization. If the background rate is set to 100%, the array administration activities have a higher priority than any other array activity. At 0%, the array administration activity is done only if there is no other activity on the array controller.
caching	Allows data to be stored in a pre-designated area of a disk or RAM (random access memory). Caching is used to speed up the operation of RAID arrays, disk drives, computers and servers, or other peripheral devices.
channel	Any path used for the transfer of data and control information between storage devices and a storage controller or I/O adapter. Also refers to one SCSI bus on a disk array controller. Each disk array controller provides at least one channel.
fabric	Fibre Channel network built around one or more switches.
fabric switch	A fabric switch functions as a routing engine, which actively directs data transfer from source to destination and arbitrates every connection. Bandwidth per node via a Fabric switch remains constant when more nodes are added, and a node on a switch port uses an up-to-100MBps data path to send or receive data.
failover	A mode of operation for failure-tolerant arrays in which a component has failed and its function has been assumed by a redundant component.
fault tolerance	The capacity to cope with internal hardware problems without interrupting the array's data availability, often by using backup systems brought online when a failure is detected. Many arrays provide fault tolerance by using RAID architecture to give protection against loss of data when a single disk drive fails. Using RAID 1 (mirroring), RAID 3 or 5 (striping with parity), or RAID 6 or 1+0 (mirroring and striping) techniques, the array controller can reconstruct data from a failed drive and write it to a standby or replacement drive.

fault tolerant logical drive	A logical drive that provides protection of data in the event of a single drive failure by employing RAID 1, 3, 5, or 6 (also known as RAID 1+0).
fibre channel	A cost-effective gigabit communications link deployed across a wide range of hardware.
fibre channel HBAs	Fibre channel adapters of a host computer, server, or workstation.
fibre hubs	An Arbitrated Loop Hub is a wiring concentrator. "Arbitrated" means that all nodes communicating over this Fibre loop are sharing a 100Mbps segment. Whenever more devices are added to a single segment, the bandwidth available to each node is further divided. A loop configuration allows different devices in the loop to be configured in a token ring style. With a Fibre Hub, a Fibre loop can be re-arranged in a star-like configuration for the Hub itself contains port bypass circuitry that forms an internal loop inside. Bypass circuits can automatically reconfigure the loop once a device is removed or added without disrupting the physical connection to other devices.
groups	<p>A group is a new data object that allows multiple servers to be contained under a single category, are similar in concept to domains, and allow you to organize servers within Configuration Service. Instead of a linear tree of all the managed servers, the Configuration Service operator can organize the servers into like sets or groups. In the case where many servers are present, groups allow more icons to appear at the same time in the main Configuration Service window without scrolling.</p> <p>Groups are not required. You can configure Configuration Service for no groups and fifteen servers, for example, or for one group with ten servers underneath, with an additional five at the top level. Configuration Service allows any combination.</p> <p>The number of groups permitted and the number of servers within a group is limited only by available array memory. If a server is a member of a group and a user deletes that group from the group list box, Configuration Service reassigns the servers in that group to the no group category. Configuration Service automatically remaps the tree in the main window.</p>
hot spare	A drive in a RAID 1 or RAID 5 configuration that contains no data and acts as a standby in case another drive fails.
hot-swappable	The ability of a field-replaceable unit (FRU) to be removed and replaced while the array remains powered on and operational.
initialization	The process of writing a specific pattern to all data blocks on all drives in a logical drive. This process overwrites and destroys existing data on the disks and the logical drive. Initialization is required to make the entire logical drive consistent at the onset. Initialization ensures that any parity checks performed in the future will execute correctly.

