



Sun StorEdge™ 3000 Family RAID Firmware 3.25 and 3.27 User's Guide

Sun StorEdge 3310 SCSI Array

Sun StorEdge 3510 FC Array

Sun StorEdge 3511 FC Array with SATA

Sun Microsystems, Inc.
www.sun.com

Part No. 817-3711-11
May 2004, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright © 2002–2004 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, California 92009, USA. All rights reserved.

Sun Microsystems, Inc. and Dot Hill Systems Corporation might have intellectual property rights relating to technology embodied in this product or document. In particular, and without limitation, these intellectual property rights might 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 might 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 might 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, Sun StorEdge, AnswerBook2, docs.sun.com, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

Oracle is a registered trademark of Oracle Corporation.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

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 © 2002–2004 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, Californie 92009, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. et Dot Hill Systems Corporation peuvent avoir les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. 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 États-Unis et dans 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 États-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Oracle est une marque déposée registre de Oracle Corporation

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

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES 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.



Please
Recycle



Adobe PostScript

Contents

1. Basic RAID Concepts	1
Before You Begin	2
Determining Which Version of the RAID Firmware Applies to Your Array	2
Fibre Channel and SCSI Firmware Illustrations in This Guide	3
RAID Terminology Overview	3
Logical Drives	3
Spare Drives	4
Logical Volumes	4
SCSI Channels	5
RAID Levels	7
RAID 0	9
RAID 1	10
RAID 1+0	11
RAID 3	12
RAID 5	13
Advanced RAID Levels	14
Local and Global Spare Drives	15
Local Spare Drives	15
Global Spare Drives	16

Using Both Local and Global Spare Drives	17
Controller Defaults and Limitations	18
Planning for Reliability, Availability, and Serviceability	18
Dual-Controller Considerations	18
Single-Controller Considerations	20
Battery Operation	20
Battery Status	20
Write-Back Versus Write-Through Cache Options	22
RAID Planning Considerations	22
Basic Configuration Overview	23
▼ To Create a Basic Array Configuration	23
2. Accessing the Controller Firmware	25
Setting Up the Serial Port Connection	25
▼ To Set Up the Serial Port Connection	26
Accessing the Firmware Application From a Solaris Host	27
▼ To Use the <code>tip</code> Command	27
Viewing the Initial Firmware Windows	28
Navigating Firmware Menus	30
Quick Installation (Reserved)	31
Upgrading Firmware	31
3. Viewing and Editing Logical Drives	33
Introducing the Logical Drive Commands	33
Default Logical Drives and RAID Levels	34
Viewing the Logical Drive Status Table	35
Creating Logical Drives	37
▼ To Create Logical Drives	39
Changing a Logical Drive Controller Assignment	43

▼ To Change a Logical Drive Controller Assignment	44
Partitioning a Logical Drive	44
▼ To Partition a Logical Drive	45
Using the Auto Configure Menu Option from the format Menu (Solaris Operating Environment Only)	48
▼ To Change the Label of a LUN	48
Deleting a Logical Drive	51
▼ To Delete a Logical Drive	52
Deleting a Logical Drive Partition	53
▼ To Delete a Logical Drive Partition	53
Assigning a Logical Drive Name	54
▼ To Assign a Logical Drive Name	54
Rebuilding a Logical Drive	55
▼ To Rebuild a Logical Drive	55
Performing a Logical Drive Parity Check	56
▼ To Regenerate the Parity of a Logical Drive	56
Overwriting Inconsistent Parity	57
▼ To Enable or Disable the “Overwrite Inconsistent Parity” Menu Option	57
Adding a SCSI Drive to a Logical Drive	58
▼ To Add a SCSI Drive to a Logical Drive	59
Copying and Replacing Drives With Drives of Larger Capacity	62
▼ To Copy and Replace a Drive	62
Expanding a Logical Drive	64
▼ To Expand a Logical Drive	65
4. Viewing and Editing Logical Volumes	67
Avoiding the Use of Logical Volumes	67
Understanding Logical Volumes (Multilevel RAID)	68
Logical Volume Limitations	68

Partitioning	69
RAID Expansion	69
Multilevel RAID Arrays	70
Spare Drives	70
Creating a Logical Volume	70
▼ To Create a Logical Volume	71
Expanding a Logical Volume	72
▼ To Expand a Logical Volume	72
Viewing the Logical Volume Status Table	73
5. Viewing and Editing Host LUNs	75
Mapping Logical Drive Partitions to Host LUNs	75
Planning for 128 LUNs on SCSI Arrays	78
▼ To Set Up 128 LUNs	78
Planning for 1024 LUNs on FC Arrays	79
▼ To Plan for 1024 LUNs on FC Arrays	79
Planning for 64 LUNs in Redundant FC Point-to-Point Configurations	80
Mapping a Partition to a LUN	81
▼ To Map a Logical Drive Partition to a LUN	81
Using the Map Host LUN Option	83
▼ To Use the Map Host LUN Option	84
Example Host LUN Mappings	85
Deleting a Host LUN Mapping	86
▼ To Delete a Host LUN Mapping	86
Creating Host Filter Entries (FC Only)	87
Determining the Host World Wide Name	88
▼ To Determine the WWN in the Solaris Operating Environment	88
▼ To Determine the WWN for Linux, Windows NT, or Windows 2000	89
▼ To Determine the WWN for Other Supported Operating Systems	89

Creating Host Filter Entries	89
▼ To Create Host Filter Entries	90
Viewing and Modifying Host Filter Information	95
▼ To View or Modify Host Filter Information	95
6. Viewing and Editing SCSI Drives	97
Viewing the Physical Drive Status Table	98
SCSI Drive IDs (SCSI Only)	100
FC Drive IDs (Fibre Channel Only)	102
Assigning a Local Spare Drive	102
▼ To Assign a Local Spare Drive	103
Assigning a Global Spare	103
▼ To Assign a Global Spare	103
Viewing Drive Information	104
▼ To View SCSI Drive Information	104
Viewing the Connected Drives	105
▼ To View the List of Available SCSI Drives	105
Scanning Drives	106
▼ To Scan a New SCSI Drive	106
Deleting a Spare Drive	107
▼ To Delete a Spare Drive	107
Setting Slot Numbers	108
▼ To Set or Edit a Slot Number	108
Assigning a Slot Number to an Empty Slot	109
Deleting a Slot Number	109
▼ To Delete the Slot Number of a SCSI Drive	109
Adding or Deleting Drive Entries	110
▼ To Add a Drive Entry	110
▼ To Remove an Empty Drive Entry	111

Fault Protection Measures	111
Cloning a Failing Drive	111
Replacing After Clone	112
▼ To Replace After Clone	112
Perpetual Clone	114
▼ To Enable Perpetual Clone	114
Viewing the Status of a Cloning Operation	116
▼ To View the Status of a Cloning Operation	116
Understanding SMART Technology	117
How Sun StorEdge 3000 Family Arrays Use SMART	117
Enabling SMART From Firmware Menus	118
▼ To Enable SMART	118
Detect Only	120
▼ To Enable Detect Only	120
Detect and Perpetual Clone	120
▼ To Detect and Perpetual Clone	120
Detect and Clone+Replace	121
▼ To Detect and Clone+Replace	121
Identifying a Failed Drive for Replacement	121
▼ To Identify a Failed Drive	122
Flashing Selected SCSI Drive	123
Flashing All SCSI Drives	124
Flashing All But Selected Drive	124
SCSI Drive Utilities (Reserved)	125
▼ To Use the SCSI Drive Utilities (Reserved) Menu	125
SCSI Drive Low-Level Format (Reserved)	125
▼ To Low-Level Format a SCSI Drive (Reserved)	126
SCSI Drive Read/Write Test (Reserved)	126

▼ To Perform the SCSI Drive Read/Write Test (Reserved)	126
Changing Disk Drive Reserved Space	128
▼ To Change the Allocated Reserved Space on a Drive	128
7. Viewing and Editing SCSI Channels	129
SCSI Channel Status Table	130
▼ To Check and Configure SCSI Channels	130
Configuring SCSI Channels as Host or Drive	132
Fibre Channel Default Channel Settings (FC Only)	132
SCSI Default Channel Settings (SCSI Only)	133
▼ To Change the SCSI Channel Assignments (FC and SCSI)	133
Creating Additional Host IDs	135
Default Fibre Channel Host IDs (FC Only)	135
Default SCSI Host IDs (SCSI Only)	136
Adding a Host ID	136
▼ To Select a Unique ID Number for a Host Channel	136
Deleting a Host Channel SCSI ID	137
▼ To Delete a Host Channel SCSI ID	138
Drive Channel SCSI IDs (Reserved)	138
Setting SCSI Channel Termination (SCSI Only) (Reserved)	139
▼ To Enable or Disable SCSI Channel Termination (SCSI Only)	139
Setting Transfer Clock Speed (SCSI Only)	140
▼ To View the Options for Host or Drive Transfer Clock Speed (SCSI Only)	140
Setting the SCSI Transfer Width (SCSI Only)	141
▼ To View the Transfer Speed Options (SCSI Only)	141
Viewing and Editing Drive Channel SCSI Targets (SCSI Only)	141
▼ To View or Edit the SCSI Targets for a Selected Drive Channel (SCSI Only)	142
Providing a Slot Number (SCSI Only)	142

- ▼ To Provide a Slot Number (SCSI Only) 142
- Maximum Synchronous Transfer Clock (SCSI Only) 143
- ▼ To Set Maximum Synchronous Transfer Clock (SCSI Only) 143
- Maximum Transfer Width (SCSI Only) 144
- ▼ To Set Maximum Transfer Width (SCSI Only) 144
- Enabling Parity Checking (SCSI Only) 144
- ▼ To Enable or Disable Parity Checking 145
- Allowing Disconnect Support (SCSI Only) 145
- ▼ To Allow or Disallow Disconnect Support (SCSI Only) 145
- Maximum Tag Count (Reserved) (SCSI Only) 146
- Restoring a SCSI Target to Default Settings 146
 - ▼ To Restore Default SCSI Target Settings 146
- Viewing Chip Information 147
 - ▼ To View Chip Information 147
- Viewing Channel Host-ID WWN Information (FC Only) 148
 - ▼ To View a Channel's Host-ID/WWN (FC Only) 148
- Viewing Device Port Names (WWPN) (FC Only) 149
 - ▼ To View a Channel's Device Port Name List 149
- Setting a Channel's Data Rate (FC Only) 150
 - ▼ To Set a Channel's Data Rate 150

8. Viewing and Editing Configuration Parameters 153

- Communication Parameters 154
 - Setting an IP Address 154
 - ▼ To Set an Array's IP Address 155
- Caching Parameters 156
 - Enabling and Disabling Write-Back Cache 156
 - ▼ To Change the Caching Parameter Option 156
 - Optimization Modes 157

Optimization Limitations	157
Database and Transaction-Based Applications	158
Video Recording, Playback, and Imaging Applications	158
Optimization for Random I/O (32 Kbyte block size)	158
Optimization for Sequential I/O (128 Kbyte block size)	158
Maximum Number of Disks and Maximum Usable Capacity for Random and Sequential Optimization	159
▼ To Choose the Optimization Mode for All Drives	160
Host-side SCSI Parameters Menu Options	160
Maximum Queued I/O Count	161
▼ To Set the Maximum Queued I/O Count	161
LUNs Per Host SCSI ID	162
▼ To Change the Number of LUNs Per Host SCSI ID	162
Maximum Number of Concurrent Host-LUN Connections	163
▼ To Change the Maximum Number of Concurrent Host-LUN Connections	163
Number of Tags Reserved for Each Host LUN Connection	164
▼ To Modify the Tag Command Queuing on the Host-LUN Connection	164
Host Cylinder/Head/Sector Mapping Configuration	165
▼ To Configure Sector Ranges, Head Ranges, and Cylinder Ranges	166
Preparing for Logical Drives Larger Than 253 Gbytes on Solaris Systems	166
▼ To Prepare Logical Drives Larger than 253 Gbytes	167
Peripheral Device Type Parameters (Reserved)	168
Fibre Connection Options (FC Only)	168
▼ To Confirm or Change the Fibre Connection for the Array	169
Drive-side SCSI Parameters Menu	170
▼ To Access the Drive-side Parameter Menu	171
SCSI Motor Spin-Up (Reserved)	171
▼ To Spin-Up SCSI Hard Drives (Reserved)	171

SCSI Reset at Power-Up (Reserved)	172
▼ To Reset SCSI at Power-Up (Reserved)	172
Disk Access Delay Time	172
▼ To Set Disk Access Delay Time	172
SCSI I/O Timeout	173
▼ To Choose SCSI I/O Timeout	173
Maximum Tag Count (Tag Command Queuing)	173
▼ To Change the Default Maximum Tag Count Setting	174
Periodic Drive Check Time	175
▼ To Set the Periodic Drive Check Time	175
Periodic SAF-TE and SES Device Check Time	175
▼ To Set the Periodic SAF-TE and SES Device Check Time	175
Periodic Auto-Detect Failure Drive Swap Check Time	176
▼ To Set the Auto-Detect Failure Drive Swap Check Time	176
Drive Predictable Failure Mode (SMART)	177
Auto-Assign Global Spare Drive	177
▼ To Automatically Assign Replacements to Faulty Drives	177
Disk Array Parameters Menu	178
▼ To Display the Disk Array Parameters Menu	178
Setting Rebuild Priority	178
▼ To Set the Rebuild Priority	178
Verification on Writes	179
▼ To Enable and Disable Verification Methods	180
Redundant Controller Parameters Menu (Reserved)	181
▼ To Display the Redundant Controller Parameters Menu (Reserved)	181
Controller Parameters	181
Controller Name	181
▼ To View and Display the Controller Name	182

LCD Title Display - Controller Logo (Reserved)	182
Password Validation Timeout	182
▼ To Set a Password Validation Timeout	183
Controller Unique Identifier (Reserved)	184
▼ To Set the Controller Unique Identifier	184
SDRAM ECC Function (Reserved)	184
Controller Failover	185
Rebuilding Logical Drives	185
Automatic Logical Drive Rebuild	185
Manual Rebuild	188
Concurrent Rebuild in RAID 1+0	189
Recovering From Fatal Drive Failure	190
▼ To Recover From Fatal Drive Failure	190
9. Viewing and Editing Peripheral Devices	193
Viewing Peripheral Device Controller Status	193
Viewing SES Status (FC Only)	194
▼ To Check the Status of SES Components (FC Only)	195
Identifying Fans	197
▼ To View the Status of Each Fan	198
SES Temperature Sensor Locations (FC Only)	199
Viewing Peripheral Device SAF-TE Status (SCSI Only)	201
▼ To Check the Status of SAF-TE Components (SCSI Only)	201
SAF-TE Temperature Sensor Locations (SCSI Only)	203
Setting Peripheral Device Entry	203
Redundant Controller Mode	204
Enabling UPS Status	204
▼ To Enable the UPS Status	204
Setting the UPS Power Fail Signal	205