logical drive	A section of disk storage space, also referred to as a LUN, that is presented to the host operating environment as a single physical drive. A logical drive may be located on one or more physical drives. Each array controller can manage one to eight logical drives
LUN mapping	The ability to change the virtual LUN as presented to the server from storage. This enables such benefits as the ability of a server to boot from the SAN without requiring of a local disk drive. Each server requires LUN 0 to boot.
LUN masking	The characteristic that enables an administrator to dynamically map an HBA to a specified LUN. This provides an individual server or multiple servers access to an individual drive or to multiple drives, and prohibits unwanted server access to the same drives.
mirroring (RAID 1)	<p>Data written to one disk drive is simultaneously written to another disk drive. If one disk fails, the other disk can be used to run the array and reconstruct the failed disk. The primary advantage of disk mirroring is 100% data redundancy. Since the disk is mirrored, it does not matter if one of the disks fails. Both disks contain the same data at all times and either can act as the operational disk.</p> <p>Disk mirroring provides 100% redundancy, but is expensive because each drive in the array is duplicated.</p>
N port	A Fibre Channel port in a point-to-point or Fabric connection.
out-of-band	Refers to the connections and devices that are not in the data path.
parity check	A process whereby the integrity of the redundant data on fault tolerant arrays (RAID 1, 3, 5, and 6 or 1+0) is checked. The parity checking procedure on a logical drive recalculates the parity of data stripes in each of the logical drive's RAID stripe sets and compares it with the stored parity. If a discrepancy is found, an error will be reported and the new correct parity will be substituted for the stored parity.
partner group	A pair of interconnected controller units. Expansion units interconnected to the pair of controller units can also be part of the partner group.
physical array	<p>A physical array is a group of physical drives in Configuration Service that participate in one or more logical drives. If a group of physical drives is configured into a logical drive without using the entire capacity of the physical drives, Configuration Service will require that the same physical array of drives be used together in any remaining logical drives that are created with the remaining capacity.</p> <p>In the New Configuration window, the Add Disk button changes to Add Array when the remaining drive capacity is on multiple physical drives that have already been used to create a logical drive. Now that the physical drives have been sliced, they have to be selected as an array. They cannot be selected separately.</p>

RAID Redundant Array of Independent Disks. An arrangement of two or more disk drives configured to work together to provide more disk storage space, better performance, and/or redundant backup of data. Various combinations of these features are described by defined RAID levels. Sun StorEdge arrays can support RAID 0, 1, 3, 5, and 6 (also called 1+0).

For more detailed RAID definitions, refer to “RAID Levels” on page 1-3.

read policy Descriptions of the read policies are as follows:

When you indicate a No Cache read policy, the controller will not store data in the read cache.

Normal specifies that the current controller does not use Read ahead caching for the current drive.

Set the read policy to Read Ahead and the controller will read multiple adjacent blocks of data automatically. This is most effective for applications with sequential reads.

rebuild Rebuilding a disk is the process of reconstructing the data on a disk before it failed. Rebuilding can be done only in arrays with data redundancy, such as RAID levels 1, 3, 5, 6 or 1+0.

See “automatic rebuild” on page -3 for detailed information on rebuilding. For the Rebuild rate, refer to “background rate” on page -3.

spanning Disk spanning makes use of the firmware’s striping capability to stripe data across two otherwise independent RAID logical drives. The two spanned logical drives are presented to the operating environment as one logical drive. The advantages of spanning are:

-- Supports two simultaneous drive failures in the combined fault tolerant logical drives (provided there is one drive failure from each logical drive).

-- Improves performance because the number of spindles is increased.

The disadvantage of spanning is that the RAID overhead for redundant RAID levels is increased since each logical drive handles fault tolerance separately.

standby drive A drive that is marked as a spare to support automatic data rebuilding after a physical drive associated with a logical drive fails. For a standby drive to take the place of another drive, it must be at least equal in size to the failed drive and all of the logical drives dependent on the failed disk must be redundant—RAID 1, 3, 5, and 6 or 1+0.

state The current operational status of a disk drive, a logical drive, or redundant controllers. The arrays stores the states of drives, logical drives, and redundant controllers in its nonvolatile memory. This information is retained across array power interruptions.