- ▼ To Set the UPS Power Fail Signal (Reserved) 205
 - Setting the Temperature Exceeds Threshold Trigger 205
 - ▼ To Configure Over-Temperature Controller Shutdown 205
 - Viewing Controller Voltage and Temperature Status 206
 - ▼ To Check the Controller Voltage and Temperature Status 206
 - Fibre Channel Error Statistics 209
- 10. System Functions and Event Logs 211**
 - Muting the Beeper 211
 - ▼ To Change the Beeper Setting 212
 - Setting and Changing the Controller Password 212
 - ▼ To Set a New Password 213
 - ▼ To Disable the Password 214
 - Resetting the Controller 214
 - ▼ To Reset the Controller Without Saving Cache Contents 215
 - Shutting Down the Controller 215
 - ▼ To Shut Down a Controller 216
 - Saving Your Configuration (NVRAM) to a File 216
 - ▼ To Save Your Configuration (NVRAM) to a File 219
 - Restoring Your Configuration (NVRAM) From a File 219
 - ▼ To Restore Configuration Settings from a Saved NVRAM File 220
 - Viewing Event Logs 221
 - ▼ To View the Event Logs 221
- A. Firmware Specifications 223**
- B. Parameter Summary Tables 229**
 - Introducing Default Parameters 229
 - Basic Default Parameters 230
 - Default Configuration Parameters 231

Default Peripheral Device Parameters	238
Default System Functions	239
Keeping Specific Parameter Defaults	240
C. Event Messages	241
Controller Events	242
Alerts	242
Warnings	242
Notifications	242
SCSI Drive Events	243
Warnings	243
Notifications	244
SCSI Channel Events	244
Alerts	244
Notifications	244
Logical Drive Events	245
Alerts	245
Notifications	246
General Target Alerts	246
SAF-TE Device	246
Controller On-Board	247
I ² C Device	247
SES Devices	248
General Peripheral Device	249
Glossary	251
Index	259

Preface

This guide provides an overview of RAID terminology and explains the use of controller firmware and VT100 commands to configure and monitor Sun StorEdge™ 3000 family arrays.

Note – Different versions of controller firmware apply to different Sun StorEdge 3000 family arrays. Before downloading new firmware, be sure to check the README file or appropriate release notes to make sure you are upgrading a supported version of the firmware for your array.

This guide is written for experienced system administrators who are familiar with Sun Microsystems hardware and software products.



Caution – Read the *Sun StorEdge 3000 Family Safety, Regulatory, and Compliance Manual* before beginning any procedure in this guide.

How This Book Is Organized

This book covers the following topics.

[Chapter 1](#) introduces RAID terminology and concepts.

[Chapter 2](#) explains 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 that should not be changed.

[Appendix C](#) lists the event messages.

[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 the following for this information:

- Software documentation that you received with your system
- Solaris™ operating environment documentation, which is at <http://docs.sun.com>

Shell Prompts

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	#

Typographic Conventions

Typeface ¹	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> .

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

Accessing Sun Documentation

All Sun StorEdge 3510 FC Array documentation is available online at the following locations:

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

<http://docs.sun.com/db/coll/3310SCSIarray>

http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3510

<http://docs.sun.com/db/coll/3510FCarray>

http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3511

<http://docs.sun.com/db/coll/3511FCarray>

Contacting Sun Technical Support

For late-breaking news and troubleshooting tips, review the release notes for your array, available at the locations shown in “[Accessing Sun Documentation](#)” on page xx.

If you have technical questions about this product that are not answered in the documentation, go to:

<http://www.sun.com/service/contacting>

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

800-USA4SUN

To obtain international technical support, contact the sales office of each country at:

<http://www.sun.com/service/contacting/sales.html>

508 Accessibility Features

The Sun StorEdge documentation is available in Section 508-compliant HTML files that can be used with assistive technology programs for visually impaired personnel. These files are provided on the Documentation CD for your product, as well as on the websites identified in “[Accessing Sun Documentation](#)” on page xx. Additionally, the software and firmware applications provide keyboard navigation and shortcuts that are documented in the user's guides.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

<http://www.sun.com/hwdocs/feedback>

Please include the title and part number of your document with your feedback: *Sun StorEdge 3000 Family RAID Firmware 3.25 and 3.27 User's Guide*, part number 817-3711-11.

Basic RAID Concepts

A redundant array of independent disks (RAID) offers major benefits in availability, capacity, and performance. Sun StorEdge 3000 family arrays provide complete RAID functionality and enhanced drive failure management.

This chapter covers the following concepts and planning guidelines:

- [“Before You Begin” on page 2](#)
 - [“Determining Which Version of the RAID Firmware Applies to Your Array” on page 2](#)
 - [“Fibre Channel and SCSI Firmware Illustrations in This Guide” on page 3](#)
- [“RAID Terminology Overview” on page 3](#)
 - [“Logical Drives” on page 3](#)
 - [“Spare Drives” on page 4](#)
 - [“Logical Volumes” on page 4](#)
 - [“SCSI Channels” on page 5](#)
- [“RAID Levels” on page 7](#)
 - [“RAID 0” on page 9](#)
 - [“RAID 1” on page 10](#)
 - [“RAID 1+0” on page 11](#)
 - [“RAID 3” on page 12](#)
 - [“RAID 5” on page 13](#)
 - [“Advanced RAID Levels” on page 14](#)
- [“Local and Global Spare Drives” on page 15](#)
 - [“Local Spare Drives” on page 15](#)
 - [“Global Spare Drives” on page 16](#)
 - [“Using Both Local and Global Spare Drives” on page 17](#)
- [“Controller Defaults and Limitations” on page 18](#)
 - [“Dual-Controller Considerations” on page 18](#)
 - [“Single-Controller Considerations” on page 20](#)
- [“Battery Operation” on page 20](#)
 - [“Battery Status” on page 20](#)
 - [“Write-Back Versus Write-Through Cache Options” on page 22](#)
- [“RAID Planning Considerations” on page 22](#)
- [“Basic Configuration Overview” on page 23](#)

Before You Begin

The firmware of the Sun StorEdge 3000 family arrays is software that is installed or “flashed” into the array hardware before it is shipped. Later versions of the firmware can be downloaded and flashed at the customer site.

Different versions of controller firmware apply to various Sun StorEdge 3000 family arrays. Before downloading new firmware, be sure to check the README file or appropriate release notes to make sure you are upgrading to a supported version of the firmware for your array.

Determining Which Version of the RAID Firmware Applies to Your Array

It is important that you run a version of firmware that is supported for your array. This manual covers the functionality for RAID firmware as follows:

- Firmware version 3.25 for Sun StorEdge 3310 SCSI arrays
- Firmware version 3.27 for Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 FC arrays with SATA.

Note – Throughout this document, the terms “FC arrays” and “Fibre Channel arrays” refer collectively to Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 FC arrays with SATA. Similarly, any menu option designated “FC Only” refer collectively to Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 FC arrays with SATA.

The firmware versions share most of the same functions; however, the values might differ. Features which are only for SCSI or FC arrays are noted in the manual.

If you are downloading a Sun Microsystems patch that includes a firmware upgrade, the README file associated with the patch tells you which Sun StorEdge 3000 family arrays support that firmware release.

Fibre Channel and SCSI Firmware Illustrations in This Guide

Illustrations in this guide demonstrate the steps you follow to use the firmware menu options and the results of those steps as they are displayed. The firmware menu options are the same for both SCSI and Fibre Channel (FC) arrays, so some of the illustrations describe SCSI arrays and other illustrations describe FC arrays. As a result, some of the device information you see on the screen differs slightly from what you see for your array.

RAID Terminology Overview

Redundant array of independent disks (RAID) is a storage technology used to improve the processing capability of storage systems. This technology is designed to provide reliability in disk array systems and to take advantage of the performance gains offered by an array of multiple disks over single-disk storage.

RAID's two primary underlying concepts are:

- Distributing data over multiple hard drives improves performance.
- Using multiple drives properly allows for any one drive to fail without loss of data and without system downtime.

In the event of a disk failure, disk access continues normally and the failure is transparent to the host system.

Logical Drives

Increased availability, capacity, and performance are achieved by creating logical drives. A logical drive is created by combining independent physical drives. To the host, the logical drive appears the same as a local hard disk drive.




FIGURE 1-1 Logical Drive Including Multiple Physical Drives

Logical drives can be configured to provide several distinct RAID levels, described in the remainder of this section.

Spare Drives

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 global spare drive is not reserved for a single logical drive. When a member drive from any of the logical drives fails, the global spare drive joins that logical drive and automatically starts to rebuild.

Logical Volumes

The concept of a logical volume is very similar to that of a logical drive. A logical volume is composed of one or more logical drives. The logical drives in a logical volume do not have to be composed of the same RAID level.

While the ability to create and manage logical volumes remains a feature of Sun StorEdge 3000 Family FC and SCSI RAID arrays for legacy reasons, the size and performance of physical and logical drives have made the use of logical volumes obsolete. Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Avoid using them and use logical drives instead. For more information about logical drives, see [“Viewing and Editing Logical Drives” on page 33](#).

A logical volume can be divided into a maximum of 32 partitions for SCSI arrays and 128 partitions for Fibre Channel arrays.

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

SCSI Channels

A SCSI channel can connect up to 15 devices (excluding the controller itself) when the Wide function is enabled (16-bit SCSI). Fibre Channel enables the connectivity of up to 125 devices in a loop. Each device has one unique ID.

A logical drive consists of a group of SCSI or Fibre Channel drives. Physical drives in one logical drive do not have to come from the same SCSI channel. Also, each logical drive can be configured for 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. A spare is not available for logical drives that have no data redundancy (RAID 0).

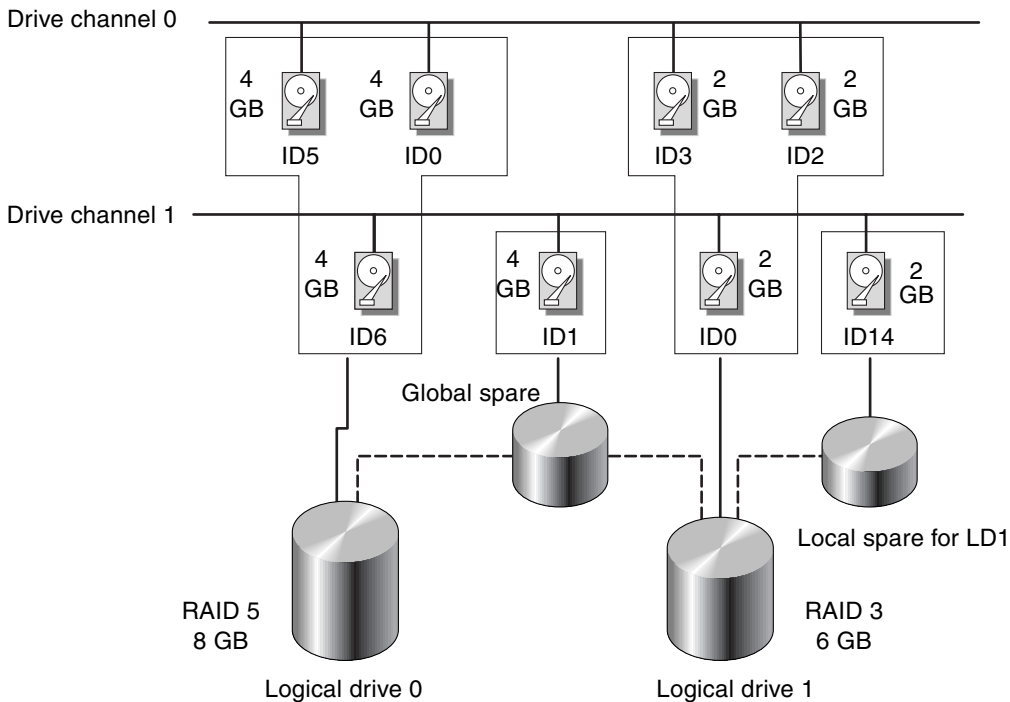


FIGURE 1-2 Allocation of Drives in Logical Drive Configurations

You can divide a logical drive or logical volume into several partitions or use the entire logical drive as a single partition.




FIGURE 1-3 Partitions in Logical Drive Configurations

Each partition is mapped to LUNs under host FC or SCSI IDs, or IDs on host channels. Each FC or SCSI ID/LUN is seen as an individual hard drive by the host computer.

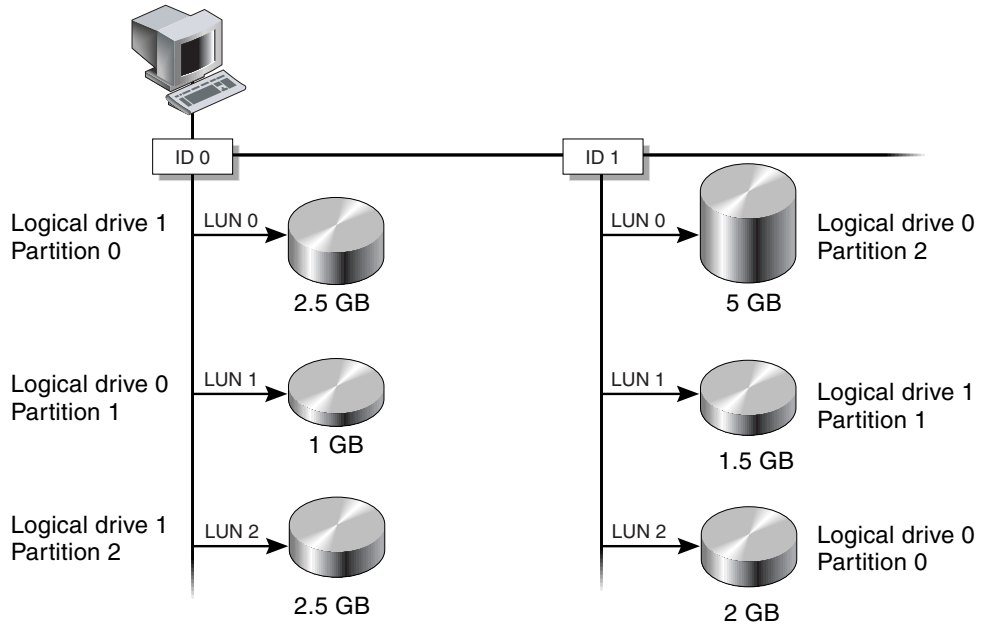


FIGURE 1-4 Mapping Partitions to Host ID/LUNs

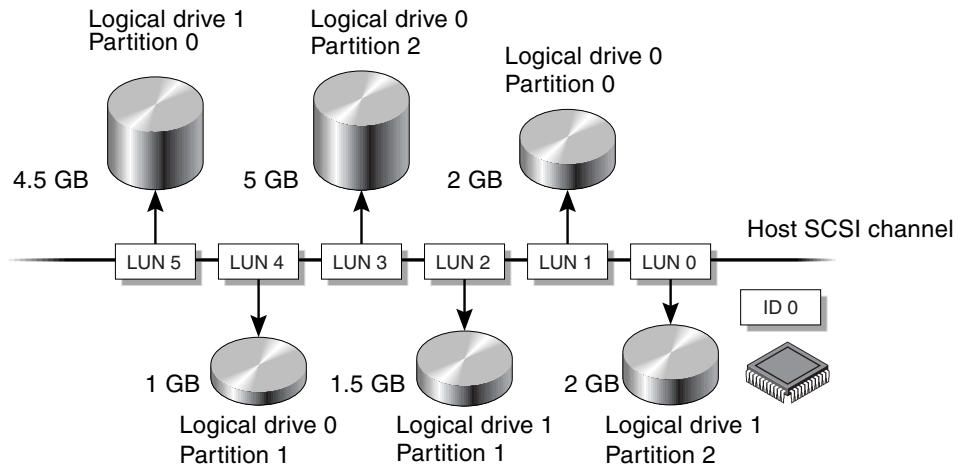


FIGURE 1-5 Mapping Partitions to LUNs Under an ID

RAID Levels

A RAID array has several advantages over non-RAID disk arrays:

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

There are several ways to implement a RAID array, using a combination of mirroring, striping, duplexing, and parity technologies. These various techniques are referred to as *RAID levels*. Each level offers a mix of performance, reliability, and cost. Each level uses a distinct algorithm to implement fault tolerance.

There are several RAID level choices: RAID 0, 1, 3, 5, 1+0, 3+0 (30), and 5+0 (50). RAID levels 1, 3, and 5 are most commonly used.

Note – The NRAID option that appears in some firmware menus is no longer used and is not recommended.

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

The following table provides a brief overview of the RAID levels.

TABLE 1-1 RAID Level Overview

RAID Level	Description	Number of Drives Supported	Capacity	Redundancy
0	Striping	2–36 physical drives	N	No
1	Mirroring	2 physical drives	N/2	Yes
1+0	Mirroring and striping	4–36 physical drives (even number only)	N/2	Yes
3	Striping with dedicated parity	3–31 physical drives	N-1	Yes
5	Striping with distributed parity	3–31 physical drives	N-1	Yes
3+0 (30)	Striping of RAID 3 logical drives	2–8 logical drives	N-# of logical drives	Yes
5+0 (50)	Striping of RAID 5 logical drives	2–8 logical drives	N-# of logical drives	Yes

Capacity refers to the total number (N) of physical drives available for data storage. For example, if the capacity is N-1 and the total number of disk drives in the logical drive is six 36-Gbyte drives, the disk space available for storage is equal to five disk drives (5 x 36 Gbyte or 180 Gbyte).

Note – The *-1* refers to the amount of striping across the example six drives, which provides redundancy of data and is equal to the size of one of the disk drives.

For RAID 3+0 (30) and 5+0 (50), *capacity* refers to the total number of physical drives (N) minus one physical drive (#) for each logical drive in the volume. For example, if the total number of disk drives in the logical drive is twenty 36-Gbyte drives and the total number of logical drives is 2, the disk space available for storage is equal to 18 disk drives (18 x 36 Gbyte or 648 Gbyte).

The advantages and disadvantages of different RAID levels are described in the following table.

TABLE 1-2 RAID Level Characteristics

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 multitasking or transaction processing. In RAID 5, an entire transfer block is placed on a single drive, but there are no dedicated data or parity drives. The data and parity are striped across each drive in the disk array, so that each drive contains a combination of data and parity 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 that: <ul style="list-style-type: none"> • It provides fault tolerance • It increases performance through the ability to perform both read and write seeks in parallel • The cost per usable megabyte of disk storage is low. RAID 5 requires at least 3 drives.
RAID 1+0	RAID 1+0 combines RAID 0 and RAID 1 to offer mirroring and disk striping. RAID 1+0 enables recovery from multiple drive failures because of the full redundancy of the hard disk drives. If four or more disk drives are selected for a RAID 1 logical drive, RAID 1+0 is performed automatically.
RAID (3+0)	A logical volume with several RAID 3 member logical drives.
RAID (5+0)	A logical volume with several RAID 5 member logical drives.

RAID 0

RAID 0 implements block striping, where data is broken into logical blocks and is striped across several drives. Unlike other RAID levels, there is no facility for redundancy. In the event of a disk failure, data is lost.

In block striping, the total disk capacity is equivalent to the sum of the capacities of all drives in the array. This combination of drives appears to the system as a single logical drive.

RAID 0 provides the highest performance. It is fast because data can be simultaneously transferred to and from every disk in the array. Furthermore, reads and writes to separate drives can be processed concurrently.




FIGURE 1-6 RAID 0 Configuration

RAID 1

RAID 1 implements disk mirroring, where a copy of the same data is recorded onto two drives. By keeping two copies of data on separate disks, data is protected against a disk failure. If, at any time, a disk in the RAID 1 array fails, the remaining good disk (copy) can provide all of the data needed, thus preventing downtime.

In disk mirroring, the total usable capacity is equivalent to the capacity of one drive in the RAID 1 array. Thus, combining two 1-Gbyte drives, for example, creates a single logical drive with a total usable capacity of 1 Gbyte. This combination of drives appears to the system as a single logical drive.

Note – RAID 1 does not allow expansion. RAID levels 3 and 5 permit expansion by adding drives to an existing array.




FIGURE 1-7 RAID 1 Configuration

In addition to the data protection that RAID 1 provides, this RAID level also improves performance. In cases where multiple concurrent I/O operations are occurring, these operations can be distributed between disk copies, thus reducing total effective data access time.

RAID 1+0

RAID 1+0 combines RAID 0 and RAID 1 to offer mirroring and disk striping. Using RAID 1+0 is a time-saving feature that enables you to configure a large number of disks for mirroring in one step. It is not a standard RAID level option that you can choose; it does not appear in the list of RAID level options supported by the controller. If four or more disk drives are selected for a RAID 1 logical drive, RAID 1+0 is performed automatically.

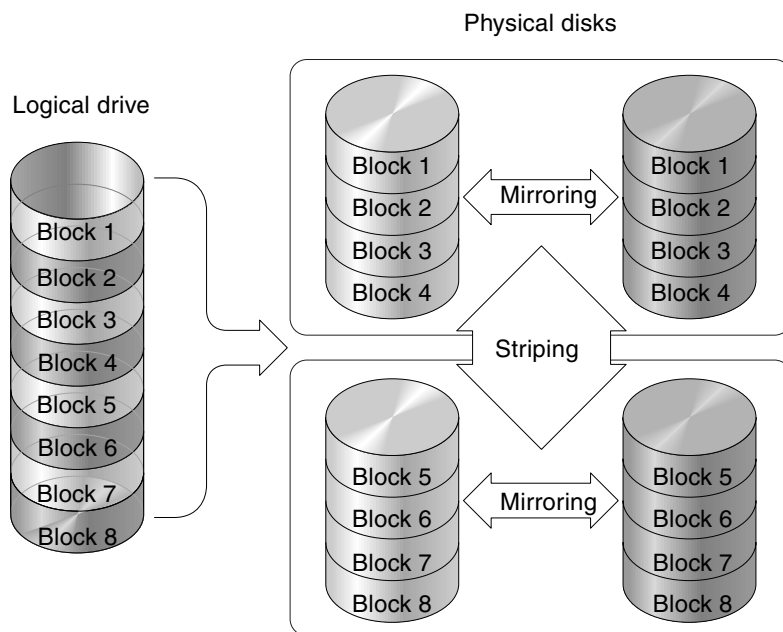


FIGURE 1-8 RAID 1+0 Configuration

RAID 3

RAID 3 implements block striping with dedicated parity. This RAID level breaks data into logical blocks, the size of a disk block, and then stripes these blocks across several drives. One drive is dedicated to parity. In the event that a disk fails, the original data can be reconstructed using the parity information and the information on the remaining disks.

In RAID 3, the total disk capacity is equivalent to the sum of the capacities of all drives in the combination, excluding the parity drive. Thus, combining four 1-Gbyte drives, for example, creates a single logical drive with a total usable capacity of 3 Gbyte. This combination appears to the system as a single logical drive.

RAID 3 improves data transfer rates when data is being read in small chunks or sequentially. However, in write operations that do not span every drive, performance is reduced because the information stored in the parity drive must be recalculated and rewritten every time new data is written, limiting simultaneous I/O.

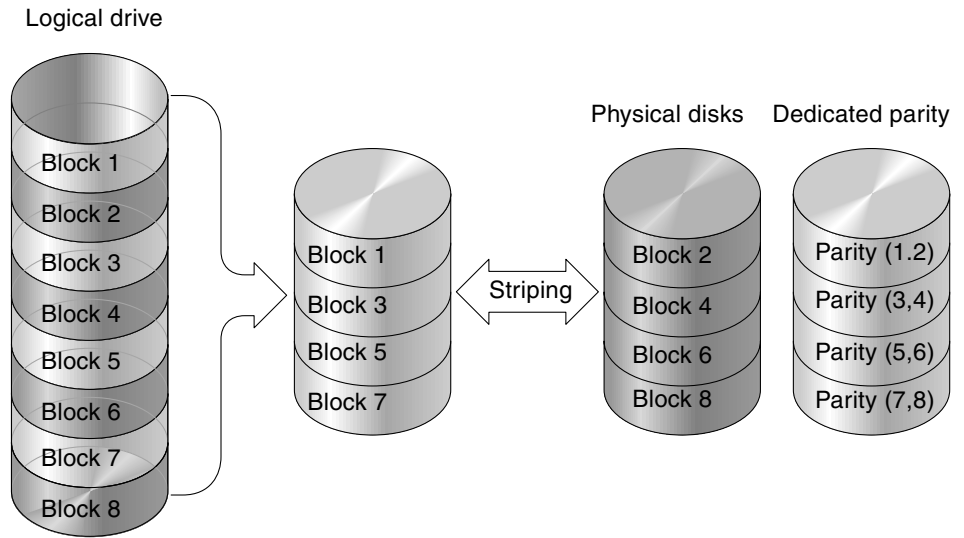


FIGURE 1-9 RAID 3 Configuration

RAID 5

RAID 5 implements multiple-block striping with distributed parity. This RAID level offers redundancy with the parity information distributed across all disks in the array. Data and its parity are never stored on the same disk. In the event that a disk fails, original data can be reconstructed using the parity information and the information on the remaining disks.

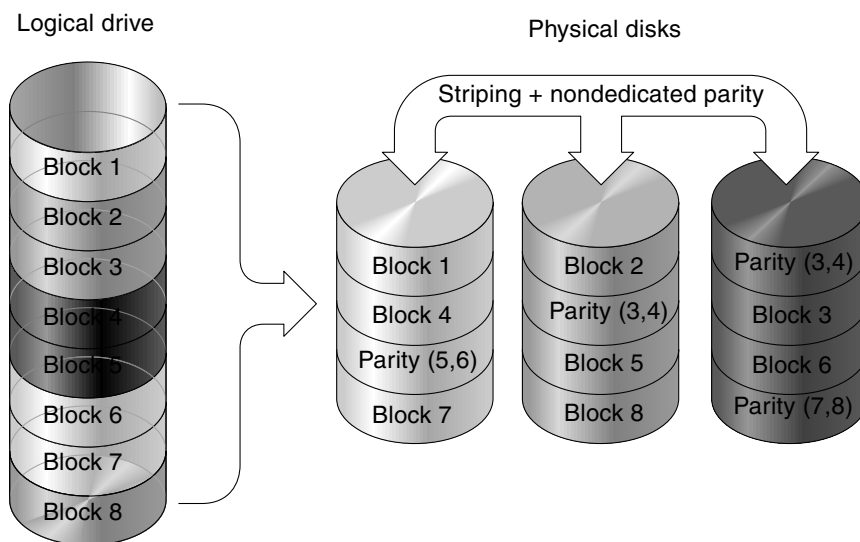


FIGURE 1-10 RAID 5 Configuration

RAID 5 offers increased data transfer rates when data is accessed randomly or in large chunks, and reduced data access time during simultaneous I/O operations.

Advanced RAID Levels

The following advanced RAID levels require the use of the array's built-in volume manager. These combination RAID levels provide the protection benefits of RAID 1, 3, or 5 with the performance of RAID 1. To use advanced RAID, first create two or more RAID 1, 3, or 5 arrays, and then join them.

The following table provides a description of the advanced RAID levels.

TABLE 1-3 Advanced RAID Levels

RAID Level	Description
RAID 3+0 (30)	RAID 3 logical drives that have been joined together using the array's built-in volume manager.
RAID 5+0 (50)	RAID 5 logical drives that have been joined together using the array's volume manager.

Local and Global Spare Drives

The external RAID controllers provide both local spare drive and global spare drive functions. A local spare drive is used only for one specified logical drive; a 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 global and local spares of sufficient capacity are both available, 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, see [“Identifying a Failed Drive for Replacement” on page 121](#).



Caution – If you mistakenly remove the wrong drive, you will no longer be able to access the logical drive because you have incorrectly failed two drives.

Local Spare Drives

A local spare drive is a standby drive assigned to serve one specified logical drive. If 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. If a drive fails and a local spare and a global spare drive are both available, the local spare drive is used.




FIGURE 1-11 Local (Dedicated) Spare

Global Spare Drives

A global spare drive is available to support all logical drives. If a member drive in any logical drive 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. If a drive fails and a local spare drive and a global spare drive of sufficient capacity are both available, the local spare drive is used.




FIGURE 1-12 Global Spare

Using Both Local and Global Spare Drives

In [FIGURE 1-13](#), the member drives in logical drive 0 are 9-Gbyte drives, and the members in logical drives 1 and 2 are all 4-Gbyte drives.