- stripe size** This is the amount of data in kilobytes that is striped across each physical drive in a logical drive. The values are in increments of 8 kilobytes and range from 8 to 64 kilobytes. Generally, large stripe sizes are more effective for arrays with primarily sequential reads.
- To change the stripe size on an existing drive, you need to back up your data, redefine the stripe size, reconfigure the storage, and restore all the data.
- striping** The storing of sequential blocks of incoming data on all the different SCSI drives in a logical drive. For example, if there are three SCSI drives in a logical drive, data will be stored as follows:
- block 1 on SCSI drive 1
 - block 2 on SCSI drive 2
 - block 3 on SCSI drive 3
 - block 4 on SCSI drive 1
 - block 5 on SCSI drive 2, etc.
- This method of writing data increases the disk array throughput because multiple drives are working simultaneously, retrieving and storing. RAID 0, 3, 5, and 6 or 1+ 0 all use striping.
- terminator** A part used to end a SCSI bus. Terminators prevent energy from reflecting back into a cable plant by absorbing the radio frequency signals.
- volume** Also called a logical unit number or LUN, a volume is one or more drives that can be grouped into a unit for data storage.
- write-back cache** A cache-writing strategy in which the array controller receives the data to be written to disk, stores it in the memory buffer, and immediately sends the host operating environment a signal that the write operation is complete, without waiting until the data is actually written to the disk drive. Within a short time, the controller, when not busy, writes the data to the disk drive.
- Write-back caching improves the performance of write operations and the throughput of the controller card. However, because there is a danger of data loss in the case of a power failure, arrays with write-back caching should be equipped with a UPS or battery backup cache. A UPS will provide power long enough to allow any data in the cache memory to be written to the disk drive. With battery backup cache, the battery will provide power to retain the memory data for up to 48 hours.
- write policy** A cache-writing strategy used to control write operations. The write policy options are write-back and write-through cache.

write-through cache A cache-writing strategy in which the array controller writes the data to the disk drive before signaling the host operating environment that the process is complete. Write-through cache has lower write operation and throughput performance than write-back cache, but it is the safer strategy, with minimum risk of data loss on power failure.

Index

SYMBOLS

#FL, 3-5
#LN, 3-5
#SB, 3-5

A

Add Channel SCSI ID command, 7-9
add drive Entry command, 6-11
Add Global spare drive command, 6-5
add Local spare drive command, 6-4
add SCSI drives command, 3-23
adding
 drive entry, 6-11
 global spare drive, 6-5
 host channel SCSI ID, 7-9
 local spare drive, 6-4
 SCSI drives, 3-23
automatic rebuild, 8-6
 definition, Glossary-3

B

background rate
 definition, Glossary-3
BAD drive status, 6-3
battery support, 1-14
baud rate, 2-3
beeper muting, 10-2
bus configurations, 7-7

C

cabling
 single vs dual bus configurations, 7-7
cache parameters, 8-2
cache status, 2-4
Caching Parameters command, 8-4
check time
 periodic auto-detection of failed drive swap, 8-23
Clear drive status command, 6-11
clone failing drive, 6-16
 perpetual clone, 6-18
comments
 sending to manufacturer, Preface-xv
concurrent rebuild, 8-9
configuration
 minimum requirements, 1-16
 saving to disk, 10-6
controller
 defaults, 1-13
 muting beeper, 10-2
 name, 8-12
 naming, 8-12, 8-13
 optimization mode, 1-15
 parameter settings, B-1
 fault management, B-1
 fibre channel, B-1
 optimization mode, B-1
 SCSI, B-1
 spin-up, B-1
parameters
 name, 8-12
 password validation timeout, 8-13

- SDRAM ECC, 8-15
 - unique identifier, 8-15
- resetting, 10-4
- shutting down, 10-5
- voltage and temperature
 - checking, 9-8
- controller assignment, 3-11
- controller failure, 8-5
- controller IDs, 7-9
- controller NAME, 3-12
- Controller Name command, 8-12
- Controller Parameters command, 8-13
- copy and replace
 - logical drive, 3-25
- copy and replace drive command, 3-25
- create
 - logical drive, 3-8
- creating
 - logical volume, 4-5
- cylinder/head/sector mapping, 8-31

D

- defaults
 - controller, 1-13
 - drive-swap check time, 8-23
 - enable UPS status, 9-6
 - host LUN connections, 8-28
 - logical drives, 3-3
 - number of LUNs per host SCSI ID, 8-30
 - password validation, 8-13
 - RAID levels, 3-3
 - SCSI I/O timeout, 8-19
 - SDRAM ECC, 8-15
 - UPS alert notification level, 9-7
- Define Peripheral Device Active Signal
 - command, 9-7
- Delete Channel SCSI ID command, 7-10
- Delete global/local spare drive command, 6-9
- deleting
 - drive entries, 6-11
 - host channel SCSI ID, 7-10
 - logical drive, 3-16
 - partition of a logical drive, 3-17
 - partition of logical drive, 3-17
 - caution, 4-3
 - SCSI drive table
 - slot number, 6-11