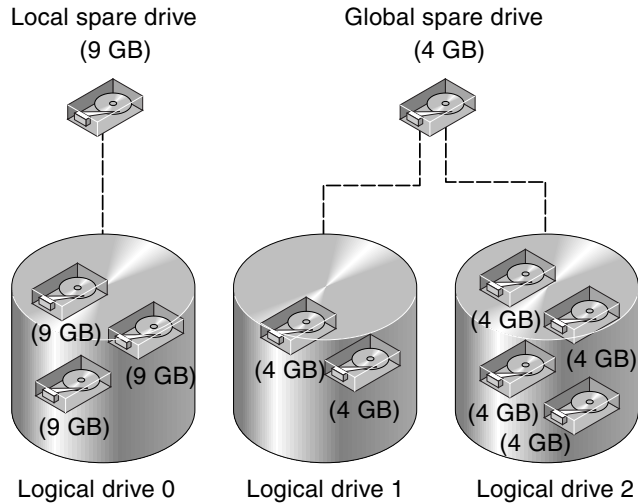


FIGURE 1-13 Mixing Local and Global Spares

A local spare drive always has higher priority than a global spare drive. If a drive fails and a local spare and a global spare drive of sufficient capacity are both available, the local spare drive is used.

In [FIGURE 1-13](#), it is not possible for the 4-Gbyte global spare drive to join logical drive 0 because of its insufficient capacity. The 9-Gbyte local spare drive is used for logical drive 0 once a drive in this logical drive fails. If the failed drive is in logical drive 1 or 2, the 4-Gbyte global spare drive is used immediately.

Controller Defaults and Limitations

This section describes default configurations and certain controller limitations.

Planning for Reliability, Availability, and Serviceability

The entry-level configuration for an FC array uses only one controller. You can mirror two single-controller arrays using volume manager software on attached servers to ensure high reliability, availability, and serviceability (RAS).

You can also use dual-controller arrays to avoid a single point of failure. A dual-controller FC array features a default active-to-active controller configuration. This configuration provides high reliability and high availability because, in the unlikely event of a controller failure, the array automatically fails over to a second controller, resulting in no interruption of data flow.

Other dual-controller configurations can be used as well. For instance, at a site where maximum throughput or connecting to the largest possible number of servers is of primary importance, you could use a high-performance configuration. Refer to the *Sun StorEdge 3000 Family Best Practices Manual* for the Sun StorEdge 3510 FC Array for information about array configurations.

Be aware, however, that departing from a high-availability configuration can result in a significant decrease in the mean time between data interruptions. System downtime, however, is not impacted as severely because the time required to replace a controller, if one is available, is only about five minutes.

Regardless of configuration, customers requiring high availability should stock field-replaceable units (FRUs) such as disk drives and controllers on-site. Your FC array has been designed to make replacing these FRUs easy and fast.

Dual-Controller Considerations

The following controller functions describe the redundant controller operation.

- Both controllers must be exactly the same. They must operate with the same firmware version, the same size of memory, and the same number of host and drive channels. If one controller in a dual-controller configuration is replaced with

a field-replaceable unit (FRU) controller, the array automatically compares the firmware versions of the two controllers. If they differ, the FRU controller firmware is automatically changed to that of the existing controller in the array.

- The controller firmware assumes that two controllers are available, or might be made available during operation at any time. In a 1U single-controller configuration, a 2U single-controller configuration, or a 2U dual-configuration, once the primary controller (which might be the only controller) is powered on, it begins to scan for a second controller. Until a second controller is discovered, which does not happen in a 1U single-controller configuration or a 2U single-controller configuration, the "Peripheral Device Status" for the redundant controller will show a Status of "Scanning." This is correct behavior and enables the firmware to discover a second controller whenever it is added without the necessity of rebooting the primary controller.
- After booting 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 logical drive to either of the controllers, and then map the logical configurations to host channel IDs/LUNs. I/O requests from a host computer are directed to the primary or the secondary controller accordingly. The total drive capacity can be grouped into several logical configurations and assigned to both controllers so that they share the workload. This active-to-active configuration engages all array resources to actively maximize performance.

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

Single-Controller Considerations

In a single-controller configuration, it is important to keep the controller as the primary controller at all times and to assign all logical drives to the primary controller. The primary controller controls all logical drive and firmware operations. In a single-controller configuration, the controller must be the primary controller or the controller cannot operate.

The secondary controller is only used in dual-controller configurations for redistributed I/O and for failover.

The Redundant Controller feature (reached by choosing “View and Edit Configuration Parameters → View and Edit Peripheral Devices → Set Peripheral Device Entry”) must remain enabled for single-controller configurations. This preserves the default primary controller assignment of the single controller.



Caution – Do not disable the Redundant Controller setting and do not set the controller as a secondary controller. If you disable the Redundant Controller setting and reconfigure the controller with the Autoconfigure option or as a secondary controller, the controller module becomes inoperable and must be replaced.

Note – For a single-controller configuration, the controller status shows “scanning” which indicates that the firmware is scanning for primary and secondary controller status and redundancy is enabled even though it is not used. There is no performance impact.

Battery Operation

The battery LED (on the far right side of the controller module) is an amber LED if the battery is bad or missing. The LED blinks green if the battery is charging and is solid green when the battery is fully charged.

Battery Status

The initial firmware screen also displays the battery status at the top of the initial screen where BAT: status displays somewhere in the range from BAD to ---- (charging), or +++++ (fully charged).

For maximum life, lithium ion batteries are not recharged until the charge level is very low, indicated by a status of -----. Automatic recharging at this point takes very little time.

A battery module whose status shows one or more + signs can support cache memory for 72 hours. As long as one or more + signs are displayed, your battery is performing correctly.

TABLE 1-4

Battery Display	Description
-----	Discharged; the battery is automatically recharged when it reaches this state.
+----	Adequately charged to maintain cache memory for 72 hours or more in case of power loss. Automatic recharging occurs when the battery status drops below this level.
++---	Over 90% charged; adequate to maintain cache memory for 72 hours or more in case of power loss.
+++--	Over 90% charged; adequate to maintain cache memory for 72 hours or more in case of power loss.
++++-	Over 90% charged; adequate to maintain cache memory for 72 hours or more in case of power loss.
+++++	Fully charged; adequate to maintain cache memory for 72 hours or more in case of power loss.

Your lithium ion battery should be changed every two years if the unit is continuously operated at 25 degrees C. If the unit is continuously operated at 35 degrees C or higher, it should be changed every year. The shelf life of your battery is three years.

Note – A safety precaution designed into your battery circuitry causes the battery to stop charging when the temperature of your array exceeds certain limits. When this happens, the battery status might be reported as BAD, but no alarm is written to the event log since no actual battery failure has occurred. This behavior is normal. As soon as the temperature returns to the normal range, battery charging resumes and the battery status is reported correctly. It is not necessary to replace or otherwise interfere with the battery in this situation.

For information on the date of manufacture and how to replace the battery module, refer to the *Sun StorEdge 3000 Family FRU Installation Guide*.

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 is not 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 can disable the Write-Back cache option and switch to the Write-Through cache option by choosing “view and edit Configuration parameters → Caching Parameters.”

RAID Planning Considerations

Here are some questions that can help you plan your RAID array.

- How many physical drives do you have?

You have from 5 drives to 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, see [“Default Logical Drives and RAID Levels” on page 34](#).

- What kind of host applications will you be using?

The frequency of read/write activities can vary from one host application to another. The application can be an SQL server, Oracle server, Informix server, or other database server of a transaction-based nature. Applications like video playback and video postproduction editing require read/write operations involving very large files in a sequential order.

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

The controller optimization mode can be changed only when there are no logical configurations. Once the controller optimization mode is set, the same mode is applied to all logical drives. Data stripe size is changed once the optimization method is changed. You cannot proceed with changing the optimization mode until data is backed up, all logical drives are deleted, and the array is restarted. Therefore, be careful in choosing an optimization mode for your controller.

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

- How many logical drives do you want, and at what RAID level?

A logical drive is a set of drives that have been combined into one logical drive to operate with a specified RAID level. It appears as a single contiguous storage volume. The controller is capable of grouping drives into eight logical drives, each configured with the same or different RAID levels. Different RAID levels provide varying degrees of performance and fault tolerance.

- Do you want to reserve any spare drives?

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

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

Basic Configuration Overview

This section briefly outlines steps you can take to configure your array.

▼ To Create a Basic Array Configuration

1. (Optional) Optimize the controller's parameters for your applications. For details on optimization modes, refer to ["Optimization Modes" on page 157](#).
2. If a hard drive was connected after the controller completes initialization, choose "view and edit scsi Drives → Scan scsi drive" to enable the controller to recognize the newly added hard drive and to make it available to be configured as a member of a logical drive.
3. (Optional) Define any additional partitions for each logical drive. See ["Partitioning a Logical Drive" on page 44](#).
4. (Optional) Add host IDs and more logical drives to create a maximum number of LUNs for your configuration. For more information see:
 - ["Planning for 128 LUNs on SCSI Arrays" on page 78](#)
 - ["Planning for 1024 LUNs on FC Arrays" on page 79](#)

- [“Planning for 64 LUNs in Redundant FC Point-to-Point Configurations”](#) on page 80
 - [“Creating Additional Host IDs”](#) on page 135
 - [“Creating Logical Drives”](#) on page 37
5. Map each logical drive and storage partition to a host ID/LUN. Refer to [“Mapping Logical Drive Partitions to Host LUNs”](#) on page 75. The host adapter recognizes the system drives after reinitializing the host bus.
 6. Save your configuration profile to disk.

Note – The controller is totally independent from the host operating environment. The host operating environment cannot determine 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 using a VT100-compatible terminal or a workstation running a terminal emulation program. Topics covered in this chapter include:

- “Setting Up the Serial Port Connection” on page 25
 - “To Set Up the Serial Port Connection” on page 26
- “Viewing the Initial Firmware Windows” on page 28
 - “Accessing the Firmware Application From a Solaris Host” on page 27
 - “To Use the `tip` Command” on page 27
- “Navigating Firmware Menus” on page 30
- “Quick Installation (Reserved)” on page 31
- “Upgrading Firmware” on page 31

Setting Up the Serial Port Connection

The RAID controller can be configured by means of a Solaris workstation running a VT100 terminal emulation program, or by a Microsoft Windows terminal emulation program such as HyperTerminal.

Note – You can also monitor and configure a RAID array over an IP network with the Sun StorEdge Configuration Service program after you assign an IP address to the array. For details, refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array or the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

▼ To Set Up the Serial Port Connection

1. Use a serial cable to connect the COM port of the RAID array to an unused serial port on your server.

A DB-9-to-DB-25 serial cable adapter is provided to connect the serial cable to DB-25 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.

It can take two or three minutes for both controllers to be initialized before communication is possible over the connection.

3. Set these serial port parameters on the workstation:

- 38,400 baud
- 8 bit
- 1 stop bit
- No parity

Note – Refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array for more detailed information about cabling and configuration procedures. Appendices provide specific information for each of the supported operating systems you can use to connect to the serial port of the array.

Accessing the Firmware Application From a Solaris Host

You can specify the baud rate of 38,400 when you type the `tip` command.

▼ To Use the `tip` Command

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

Use the `tip` command to access the array locally.

```
# tip -38400 /dev/ttyn
```

where *n* is the COM port identifier. For instance, if you have connected the array to the COM port identified as `ttyb`, use this command:

```
# tip -38400 /dev/ttyb
```

2. Refresh your screen by holding down the Control key on your keyboard (this is abbreviated Ctrl on some keyboards) and pressing the letter L key on your keyboard.

Viewing the Initial Firmware Windows

You see the initial controller screen (FIGURE 2-2) when you first access the RAID controller firmware.




FIGURE 2-2 Initial Screen of the Terminal Emulation Application.

The initial screen is displayed when the RAID controller is powered on. Use the up and down arrow keys to choose the VT100 terminal emulation mode, and then press Return to enter the Main Menu.

TABLE 2-1 Components of the Controller Firmware Screen

Component	Description
Cursor	Move the cursor to a desired item, and then press Return to select.
Controller name	Identifies the type of controller.
Progress indicator	Indicates the current data transfer rate or progress of an event.
Gauge range	Use + or – keys to change the gauge range in order to view the progress indicator.
Cache status	Indicates the percentage in the controller cache that differs from what is saved to disk.
PC graphic (ANSI mode)	Enters the Main Menu and operates in ANSI mode.

TABLE 2-1 Components of the Controller Firmware Screen (Continued)

Component	Description
(VT100 mode)	Enters the Main Menu and operates in VT100 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.

The progress indicator is displayed when necessary to indicate the percentage of completion of a particular task or event. Sometimes the event is represented by a descriptive title, such as "Drive Copying:"

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	488FFB88	NA	RAID5	103428	GOOD				7	B	4	0	0	
P1	34E	Drive Copying												
2		-----												
3		87% Completed												

Event messages showing full descriptive titles for the progress indicator include:

- Drive Copying
- Flash Erasing
- Flash Programming

For other events, the progress indicator merely shows a two-letter code in front of the percentage completed. These codes and their meanings are shown in the following table:

TABLE 2-2 Progress Indicator Prefix Meanings

Prefix	Description
IX:	Logical Drive Initialization
PX:	Parity Regeneration
EX:	Logical Drive Expansion
AX:	Add SCSI Drives

Navigating Firmware Menus

The firmware screens described in this document, along with the steps you follow, are the same regardless of whether you have connected to the controller through an IP address or through a serial port connection.

Note – The firmware functionality is similar for both Fibre Channel and SCSI arrays, but the contents of some informational screens vary between platforms.

Once you have chosen the VT 100 Terminal Emulation display mode, 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
```

FIGURE 2-3 Firmware Main Menu

Note – Since Fibre Channel and SCSI arrays share the same controller firmware, most menu options are the same. Parameter values might vary according to the product.

Use the following keys to navigate within the Main Menu and all its submenus.

Arrow keys	To select menu options
Return or Enter	To perform the selected menu option or display a submenu
Esc	To return to the previous menu without performing the selected menu option
Ctrl-L (Ctrl key and L key together)	To refresh the screen information
Press a letter as a keyboard shortcut for commands which have a boldface capital letter	To access a Main Menu command quickly

The firmware procedures use the term “Choose” as a shortcut description. Quotation marks are used to indicate a specific menu option or a series of menu options.

Procedure	Meaning
Choose “menu option”	Highlight the menu option and press Return. or Press the key that corresponds to the capitalized letter in the menu option if one is available.
Choose “menu option 1 → menu option 2 → menu option 3”	This represents a series of nested menus options which are selected with arrow keys. Press Return after each selection in order to access the next menu item and to complete the series.

Quick Installation (Reserved)



Caution – Do not use this command unless directed by Technical Support.

Upgrading Firmware

Firmware upgrades are performed by downloading firmware patches. For firmware upgrade procedures and any prerequisite upgrades, refer to the release notes for your array.

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 33
- “Default Logical Drives and RAID Levels” on page 34
- “Viewing the Logical Drive Status Table” on page 35
- “Creating Logical Drives” on page 37
- “Changing a Logical Drive Controller Assignment” on page 43
- “Partitioning a Logical Drive” on page 44
- “Using the Auto Configure Menu Option from the `format` Menu (Solaris Operating Environment Only)” on page 48
- “Deleting a Logical Drive” on page 51
- “Deleting a Logical Drive Partition” on page 53
- “Assigning a Logical Drive Name” on page 54
- “Rebuilding a Logical Drive” on page 55
- “Performing a Logical Drive Parity Check” on page 56
- “Overwriting Inconsistent Parity” on page 57
- “Adding a SCSI Drive to a Logical Drive” on page 58
- “Copying and Replacing Drives With Drives of Larger Capacity” on page 62
- “Expanding a Logical Drive” on page 64

Introducing the Logical Drive Commands

Use the “view and edit Logical drives” command to combine 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	O	#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
copy and replace drive

```

Under this menu command, you can create multiple logical drives. Once logical drives are created, you can manipulate them by choosing a desired logical drive.