- SCSI ID, 7-10
 - spare drive, 6-9
 - global or local, 6-9
- detection
 - idle drive failure, 8-23
- device supports removable media
 - peripheral device parameters, 8-33
- Disconnect support, 7-17
- disconnecting support, 7-17
 - SCSI target/drive channel, 7-17
- disk
 - access delay time, 8-19
 - array parameters
 - advanced config, 8-24
 - mirroring, 1-6
 - striping
 - mirroring, 1-7
 - with dedicated parity disk, 1-8
- documentation
 - how book is organized, Preface-xi
 - using UNIX commands, Preface-xii
- documentation available, Preface-xiv
- documentation comments, Preface-xv
- documentation online, Preface-xv
- drive
 - configuring, 1-16
 - entries
 - adding, 6-11
 - deleting, 6-11
 - failed
 - checking, 8-23
 - spares
 - global, 1-11
 - local, 1-11
- drive channel commands, 7-4
- drive-side parameters
 - advanced config, 8-16
- Drive-side SCSI Parameters command, 8-17, 8-18
- DRV FAILED status, 3-5
- dual bus configurations, 7-7

E

- ECC
 - drives, 1-3
- ECC SDRAM function, 8-15
- error correction code
 - drives, 1-3

- event logs
 - viewing on screen, 10-7
- Execute Drive Testing command, 6-22
- Expand logical drive command, 3-29
- Expand logical volume command, 4-7
- expanding
 - limitations, 3-21
 - logical drive, 3-28
 - copy and replace, 3-25
 - logical volume, 4-4, 4-7
- expansion
 - limitations, 3-21

F

- fabric
 - definition, Glossary-3
- fault prevention, 6-16
- fault tolerance
 - definition, Glossary-3
- fault-tolerance, 8-5
- fault-tolerant logical drive
 - definition, Glossary-4
- FC-AL
 - definition, Glossary-1
- fibre channel
 - definition, Glossary-4
- firmware
 - cache status, 2-3
 - controller failure/rebuild/replacement, 8-5
 - controller voltage and temperature, 9-8
 - cursor bar, 2-3
 - gauge range, 2-3
 - logical drive status, 3-4
 - main menu, 2-5
 - manual rebuild, 8-6
 - quick installation, 2-5
 - SCSI channel
 - status, 7-2
 - SCSI drive status, 6-2
 - transfer rate indicator, 2-3
 - VT-100 screen selection, 2-3
- flash all but selected drive command, 6-14
- Flash All Drives command, 6-13
- flash drive time command, 6-13
- flash selected drive command, 6-14
- flashing
 - all SCSI drives, 6-15

- selected SCSI drives, 6-14

format

- low-level, 6-21

G

- gauge range, 2-4
- GBIC
 - definition, Glossary-1
- global spare drive, 1-11
 - creating, 6-5
 - deleting, 6-9
 - explained, 1-11
- groups
 - definition, Glossary-4

H

- head
 - mapping, 8-31
- host application, 1-15
- host channel commands, 7-4
- host channel SCSI ID
 - adding, 7-9
 - deleting, 7-10
- Host Cylinder/Head/Sector Mapping Configuration command, 8-31
- host IDs (adding), 7-9
- host LUN
 - connection
 - number of tags reserved, 8-29
 - example mappings, 5-6
 - mapping
 - deleting, 5-7
- host LUN connections, 8-28
- host-side
 - parameters, 8-27
- hot spare
 - definition, Glossary-4
- hot-swappable
 - definition, Glossary-4

I

- I/O
 - maximum queued count, 8-29

- random or sequential optimization, 8-3
- SCSI timeout, 8-19
- ID
 - SCSI, deleting, 7-10
- identifying drives, 1-10, 6-13, 8-9
- Identifying SCSI drive command, 6-13
- idle drive failure detection, 8-23
- INCOMPLETE status, 3-5
- initial screen
 - main menu, 2-5
- initialization
 - definition, Glossary-4
- INITING status, 3-4
- international support, Preface-xiv
- INVALID status, 3-4

J

- JBOD, 1-5
 - explained, 1-5
 - single drive control, 1-5

L

- LCD title display controller name
 - not applicable, 8-13
- limitations
 - expanding, 3-21
 - logical volume, 4-3
 - redundant config, 1-13
- local spare assignment, 3-9
- local spare drives
 - deleting, 6-9
 - explained, 1-11
- logical drive
 - 128-LUN limit, 3-3
 - adding a SCSI drive, 3-21
 - assign local spare, 3-9
 - assigning name, 3-18
 - change assignment, 3-11
 - commands, 3-2
 - logical drive Assignments, 3-12
 - view and edit logical drives, 3-4, 3-16
 - copying and replacing, 3-25
 - creating, 3-2, 3-5, 3-8
 - defaults, 3-3
 - definition, Glossary-5

- deleting, 3-16
- expanding, 3-28
- explained, 1-2
- ID, 3-4
- LG number, 3-4
- maximum physical capacity, 3-9
- maximum physical drive capacity, 3-9
- minimum requirements, 1-16
- parity
 - checking, 3-20
- partitions
 - maximum, 1-16
- RAID level, 3-4
- RAID level selected, 3-8
- RAID levels
 - default, 3-3
- rebuilding, 3-19
- size, 3-4
- status, 3-4
- logical drive Assignments command, 3-12
- logical drive name command, 3-18
- logical drive preferences, 3-9
- logical drives, 8-24
- logical volume, 4-2, 4-3, 4-4
 - avoiding failure, 4-3
 - creating, 4-5
 - example
 - multi-level, 4-2
 - expanding, 4-7
 - using RAID expansion, 4-4
 - explained, 1-2, 4-2
 - limitations, 4-3
 - maximum partitions, 1-2, 4-2
 - multi-level RAID, 4-4
 - partitioning, 4-3
 - spare drives, 4-3
- logical volume status, 4-8
- low-level format, 6-21
- LUN
 - described, 5-3
 - explained, 8-27
 - mapping
 - definition, Glossary-5
 - masking
 - definition, Glossary-5
 - per host SCSI ID, 8-30
 - changing, 8-30
- LUN Applicability parameter, 8-33
- LUN defined, 8-27

M

- main menu, 2-5
 - navigating, 2-5
 - quick installation, 2-5
- manual rebuild, 8-6
- mapping
 - cylinder/head/sector, 8-31
- maximum
 - queued I/O count, 8-29
 - tag count, 7-18
 - tag count (tag command queuing), 8-21
 - transfer width, 7-16
- maximum concurrent host LUN connections, 8-28
- maximum drive capacity, 3-9
- Maximum Queued I/O Count command, 8-29
- maximum sync. xfer Clock command, 7-16
- Maximum Tag Count (tag command queuing)
 - command, 8-21
- maximum Tag count command, 7-18
- maximum xfer Width command, 7-16
- mirroring, 1-6
- mirroring (RAID 1)
 - definition, Glossary-5
- MISSING drive status, 6-3
- motor spin-up, 8-17
 - SCSI, 8-17

N

- N port
 - definition, Glossary-5
- NAME (controller), 3-12
- naming
 - controller, 8-12, 8-13
- narrow transfer, 7-14
- NEW DRV drive status, 6-2
- nexus (SCSI), 8-28
- NRAID, 1-4
 - disk spanning, 1-4
- Number of Tags Reserved for each Host-LUN Connection Command, 8-29
- NVRAM
 - saving to disk, 10-6

O

- Optimization for Random I/O command, 8-3

- Optimization for Sequential I/O command, 8-3
- optimization mode, 1-15, 8-2
 - limitations, 8-2
 - random or sequential, 8-3
- out-of-band
 - definition, Glossary-5