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.

Default Logical Drives and RAID Levels

A logical drive is a set of drives grouped together to operate under a particular RAID level. Each RAID array is capable of supporting as many as eight logical drives.

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.

Spare can be part of an automatic array rebuild.

Note – A spare is not available for logical drives with no data redundancy (RAID 0).

The logical drives can have the same or different RAID levels.

For a 12-drive array, the RAID array is preconfigured as:

- Two RAID 5 logical drives, consisting of five physical drives in each logical drive
- Two global spares

For a 5-drive array, the RAID array is preconfigured as:

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

For more information about spares and RAID levels, refer to [“RAID Levels” on page 7](#) and [“Local and Global Spare Drives” on page 15](#).

Note – To reassign drives and add additional local or global spares on your preconfigured array, you must first unmap and then delete the existing logical drives, and then create new logical drives.

Viewing the Logical Drive Status Table

To check and configure logical drives, choose “view and edit Logical drives” from the Main Menu.

The status of all logical drives is displayed in a table.

Q	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
U	P0	27D0811C	NA	RAID0	206856	GOOD	S	6	-	0	
U	S1	710A07D9	NA	RAID0	206856	GOOD	S	6	-	0	
U	P2	1E6B7F1C	NA	RAID0	198000	GOOD	S	6	-	0	
U	S3	5BA0BD22	NA	RAID0	192000	GOOD	S	6	-	0	
S	4			NONE							
U	5			NONE							
	6			NONE							
	7			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.
DRV ABSENT	One of the disk drives cannot be detected.
INCOMPLETE	Not enough drives are available to support the current RAID configuration. This status is displayed only when the array is powering up.
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.

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

Parameter	Description
#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 members in the logical drive.
Name	Logical drive name (user-configurable).

Note – The SIZE (MB) parameter for a logical drive might not correspond exactly with the total size reported for each of the physical drives that make up the logical drive when using the “view and edit Logical drives” menu option. Any discrepancy is minor and is a result of how the drive manufacturers report their device size, which varies among manufacturers.

Creating Logical Drives

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

If you prefer a different configuration, use the procedure described in this section to modify the RAID level or to add more logical drives. 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 create the maximum number of LUNs for your array, you need a minimum of 8 logical drives per array and 32 LUNs per logical drive. SCSI arrays can have a maximum of 128 LUNs. FC arrays in loop configurations can have a maximum of 1024 LUNs, while FC arrays in redundant point-to-point configurations can have a maximum of 64 LUNs.

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.

In a RAID 1 or RAID 1+0 configuration the order in which you select the physical disks for a logical drive determines the channels to which they are assigned. If you wish drives to be mirrored over two channels, select them in the appropriate order. For example:

1. The first drive you select is assigned to channel 2, ID 0.
2. The second drive you select is assigned to channel 0, ID 0.
3. The third drive you select is assigned to channel 2, ID 1.
4. The fourth drive you select is assigned to channel 0, ID 1.

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

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

Note – This procedure is applicable to both SCSI and FC arrays. The illustrations show FC examples.

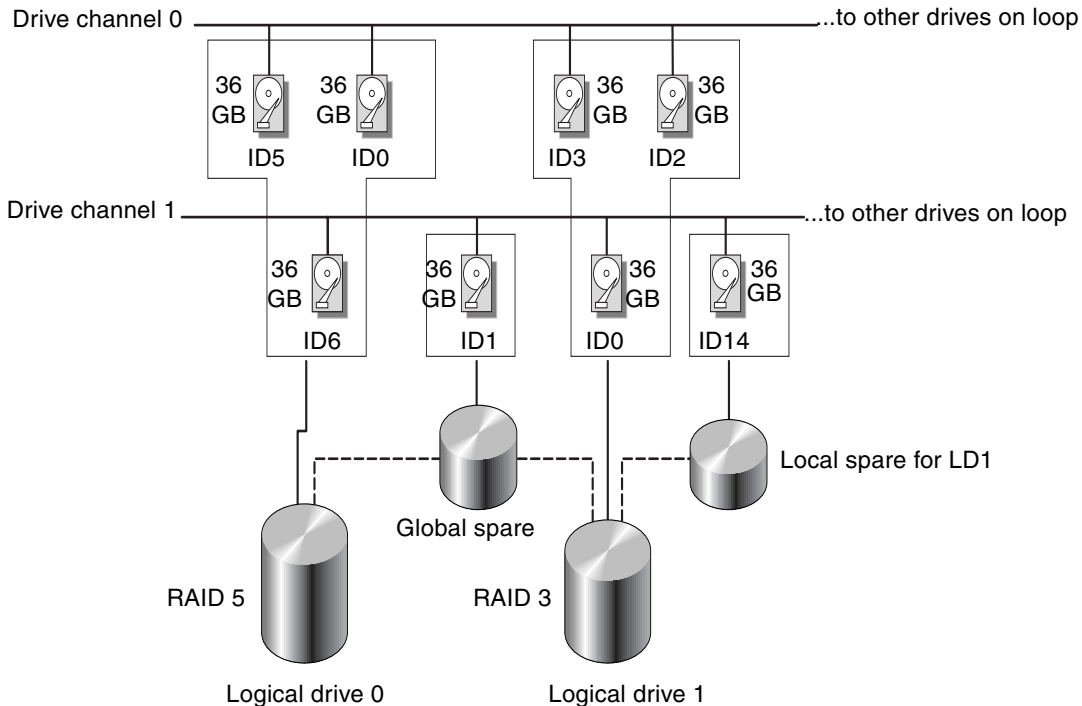


FIGURE 3-1 Example of an Allocation of Local and Global Spare Drives

▼ To Create Logical Drives

1. View the connected drives.

Before configuring disk drives into a logical drive, you must know the status of physical drives in your enclosure.

a. Choose “view and edit Scsi drives.”

A table shows information about all the physical drives that are installed.

	Slot	Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
Quick view		2<3>	0	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G
view		2<3>	1	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G
view		2<3>	2	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G
view		2<3>	3	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G
view		2<3>	4	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G
view		2<3>	5	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G
view		2<3>	6	34732	200MB	1	ON-LINE	SEAGATE ST336752FSUN36G
view		2<3>	7	34732	200MB	1	ON-LINE	SEAGATE ST336752FSUN36G

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. Refer to the *Sun StorEdge 3000 Family FRU Installation Guide for 2U Arrays* for information about scanning and replacing drives.

When the power is on, the controller scans all physical drives that are connected through the drive channels. If a physical drive was connected after the controller completed initialization, use the “Scan scsi drive” menu on the “view and edit scsi Drives” menu so the controller can recognize the newly added hard drive. Then you can configure the new drive 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 is lost.

2. Create a logical drive.

a. Choose “view and edit logical drive” from the Main Menu.

- b. Select the first available unassigned logical group (LG) to proceed.

Q	LG	ID	LU	RAID	Size<MB>	Status	O	#LN	#SB	#FL	NAME
U	P0	27D0811C	NA	RAID0	206856	GOOD	S	6	-	0	
U	S1	710A07D9	NA	RAID0	206856	GOOD	S	6	-	0	
U	P2	1E6B7F1C	NA	RAID0	198000	GOOD	S	6	-	0	
U	S3	5BA0BD22	NA	RAID0	192000	GOOD	S	6	-	0	
U	4			NONE							
U	5			NONE							
U	6			NONE							
U	7			NONE							

3. When prompted to “Create Logical Drive?” choose Yes to proceed.

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

4. Choose a RAID level for this logical drive.

Note – RAID 5 is used in this example.

For brief descriptions of RAID levels, refer to [“Default Logical Drives and RAID Levels” on page 34](#). See [“Basic RAID Concepts” on page 1](#) for more detailed information about RAID levels.

5. Select your member drives from the list of available physical drives.

The drives can be tagged for inclusion by highlighting the drive and then pressing Return. An asterisk (*) is displayed on the tagged physical drives.

To deselect the drive, move the cursor bar back to the physical drive and press Return again. The “*” disappears.

Note – You must select at least the minimum number of drives required for the RAID level you’ve chosen. See [“RAID Levels” on page 7](#) for information about the minimum numbers of drives required for each RAID level.

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

LG	ID	LV	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 the Esc key to continue to the next option.

After member physical drives are selected, a menu of configuration options is displayed.

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

6. (Optional) Set Maximum Physical Drive Capacity and assign spares.
 a. (Optional) Choose “Maximum Drive Capacity” from the menu.

Note – Changing the maximum drive capacity reduces the size of the logical drive and leaves 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. Replace the value in the Maximum Drive Capacity<MB> text area with the desired capacity value and press Return.
 c. (Optional) 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 local spare selected here automatically replaces any failed physical drive in this logical drive. The local spare is not 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	FRM DRV	SEAGATE ST336605LSUN36G
	0	13	34732	160MB	NONE	NEW DRV	SEAGATE ST336605LSUN36G

Note – A logical drive created in RAID level 0, which has no data redundancy, does not support spare drive rebuilding.

7. (Optional) Choose “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 ?
 Yes 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.



Caution – Do not assign logical drives to secondary controllers in single-controller arrays. Only the primary controller works in these arrays.

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

- b. Choose **Yes**, press **Return** to confirm, and then press the **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 indicates that the logical drive initialization has begun.

- d. Press the **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.
8. Choose “view and edit Logical drives” to view details of the created logical drive.

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 is 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 75](#)).



Caution – Do not assign logical drives to secondary controllers in single-controller arrays. Only the primary controller works in these arrays.

▼ To Change a Logical Drive Controller Assignment

1. To change a logical drive controller assignment, choose “view and edit Logical drives” from the Main Menu.
2. Select the drive you want to reassign.
3. Choose “logical drive Assignments.”

Q	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
Q	P0	27D0811C	NR	RAID0	206856	GOOD	S	6	-	0	
U											
V	S	View scsi drives				GOOD	S	6	-	0	
V		Delete logical drive									
V	P	Partition logical drive				GOOD	S	6	-	0	
V		logical drive Name									
V	S	logical drive assignments				GOOD	S	6	-	0	
V		Expand logical drive									
V		add Scsi drives									
V		copy and replace drive									
	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.

4. **Reassign the controller by choosing Yes.**
You must reset the controller for the changes to take effect.
5. **Press Esc to display the Main Menu.**
6. **From the Main Menu, choose “system Functions → Reset controller.”**

Partitioning a Logical Drive

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

For guidelines on setting up 128 LUNs on SCSI arrays, refer to [“Mapping Logical Drive Partitions to Host LUNs”](#) on page 75.

For guidelines on setting up 1024 LUNs on Fibre Channel arrays, refer to [“Planning for 1024 LUNs on FC Arrays”](#) on page 79.

Note – If you plan to map hundreds of LUNs, the process is easier if you use the Sun StorEdge Configuration Service program. Refer to *Sun StorEdge 3000 Family Configuration Service User’s Guide* for more information.



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

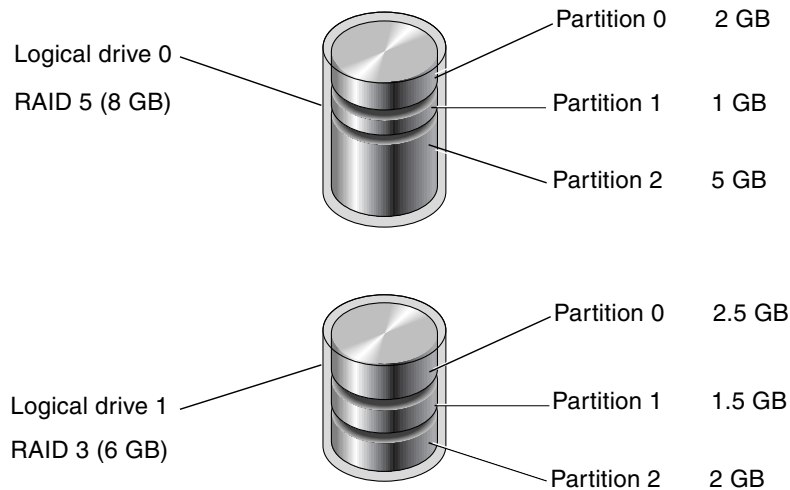


FIGURE 3-2 Partitions in Logical Configurations

▼ To Partition a Logical Drive

1. From the Main Menu, choose “view and edit Logical drives.”

2. Select the logical drive you want to partition.

Q	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
U	P0	27D0811C	NA	RAID0	206856	GOOD	S	6	-	0	
U	S1	710A07D9	NA	RAID0	206856	GOOD	S	6	-	0	
U	P2	1E6B7F1C	NA	RAID0	198000	GOOD	S	6	-	0	
U	S3	5BA0BD22	NA	RAID0	192000	GOOD	S	6	-	0	
U	4			NONE							
U	5			NONE							
U	6			NONE							
U	7			NONE							

3. Choose "Partition logical drive."

The following message is displayed:

Partitioning the Logical Drive will make it no longer eligible for membership in a logical volume.

Continue Partition Logical Drive?

4. Choose Yes to confirm that you want to partition the logical drive and do not want to include it in a logical volume.

Q	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
U	P0	27D0811C	NA	RAID0	206856	GOOD	S	6	-	0	
U	S	WARNING Partitioning the logical drive will make it no longer eligible for membership in a logical volume. Continue Partition Logical Drive ? Yes <input checked="" type="checkbox"/> No									
U	5			NONE							
U	6			NONE							
U	7			NONE							

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 is listed as "partition 0."

5. Select from the list of undefined partitions.

- Type the desired size for the selected partition and press Return.

Q	LG	ID	LU	RAID	Size<MB>	Partition	Offset<MB>	Size<MB>	NAME
U	P0	B61E5AB	NA	RAID5	30000	0	0	30000	
U	1			NONE					
U	2			NONE					

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?



Caution – Make sure any data on this partition that you want to save has been backed up before you partition the logical drive.

- Choose Yes to confirm.

The remaining capacity of the logical drive is automatically allotted to the next partition. In the following figure, a partition size of 3000 MB was entered; the remaining storage of 27,000 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	B61E5AB	NA	RAID5	30000	0	0	3000	
U	1			NONE		1	3000	27000	
U	2			NONE		2			
S	3			NONE		3			

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

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

Note – When you modify a partition or logical drive size, you must reconfigure all host LUN mappings. All the host LUN mappings are removed with any change to partition capacity. See [“Mapping a Partition to a LUN” on page 81](#).

Using the Auto Configure Menu Option from the `format` Menu (Solaris Operating Environment Only)

If you are using a Solaris host, you should use the “Auto configure” menu option from the `format` menu whenever you resize a logical drive.

For the Solaris operating environment to recognize a LUN, it must be labeled with the `format` command. Whenever you create or reconfigure a LUN, label it using the following procedure. For more information on the `format` command, see the `format(1M)` man page.

Note – This example applies to the Sun StorEdge 3310 SCSI array. Follow the same steps for Sun StorEdge 3510 FC arrays as well, using the appropriate drive type and label.

▼ To Change the Label of a LUN

1. **On the data host, type `format` at the root prompt.**

```
# format
```

2. **Specify the disk number when prompted.**

In this example, the array is shown as disk number 2. It can be identified by the `SUN-StorEdge3310-0325` label.

```

# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
    0. c0t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@1f,4000/scsi@3/sd@0,0
    1. c7t0d0 <SUN-StorEdge3310-0325 cyl 43774 alt 2 hd 127 sec 127>
       /pci@4,2000/pci@1/scsi@4/sd@0,0
    2. c7t1d0 <SUN-StorEdge3310-0325 cyl 43774 alt 2 hd 127 sec 127>
       /pci@4,2000/pci@1/scsi@4/sd@1,0
    3. c8t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@4,2000/pci@1/scsi@5/sd@0,0
    4. c8t1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@4,2000/pci@1/scsi@5/sd@1,0
    5. c8t2d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@4,2000/pci@1/scsi@5/sd@2,0
    6. c8t3d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@4,2000/pci@1/scsi@5/sd@3,0
    7. c8t4d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@4,2000/pci@1/scsi@5/sd@4,0
    8. c8t5d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@4,2000/pci@1/scsi@5/sd@5,0
    9. c8t8d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@4,2000/pci@1/scsi@5/sd@8,0
   10. c8t9d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
       /pci@4,2000/pci@1/scsi@5/sd@9,0
Specify disk (enter its number): 2
selecting c7t1d0
[disk formatted]

```

- If the volume had been labeled previously using the format command, the next display is the FORMAT MENU. If the FORMAT MENU is displayed, continue to the next step.
- If the volume had not been labeled previously using the format command, you are prompted with the question: Disk not labeled. Label it now? If so, answer y at the prompt and press Return.

The FORMAT MENU is displayed.

3. Type type to select a drive type.

```
FORMAT MENU:
disk - select a disk
type - select (define) a disk type
partition - select (define) a partition table
current - describe the current disk
format - format and analyze the disk
repair - repair a defective sector
label - write label to the disk
analyze - surface analysis
defect - defect list management
backup - search for backup labels
verify - read and display labels
save - save new disk/partition definitions
inquiry - show vendor, product and revision
volname - set 8-character volume name
!<cmd> - execute <cmd>, then return 0
quit
format> type
```

4. Type 0 to choose the Auto configure menu option.

Select the Auto configure menu option regardless of which drive types are displayed by the type option. For further information on partitioning, see the Solaris format(1M) man page.


```
AVAILABLE DRIVE TYPES:
    0. Auto configure
    1. Quantum ProDrive 80S
    2. Quantum ProDrive 105S
    3. CDC Wren IV 94171-344
    4. SUN0104
    5. SUN0207
    6. SUN0327
    7. SUN0340
    8. SUN0424
    9. SUN0535
   10. SUN0669
   11. SUN1.0G
   12. SUN1.05
   13. SUN1.3G
   14. SUN2.1G
   15. SUN2.9G
   16. Zip 100
   17. Zip 250
   18. SUN36G
   19. SUN-StorEdge3310-0325
   20. other

Specify disk type (enter its number)[19]: 0
c7t1d0: configured with capacity of 336.66GB
<SUN-StorEdge3310-0325 cyl 43774 alt 2 hd 127 sec 127>
selecting c7t1d0
[disk formatted]
```

5. Type label and answer y when prompted to continue.

```
format> label
Ready to label disk, continue? y
```

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 reconfigure the entire array with different logical drives and RAID levels, you must first delete the existing logical drives and reconfigure the array prior to storing data on the array.



Caution – This operation erases all data on the logical drive. Therefore, if any data exists on the logical drive, copy it to another location before deleting the current logical drive.

Note – You can only delete a logical drive if it has first been unmapped.

▼ To Delete a Logical Drive

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the logical drive to be deleted.

Q	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
U	P0	27D0811C	NA	RAID0	206856	GOOD	S	6	-	0	
U	S1	710A07D9	NA	RAID0	206856	GOOD	S	6	-	0	
U	P2	1E6B7F1C	NA	RAID0	198000	GOOD	S	6	-	0	
U	S3	5BA0BD22	NA	RAID0	192000	GOOD	S	6	-	0	
U	4			NONE							
U	5			NONE							
U	6			NONE							
U	7			NONE							

3. Choose “Delete logical drive.”

A warning notice is displayed.

4. If it is safe to delete the logical drive, choose Yes.

Q	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
U	P0	4DB84961	NA	RAID5	4000	GOOD	S	3	0	0	
U	View scsi drives										
U	Delete logical drive										
U	This operation will result in the LOSS OF ALL DATA on the logical Drive ? Delete Logical Drive ? <input checked="" type="checkbox"/> Yes <input type="checkbox"/> No										
U	6			NONE							
U	7			NONE							

Deleting a Logical Drive Partition

When you delete a partition on a logical drive, all of the capacity of the deleted partition is added to the remaining partition with the highest partition number. If a logical drive contains partitions 0 through 3 and you delete partition 3, for example, all of the capacity of partition 3 is added to partition 2.

▼ To Delete a Logical Drive Partition

1. Choose “view and edit Logical drives” from the Main Menu.
2. Select the logical drive that has a partition you want to delete.
3. Choose “Partition logical drive.”
Partitions of the logical drive are displayed in a table.
4. Select the partition you want to delete.
5. Type “0” in the Partition Size (MB) text area to delete this partition.

Q	LG	ID	LV	RAID	Size(MB)	Partition	Offset(MB)	Size(MB)	NAME
Q	P0	4149A729	NA	RAID0	3999	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			

Partition Size (MB): 0

A warning prompt is displayed:

This operation will result in the LOSS OF ALL DATA on the partition.

Partition Logical Drive?

As illustrated in [FIGURE 3-3](#), the capacity of the deleted partition 1 (200 MB) is added into the last partition (now partition 2), which combines 400 MB and 200 MB for a total of 600 MB.

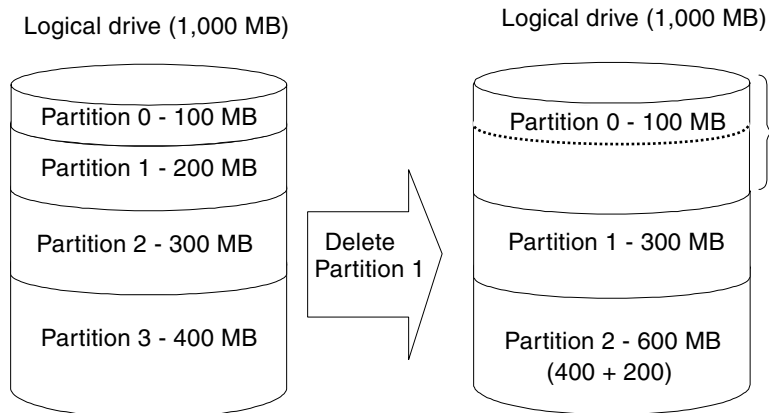


FIGURE 3-3 Example of Deleted Partitions



Caution – If a partition has been changed, you must reconfigure all host LUN mappings. All the host LUN mappings are removed with any partition change.

Assigning a Logical Drive Name

You can assign a name to a logical drive after the logical drive is created. This logical drive name is used only in the RAID firmware administration and monitoring, and does not appear anywhere on the host. You can also edit this drive name.

▼ To Assign a Logical Drive Name

1. Choose “view and edit Logical drives” from the Main Menu.
2. Select the logical drive that you want to name.

3. Choose “logical drive Name.”

Q	LG	ID	LU	RAID	Size(MB)	Statu	O	#LN	#SB	#FL	NAME
U	P0	27D0811C	NR	RAID0	206856	GOOD	S	6	-	0	
U	S	View scsi drives				GOOD	S	6	-	0	
U	P	View scsi drives				GOOD	S	6	-	0	
U	S	Delete logical drive				GOOD	S	6	-	0	
U	S	Partition logical drive				GOOD	S	6	-	0	
U		logical drive Name									
U		logical drive Assignments				GOOD	S	6	-	0	
E		Current Logical Drive Name: New Logical Drive Name: _									
a											
c											
7											

4. Type the name you want to give the logical drive in the “New Logical Drive Name” text area and press Return to save the name.

Rebuilding a Logical Drive

If no spare drive is available for logical drive rebuilding, replace the failed drive immediately with a new drive and initiate the rebuild process manually.

▼ To Rebuild a Logical Drive

1. Choose “view and edit Logical drives” from the Main Menu.
2. Select the logical drive that has a failed member drive.
3. Choose “Rebuild logical drive.”

4. **When prompted to regenerate the parity, choose Yes.**
A notification is displayed that parity regeneration has begun.
5. **Press Esc to see a progress indicator.**

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

Overwriting Inconsistent Parity

In normal operation, many regions of physical disks might not be accessed for long periods of time. The primary purpose of the parity checking operation in RAID 3 and RAID 5 configurations is to ensure that all sectors of the media can be successfully read and to provide alerts if a disk is starting to experience read or write errors.

If the result of a parity check indicates an inconsistency, it indicates that a data error exists either on one of the disk drives or on the parity drive. However, RAID algorithms such as RAID 5 that use XOR parity cannot determine if the error resides on a data disk or on the parity disk. Enabling the “Overwrite Inconsistent Parity” menu option causes the RAID controller to correct the data on the parity disk automatically whenever a parity check detects an inconsistency. In most cases it is important to correct the data on the parity disk as soon as an inconsistency is detected, in order to avoid the potential for data loss in the event of a drive failure.

However, some users prefer to check the integrity of their data before overwriting the parity disk. Disabling the “Overwrite Inconsistent Parity” menu option causes the controller to report any inconsistency found in a parity check without overwriting the parity disk. In this case you can check your data and determine if it is intact or if the error occurred on a data disk. Once you’ve made this determination, and reloaded the data from backup if necessary, you can manually regenerate the parity using the “reGenerate parity” menu option.

▼ To Enable or Disable the “Overwrite Inconsistent Parity” Menu Option

1. **Choose “view and edit Logical drives” from the Main Menu.**
2. **Select the logical drive whose automatic parity area overwriting you want to enable or disable.**

3. To toggle between enabling and disabling this menu option, choose “reGenerate parity → Overwrite Inconsistent Parity - .”

Adding a SCSI Drive to a Logical Drive

RAID expansion allows users to expand a logical drive by adding new physical 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.

Expansion can be performed only on RAID 0, 3 and 5 logical drives. Expansion cannot be performed on a 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 200-Gbyte logical drive and add a 36-Gbyte drive, the total logical drive is 236 Gbytes with two partitions (one 200-Gbyte partition and one 36-Gbyte partition).

The new partition must be mapped to a host LUN for the HBA 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 hard disk drives cannot be canceled once begun. If a power failure occurs, the expansion pauses and the controller does 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 pauses. The expansion resumes automatically after the logical drive rebuild has completed.

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

When expansion is completed, data is re-striped across the original and the newly added drives. After adding drives, the additional capacity is displayed as a new partition, as shown in [FIGURE 3-4](#).

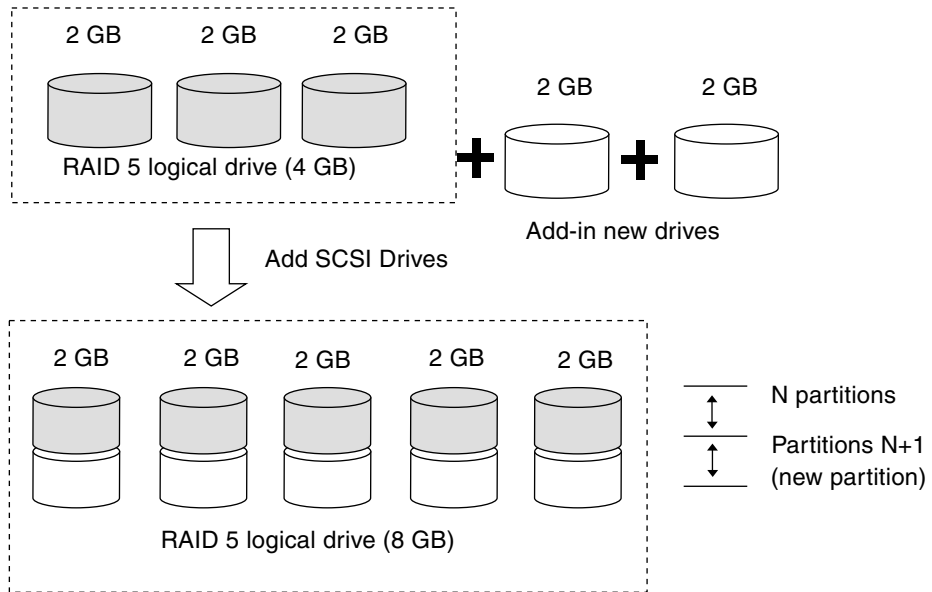


FIGURE 3-4 Expansion by Adding Drive

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

▼ To Add a SCSI Drive to a Logical Drive

1. From the Main Menu, choose "view and edit logical drive."
2. Select a logical drive.

Note – Only logical drives that are configured as RAID 0, RAID 3, or RAID 5 support expansion by adding physical drives.