P

- parameters
 - cache, 8-2
 - controller, 8-12
 - drive-side, 8-16
 - drive-side SCSI, 8-17
 - host-side
 - advanced config, 8-27
 - peripheral device, 8-33
 - physical drives, 6-1
- parity
 - logical drive
 - checking, 3-20
- parity check, 7-17
 - definition, Glossary-5
 - SCSI target/drive channel, 7-17
- partitioning
 - logical volume, 4-3
- partitions
 - deleting, 3-17
 - logical drive
 - deleting, 3-17
 - logical volume
 - maximum, 1-2
 - maximum, 1-16, 4-2
- partner group
 - definition, Glossary-5
- password
 - changing, 10-3
 - disabling, 10-4
 - setting a new, 10-3
 - setting new, 10-3
 - validation timeout, 8-13
- Password Validation Timeout command, 8-14
- periodic drive check time, 8-22
- periodic drive swap auto check, 8-23
- peripheral device
 - parameters, 8-33
 - setting, 9-5
 - viewing status, 9-2

- Peripheral Device Qualifier parameter, 8-33
- perpetual clone
 - clone failing drive, 6-18
- phone support, Preface-xiv
- physical drives listed, 3-7, 6-7
- physical array
 - definition, Glossary-5
- physical drive capacity setting, 3-9
- physical drives
 - parameters, 6-1
- primary/secondary controller ID, 7-9

Q

- quick installation, 2-5
 - warning, 2-5

R

- RAID
 - advantages, 1-3
 - controller, 8-24
 - definition, 1-3
 - expansion with logical volume, 4-4
 - glossary, Glossary-1
 - level (0+1), disk striping with mirroring, 1-7
 - level 0, disk striping, 1-5
 - level 1, disk mirroring, 1-6
 - level 3, disk striping with dedicated parity disk, 1-8
 - level 5, striping with interspersed parity, 1-9
 - NRAID, disk spanning, 1-4
 - planning considerations, 1-15
- RAID (0+1), 4-4
 - concurrent rebuild in, 8-9
- RAID (3+0), 4-4
- RAID (5+0), 4-4
- RAID (5+1), 4-4
- RAID (5+5), 4-5
- RAID 0, 1-5
- RAID 1, 1-6
- RAID 3, 1-8
- RAID 5, 1-9
- RAID levels
 - controller optimization mode, 1-15
 - described, 1-3
 - explained, Glossary-6

- planning, 1-15
 - selected, 3-8
- read ahead policy
 - definition, Glossary-6
- read policy
 - definition, Glossary-6
- Read/Write Test, 6-22
- rebuild
 - automatic, 8-6
 - concurrent in RAID (0+1), 8-9
 - definition, Glossary-6
 - manual, 8-6
- Rebuild logical drive command, 3-19
- Rebuild Priority command, 8-25
- rebuilding, 8-24
 - logical drive, 3-19
- redundant controller
 - explained, 8-5
- Regenerate Parity command, 3-20
- remote file, 2-3
- replace after clone
 - clone failing drive, 6-16
- Replace After Clone command, 6-17
- reset controller
 - controller reset, 3-12
- RS-232
 - connecting to, 2-2

S

- SAF-TE
 - enclosure monitoring, 8-22
- saving
 - configuration to disk, 10-6
- SB-MISS drive status, 6-3
- Scan SCSI drive command, 6-8
- scanning
 - new SCSI drive, 6-8
- scanning a new SCSI drive, 6-8
- SCSI
 - channel
 - explained, 8-27
 - setting termination, 7-12
 - status, 7-2
 - terminator, 7-12
 - drive-side parameters, 8-17
 - I/O timeout, 8-19
 - ID