3. Choose "add Scsi drives."
A confirmation prompt is displayed.

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. 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 newly created area becomes a new partition.

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

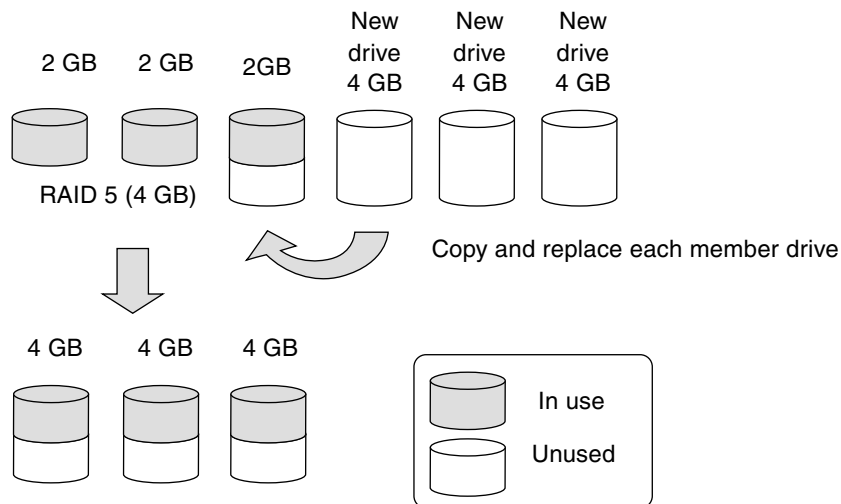


FIGURE 3-5 Expansion by Copy and Replace

▼ To Copy and Replace a Drive

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select a target logical drive.

- Repeat these steps to copy and replace every member drive with drives of higher capacity as needed.

You can 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
    [Z1A2] LG:0 Logical Drive NOTICE:CHL:1 ID:3 Copy and Replace Completed
  
```

Expanding a Logical Drive

You 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 drives to a logical drive. After a logical drive is expanded, the additional capacity is displayed as a new partition.

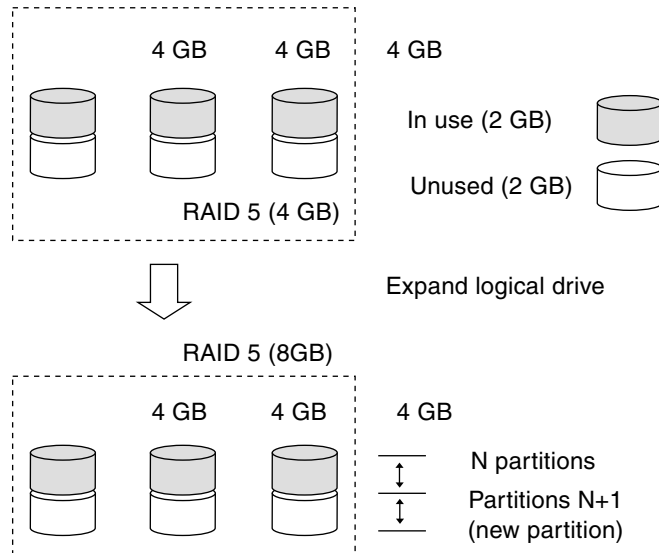


FIGURE 3-6 Logical Drive Expansion

Note – The newly created area becomes 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 to recognize its presence. To add the new partition into an existing logical drive, operating system 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 Gbyte. 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 Gbytes. The next step is to perform the “Expand logical drive” function so you can access the additional capacity of the new drives.

▼ To Expand a Logical Drive

1. Choose “view and edit Logical drives” from the Main Menu.
2. Select the logical drive with its members copied and replaced.
3. Choose “Expand logical drive.”

Drive capacity information is displayed.

4. Proceed by pressing Return, or by entering any value no larger than the “maximum drive expand capacity” and then pressing Return.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
V	P0	0499A7C9	INA	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_

A confirmation message is displayed.

Viewing and Editing Logical Volumes

This chapter describes the procedures for creating and using logical volumes.

Topics covered in this chapter include:

- [“Avoiding the Use of Logical Volumes” on page 67](#)
- [“Understanding Logical Volumes \(Multilevel RAID\)” on page 68](#)
- [“Creating a Logical Volume” on page 70](#)
- [“Expanding a Logical Volume” on page 72](#)
- [“Viewing the Logical Volume Status Table” on page 73](#)

Avoiding the Use of Logical Volumes

While the ability to create and manage logical volumes remains a feature of Sun StorEdge 3000 Family FC and SCSI RAID arrays for legacy reasons, the size and performance of physical and logical drives have made the use of logical volumes obsolete. Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Avoid using them and use logical drives instead. For more information about logical drives, see [“Viewing and Editing Logical Drives” on page 33](#).

Understanding Logical Volumes (Multilevel RAID)

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.

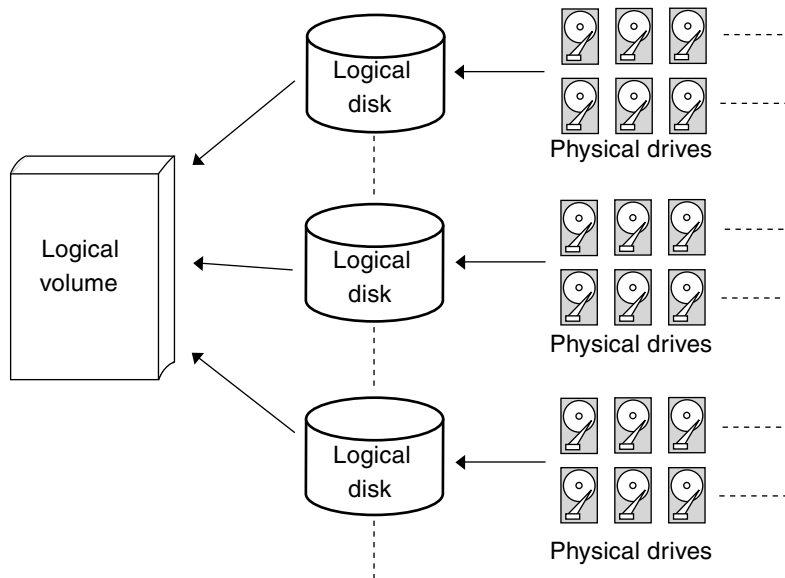


FIGURE 4-1 Logical Volume Composed of Multiple Drives

A logical volume can be divided into a maximum of 32 partitions for SCSI arrays and 128 partitions for Fibre Channel arrays. During normal operation, the host sees an unpartitioned logical volume or a partition of a partitioned logical volume as a single physical drive.

Logical Volume Limitations

A 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 starts 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:

- Configure logical drives as members of a logical volume in RAID levels that provide redundancy, such as RAID levels 1, 1+0, 3, or 5.
- Rebuild the logical drive as soon as possible whenever a drive failure occurs.
- Compose the logical drive with drives from different drive channels to avoid the fatal loss of data caused by bus failure.

Partitioning

Once a logical drive has been divided into partitions, the logical drive can no longer be used as a member of a logical volume. If you want to use a partitioned logical drive for a logical volume, delete the other partitions in this logical drive until only one partition remains that uses the entire capacity of the logical drive.



Caution – Deleting the partition of the logical drive destroys all the data. Data should be backed up before changing partition configurations.

When a logical drive is used as a member to a logical volume, this logical drive can no longer be partitioned using the “View and Edit Logical Drives” menu option. Instead, the logical volume can be partitioned using the “view and edit logical Volumes” menu option.

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

A logical volume can be expanded using the RAID expansion function. Expanding a logical volume is similar to 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 and then perform logical drive expansion to utilize the newly added capacity. To perform RAID expansion on a logical volume, expand each member logical drive and then perform RAID expansion on the logical volume.

Multilevel RAID Arrays

A multilevel RAID array includes logical drives of different RAID levels. A multilevel RAID array with logical volume support offers the following available configurations.

- **RAID 1+0.** This is a standard feature of Sun StorEdge 3000 family RAID controllers. It has the benefits of RAID 1 (high availability) and RAID 0 (enhanced I/O performance through striping). Simply choose four or more drives for a RAID 1 logical drive and the RAID controller implements RAID 1+0 automatically.
- **RAID (3+0).** A logical volume itself is a multilevel 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 controller 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.

Spare Drives

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.

Creating a Logical Volume

A logical volume consists of one or several logical drives.

▼ To Create a Logical Volume

1. Choose “view and edit logical Volumes” from the Main Menu.

The current logical volume configuration and status are displayed.

Q	LV	ID	Size(MB)	#LD	s
V	0				es
V	1				
V	2				
V	3				
V	4				
V	5				
V	6				
V	7				

parameters
choices

2. Select a logical volume number (0–7) that has not yet been defined.

A confirmation message 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. Choose Yes to confirm.

Q	LV	ID	Size(MB)	#LD	s							
V	0				es							
V	1	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
V	2	MP0	69FC88B	NA	RAID5	39996	GOOD	R	4	1	0	
V	3	MP1	55A373B	NA	RAID0	19998	GOOD	R	2	-	0	
V	4											
V	5											
V	6											
V	7											

4. Select one or more available logical drives from the list, and press Return to tag the logical drives to be included in the volume.

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

Note – Press Return to deselect a tagged logical drive.

LV	ID	Size(MB)	#LD	
0				5 es
1	Logical Volume Assignments			
2	Redundant Controller Logical Volume Assign to Secondary Controller ?			
3			<input checked="" type="checkbox"/> 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.

5. Once 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. Choose Yes to create the logical volume.

6. Press Return to display the information about the logical volume you created.

Expanding a Logical Volume

This section describes how to expand a logical volume.

▼ To Expand a Logical Volume

1. Expand the logical drives in the logical volume.
2. Choose "Expand logical volume."

A confirmation message asks if you want to expand the logical volume.

3. Choose Yes to confirm.

Viewing the Logical Volume Status Table

To check and configure logical drives, from the Main Menu choose “view and edit Logical drives.”

The status of all logical drives is displayed:

Q	LV	ID	Size(MB)	#LD	
v	P0	466C5C8D	60000	1	s
v	1				es
v	2				parameters vices
v	3				
s	4				
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 drives 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 logical drive (LD) or logical volume (LV) can be mapped more than once to achieve a redundant data path, using multipathing software.

Topics covered in this chapter include:

- [“Mapping Logical Drive Partitions to Host LUNs” on page 75](#)
- [“Planning for 128 LUNs on SCSI Arrays” on page 78](#)
- [“Planning for 1024 LUNs on FC Arrays” on page 79](#)
- [“Planning for 64 LUNs in Redundant FC Point-to-Point Configurations” on page 80](#)
- [“Mapping a Partition to a LUN” on page 81](#)
 - [“Using the Map Host LUN Option” on page 83](#)
 - [“Example Host LUN Mappings” on page 85](#)
 - [“Determining the Host World Wide Name” on page 88](#)
 - [“Creating Host Filter Entries” on page 89](#)
- [“Deleting a Host LUN Mapping” on page 86](#)
- [“Creating Host Filter Entries \(FC Only\)” on page 87](#)

Mapping Logical Drive Partitions to Host LUNs

A logical unit number (LUN) is a unique identifier used on a SCSI channel that enables a host to differentiate between separate devices.

Once you have created logical drives or logical volumes, you can map each storage partition as one system drive (host ID/LUN). The host adapter recognizes the system drives after reinitializing the host bus.

Note – The UNIX `format` and Solaris `probe-scsi-all` commands do not display all mapped LUNs if there is not a logical drive mapped to LUN 0.

An FC channel can connect up to 126 devices. Each device has one unique ID.

A SCSI bus channel 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.

FIGURE 5-1 illustrates the idea of mapping a system drive to a host ID/LUN combination.

- The FC or SCSI ID is like a cabinet, and the drawers are the LUNs
- Each cabinet (ID) can have up to 32 drawers (LUNs).
- Data can be stored into one of the LUNs of the FC or SCSI ID. Most FC host adapters treat a LUN like another FC or SCSI device.
- The maximum number of LUNs that can be created for a Sun StorEdge SCSI RAID array is 128. To create a total of 128 LUNs, see [“Planning for 128 LUNs on SCSI Arrays” on page 78](#).
- The maximum number of LUNs that can be created for a Sun StorEdge FC RAID array in loop configurations is 1024, using 4 host channels, 8 host IDs per channel, 8 logical drives, and 128 partitions for each logical drive. To create a total of 1024 LUNs, see [“Planning for 1024 LUNs on FC Arrays” on page 79](#).
- The maximum number of LUNs that can be created for a Sun StorEdge FC RAID array in redundant point-to-point configurations is 64. Refer to *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your Sun StorEdge 3510 FC array to see more information about point-to-point configurations.




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

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

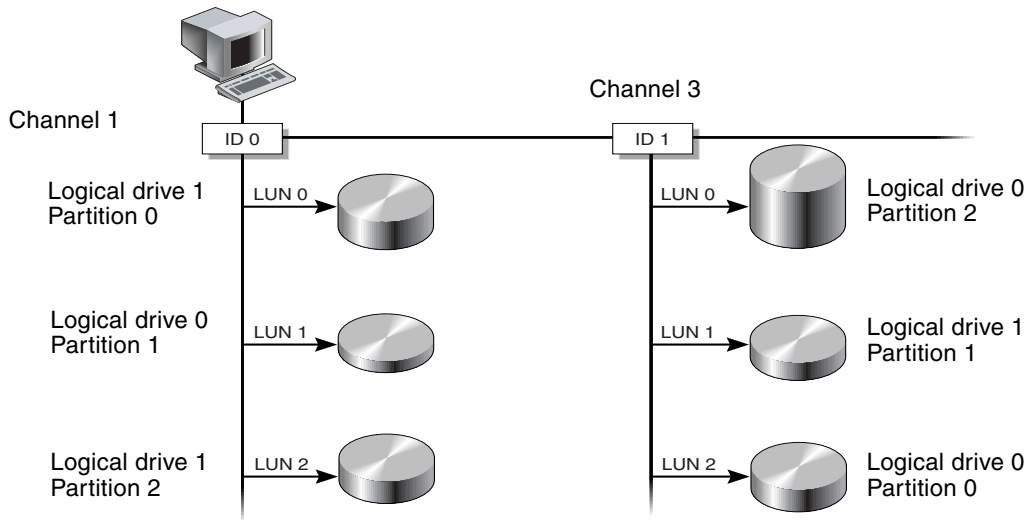


FIGURE 5-2 Mapping Partitions to Host ID/LUNs

Planning for 128 LUNs on SCSI Arrays

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

- Create four host IDs and 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.
- Create six host IDs (this requires three host drives), perform one of the following steps, and then map the 128 partitions to the six 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 135](#).

Note – For an overview of how partitions, LUNs, and host IDs work, refer to [“Mapping Logical Drive Partitions to Host LUNs” on page 75](#).

▼ To Set Up 128 LUNs

1. Create a minimum of four host IDs.

By default, you have two host IDs:

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

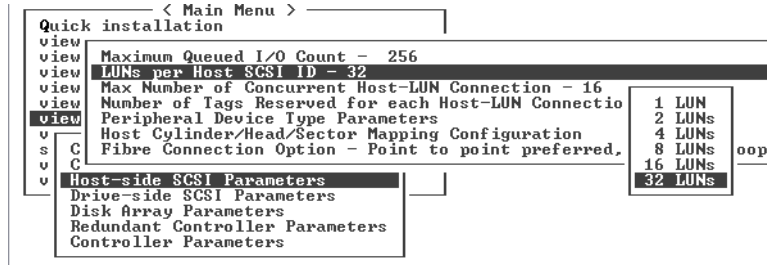
You can have a total of two IDs per channel, one for the primary controller and one for the secondary controller.

For the detailed procedure, refer to [“Creating Additional Host IDs” on page 135](#).

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

- a. Choose **“view and edit Configuration parameters → Host-side SCSI Parameters.”**

- b. If “LUNs per Host SCSI ID” is not set to 32, choose “LUNs per Host SCSI ID -> 32 LUNS.”



3. Create at least four logical drives.

Refer to “Creating Logical Drives” on page 37 for more information.

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

For the detailed procedures, refer to “Partitioning a Logical Drive” on page 44 and “Mapping Logical Drive Partitions to Host LUNs” on page 75.

Planning for 1024 LUNs on FC Arrays

For FC arrays in loop configurations, if you want to create the maximum number of storage partitions that can be mapped for a Fibre Channel RAID array (1024), set up the following configuration:

▼ To Plan for 1024 LUNs on FC Arrays

1. If necessary, edit “Host-side SCSI Parameters” so that “LUNs Per Host SCSI ID” is set to 32.
See “Host-side SCSI Parameters Menu Options” on page 160 for more information.
2. Keep the four default host channels (CH 0, 1, 4, and 5).
3. Create eight host IDs per host channel (four primary controller IDs and four secondary controller IDs per host channel) for a total of 32 host IDs.
For the detailed procedure, see “Creating Additional Host IDs” on page 135.

4. Create eight logical drives.

For the detailed procedure, see [“Creating Logical Drives” on page 37](#).

5. Partition each logical drive into 128 partitions (8 times 128 = 1024).

6. Map the 1024 partitions to the 32 host IDs.

For the detailed procedures, see [“Partitioning a Logical Drive” on page 44](#) and [“Mapping Logical Drive Partitions to Host LUNs” on page 75](#).

TABLE 5-1 Configuration for 1024 LUNs

Configuration Item	Number
Maximum number of logical drives per RAID array	8
Maximum number of partitions per logical drive	128
Maximum number of host channels	4 (Channels 0, 1, 4, 5)
Maximum number of LUNs assigned to each host ID	32
Required number of host IDs per channel	8 (4 PIDs and 4 SIDs)

Planning for 64 LUNs in Redundant FC Point-to-Point Configurations

For FC arrays in point-to-point configurations that use redundant controllers to ensure maximum reliability, accessibility, and serviceability (RAS), you can have a maximum of 64 LUNs. To achieve this redundancy with a maximum number of LUNs, you need to use multipathing software on each host accessing the array.

Note – Multipathing for the Sun StorEdge 3510 FC array is provided by Sun StorEdge Traffic Manager software. Refer to *Sun StorEdge 3510 FC Array Release Notes* for information about which versions of Sun StorEdge Traffic Manager software are supported on which platforms.

For example, to set up 64 LUNs with multipathing enabled, you can map 32 LUNs to channel 0 on one controller and channel 1 on the other controller, and map the other 32 LUNs to channel 4 on one controller and channel 5 on the other controller.

TABLE 5-2 Example of ID Assignments for 64 LUNs With Multipathing Enabled

Channel	Controller Port	PID	SID
0	Top	40	NA
1	Bottom	41	NA
4	Top	NA	50
5	Bottom	NA	51

Refer to *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for more information about point-to-point configurations, including configuring a maximum number of LUNs with redundant controllers and automatic failover. Refer to the *Sun StorEdge 3000 Family Best Practices Manual* for your array to see more information about point-to-point and loop storage area network (SAN) and direct-attached storage (DAS) configurations.

Mapping a Partition to a LUN

Note – The illustrations in this section depict a FC array, but SCSI arrays follow the same procedure.

▼ To Map a Logical Drive Partition to a LUN

1. Choose “view and edit Host luns” from the Main Menu.

A list of available channels and their associated controllers is displayed.

2. Select the desired channel and ID to which the logical drive will be mapped.

- If you see "Logical Drive" and "Logical Volume" menu options, choose "Logical Drive."

The LUN table is displayed.

< Main Menu >						
Quick installation	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit Logical drives	0	LD	0	0	206856	RAID0
view and edit logical Volumes						
view and edit Host Luns	M 1	LD	2	0	198000	RAID0
v CHL 0 ID 40 <Primary Control						
v CHL 1 ID 42 <Secondary Control						
v CHL 4 ID 44 <Primary Control						
s CHL 5 ID 46 <Secondary Control						
v Edit Host-ID/WWN Name List						
v						
v						
v						
v						
v						
v						
v						
v						
v						

- Using the arrow keys, select the desired LUN.

A list of available logical drives is displayed.

< Main Menu >											
Quick installation	LUN	LU/LD	DRU	Partition	Size(MB)	RAID					
view and edit Logical drives	0										
view and edit logical Volumes											
v	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
v	P0	4DB84961	NA	RAID5	4000	GOOD	S	3	0	0	
v											
v	CHL 5 ID 46 <Secondary Control										
v	Edit Host-ID/WWN Name List										
v											
v											
v											
v											
v											
v											

Note – The Sun StorEdge Configuration Service program requires that a device is mapped to LUN 0 as a minimum.

- Select the desired logical drive (LD).

The partition table is displayed.

< Main Menu >											
Quick installation	LUN	LU/LD	DRU	Partition	Size(MB)	RAID					
view and edit Logical drives	0										
view and edit logical Volumes											
v	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
v	P0	4DB84961	NA	RAID5	4000	GOOD	S	3	0	0	
v											
v	CHL	Partition	Offset(MB)	Size(MB)							
v	Edi	0	0	4000							
v											
v											
v											
v											
v											
v											

6. Select the desired partition.

Two mapping options are displayed:

Main Menu >		LUN	LU/LD	DRU	Partition	Size(MB)	RAID
Quick installation		0					
view and edit Logical drives							
view and edit logical Volumes							
view and edit host luns							
u	CHL 0 ID 40 <Primary Controll						
u	CHL 1 ID 42 <Secondary Contro						
u	CHL 4 ID 44 <Primary Controll						
s	CHL 5 ID 46 <Secondary Contro	3					
u	Edit Host-ID/WWN Name List	4					

- “Map Host LUN” is used with all SCSI arrays and with Fibre Channel arrays when multiple hosts are not on the same loop.
 - “Create Host Filter Entry” is used when multiple hosts share the same Fibre Channel loop, can view all the drives, and need to be filtered so that a host only sees the logical drives that are exclusive to its use. For more information, see [“Creating Host Filter Entries \(FC Only\)” on page 87](#).
7. Choose the mapping option that is appropriate for your network and proceed with [“Using the Map Host LUN Option” on page 83](#) or [“Creating Host Filter Entries \(FC Only\)” on page 87](#).

Using the Map Host LUN Option

Each partition must be mapped to a host LUN. Multiple partitions can be mapped to the same host LUN. The “Map Host LUN” menu option is used with all SCSI arrays, and with FC arrays when multiple hosts are not on the same loop.

Note – This procedure is applicable to both SCSI and FC arrays. The illustrations show FC examples.

To map partitions with the host filter entry option, see [“Creating Host Filter Entries \(FC Only\)” on page 87](#).

Note – If you plan to map hundreds of LUNs, it is significantly easier to use the Sun StorEdge Configuration Service program. Refer to *Sun StorEdge 3000 Family Configuration Service User’s Guide* for more information.

▼ To Use the Map Host LUN Option

1. After the completion of the steps in “Mapping a Partition to a LUN” on page 81, choose “Map Host LUN.”



2. Confirm the mapping scheme by choosing Yes.

< Main Menu > Quick installation view and edit Logical drives view and edit Logical Volumes view and edit Host luns		LUN	LU/LD	DRU	Partition	Size<MB>	RAID
u	CHL 0 ID 40 <Primary Control	0					
u	CHL 1 ID 43 <Secondary Contr						
u	CHL 4 ID 44 <Primary Control						
s	CHL 5 ID 46 <Secondary Contr						
v	Edit Host-ID/WWN Name List						

The partition is now mapped to a LUN.

< Main Menu > Quick installation view and edit Logical drives view and edit Logical Volumes view and edit Host luns		LUN	LU/LD	DRU	Partition	Size<MB>	RAID
u	CHL 0 ID 40 <Primary Control	0	LD	0	0	4000	RAID5
u	CHL 1 ID 42 <Secondary Contro	1					
u	CHL 4 ID 44 <Primary Control	2					
s	CHL 5 ID 46 <Secondary Contro	3					
v	Edit Host-ID/WWN Name List	4					
		5					
		6					
		7					

3. Press the Esc key to return to the Main Menu.
4. Repeat Step 1 through Step 3 for each partition until all partitions are mapped to LUNs.
5. To verify unique mapping of each LUN (unique LUN number, unique DRU number, or unique Partition number), choose “view and edit Host luns.”
6. Select the appropriate controller and ID to review the LUN information.

Note – If you are using host-based multipathing software, map each partition to two or more controller/ID combinations in order to make multiple paths to the same LUN available to the host.

Example Host LUN Mappings

The following SCSI array example screens illustrate four channel IDs with eight LUNs per channel:

Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns v v v v s v v		LUN	LV/LD	DRV	Partition	Size(MB)	RAID
		0	LD	0	400	RAID5	
		1	LD	0	400	RAID5	
		2	LD	0	400	RAID5	
		3	LD	0	400	RAID5	
		4	LD	0	400	RAID5	
		5	LD	0	400	RAID5	
		6	LD	0	400	RAID5	
		7	LD	0	400	RAID5	

Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns v v v v s v v		LUN	LV/LD	DRV	Partition	Size(MB)	RAID
		0	LD	1	300	RAID5	
		1	LD	1	300	RAID5	
		2	LD	1	300	RAID5	
		3	LD	1	300	RAID5	
		4	LD	1	300	RAID5	
		5	LD	1	300	RAID5	
		6	LD	1	300	RAID5	
		7	LD	1	300	RAID5	

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
v
v CHL 1 ID 0 (Primary Controll
v CHL 1 ID 1 (Secondary Contro
v CHL 3 ID 2 (Secondary Contro
s CHL 3 ID 3 (Primary Controll
v Edit Host-ID/WWN Name List
v