- deleting, 7-10
 - explained, 8-27
- motor spin-up, 8-17
- reset at power-up, 8-18
- target drive channel
 - viewing and editing, 7-15
- transfer clock speed
 - setting, 7-13
- SCSI channel commands, 7-4
- SCSI channel defined, 8-27
- SCSI drive
 - adding to logical drive, 3-21
 - disk capacity, 6-6
 - fibre port name, 6-6
 - global or local spare, 6-2
 - identifying a drive
 - diagnosing system errors, 6-13
 - low-level format, 6-21
 - read/write test, 6-22
 - redundant loop ID, 6-6
 - revision number, 6-6
 - scanning new, 6-8
 - serial number, 6-6
 - slot number
 - deleting, 6-11
 - STANDBY mode, 6-2
 - table
 - adding entries, 6-11
 - clear drive status, 6-11
 - deleting drive entries, 6-11
 - removing empty drive entry, 6-12
 - setting slot numbers, 6-10
 - USED DRV mode, 6-2
 - utilities, 6-21
 - vendor ID, 6-3
 - viewing information, 6-6
- scsi Drive Low-Level Format command, 6-22
- SCSI host IDs (creating), 7-9
- SCSI ID defined, 8-27
- SCSI Motor Spin-Up command, 8-17
- SCSI Reset at Power-Up command, 8-18
- SCSI target/drive channel
 - maximum tag count, 7-18
 - parity check, 7-17
 - slot number, 7-16
- scsi Terminator command, 7-12
- SDRAM ECC, 8-15
 - default, 8-15
- sector
 - mapping, 8-31
- sending documentation comments, Preface-xv
- serial port
 - connection and setup, 2-1
- serial port parameters, 2-2
- SES
 - definition, Glossary-2
- Set Peripheral Device Entry
 - command, 9-5
- set slot Number command, 6-12
- setting a new password, 10-3
- single bus configurations, 7-7
- slot number
 - assigning to empty slot, 6-10
 - deleting, 6-11
 - SCSI target/drive, 7-16
 - setting, 6-10
- Solaris
 - reset baud rate, 2-3
- spanning, 1-4
 - definition, Glossary-6
- spare (local for logical drive), 3-9
- spare drives, 1-10, 1-16, 8-9
 - assigning, 6-4
 - deleting, 6-9
 - global
 - creating, 6-5
 - local
 - explained, 6-4
 - logical volume, 4-3
- standby drive
 - definition, Glossary-6
- STAND-BY drive status, 6-2
- status
 - logical drive, 3-4
 - peripheral device, 9-2
 - temperature, 9-8
 - UPS, 9-6
 - voltage, 9-8
- stripe size
 - definition, Glossary-7
- striping, 1-5
 - definition, Glossary-7
 - with interspersed parity, 1-9
 - with mirroring, 1-7
 - with parity disk, 1-8
- sync transfer clock, 7-13
- system functions, 10-2
 - controller

- resetting, 10-4
- shutting down, 10-5
- muting beeper, 10-2
- password
 - changing, 10-3
 - disabling, 10-4
- saving NVRAM to disk, 10-6
- setting new password, 10-3

T

- tag count
 - maximum, 7-18, 8-21
- technical support, Preface-xiv
- temperature
 - check status, 9-8
- terminator
 - SCSI channel, 7-12
- tip command, 2-3
- transfer clock
 - maximum synchronous, 7-16
- transfer clock speed
 - options, 7-14, 7-15
 - setting, 7-13
- transfer rate indicator, 2-4
- transfer speed
 - setting, 7-13
- transfer width
 - maximum, 7-16
 - setting, 7-14
- troubleshooting support, Preface-xiv

U

- uninterruptible power supply
 - enable status, 9-6
- UPS
 - enable status, 9-6
 - fail signal, 9-7
 - status, 9-6
- UPS power fail signal, 9-7
- UPS status, 9-6
- USED DRV drive status, 6-2

V

- Verification on LD Initialization Writes Disabled
 - command, 8-26
- verification on writes, 8-25
- view
 - connected drives, 3-7, 6-7
- view and edit Configuration parameters
 - command, 8-3, 8-4
- view and edit Event logs command, 10-7
- view and edit Host luns command, 5-7
- view and edit logical drives command, 3-4, 3-16
- view and edit logical Volumes command, 4-5
- view and edit scsi drives command, 6-1, 6-4, 6-5
- view peripheral device status command, 9-2, 9-8
- voltage
 - check status, 9-8
- volume
 - definition, Glossary-7
- VT-100
 - connection and setup, 2-1

W

- warnings
 - quick installation, 2-5
- world wide name
 - definition, Glossary-2
- write error
 - avoiding, 8-25
- write policy
 - definition, Glossary-7
- write-back cache
 - definition, Glossary-7
 - disabling, 8-4
 - enabling, 8-4
- write-through
 - disabling, 8-4
 - enabling, 8-4
- write-through cache
 - definition, Glossary-8
- WWN
 - definition, Glossary-2