```

LUN	LV/LD	DRV	Partition	Size(MB)	RAID
0	LD	3	0	350	RAID5
1	LD	3	1	350	RAID5
2	LD	3	2	350	RAID5
3	LD	3	3	350	RAID5
4	LD	3	4	350	RAID5
5	LD	3	5	350	RAID5
6	LD	3	6	350	RAID5
7	LD	3	7	350	RAID5

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
v
v CHL 1 ID 0 (Primary Controll
v CHL 1 ID 1 (Secondary Contro
v CHL 3 ID 2 (Secondary Contro
s CHL 3 ID 3 (Primary Controll
v Edit Host-ID/WWN Name List
v

```

LUN	LV/LD	DRV	Partition	Size(MB)	RAID
0	LD	2	0	400	RAID5
1	LD	2	1	400	RAID5
2	LD	2	2	400	RAID5
3	LD	2	3	400	RAID5
4	LD	2	4	400	RAID5
5	LD	2	5	400	RAID5
6	LD	2	6	400	RAID5
7	LD	2	7	400	RAID5

Deleting a Host LUN Mapping

This section provides the procedure for deleting a host LUN mapping

▼ To Delete a Host LUN Mapping

1. Choose “view and edit Host luns” from the Main Menu.
2. Highlight the channel and ID that are mapped to the host LUN.
3. Highlight the specific host LUN.
4. Choose Yes to delete the host LUN.

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

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

Creating Host Filter Entries (FC Only)

For multiple servers connected to the same Sun StorEdge 3510 FC array, LUN filtering organizes how the array devices are accessed and viewed from host devices. LUN filtering is used to provide exclusive access from a server to a logical drive and exclude other servers from seeing or accessing the same logical drive.

LUN filtering also enables multiple logical drives or partitions to be mapped to the same LUN number, allowing different servers to have their own LUN 0. LUN filtering is valuable in clarifying mapping when each HBA typically sees twice the number of logical drives when viewed through a hub.




FIGURE 5-3 Example of LUN Filtering

An advantage of LUN filtering is that it allows many hosts to attach to an array through a common Fibre Channel port and still maintain LUN security.

Each Fibre Channel device is assigned a unique identifier called a world wide name (WWN). A WWN is assigned by the IEEE and stays with the device for its lifetime. LUN filtering uses the WWN to specify which server is to have exclusive use of a specific logical drive.

As shown in [FIGURE 5-3](#), when you map LUN 01 to host channel 0 and select WWN1, server A has a proprietary path to that logical drive. All servers will continue to see and access LUN 02 and LUN 03 unless filters are created on them.

Determining the Host World Wide Name

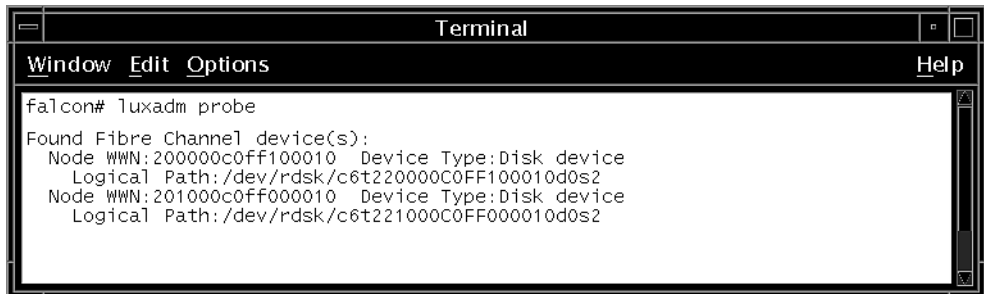
Prior to using LUN filtering, you need to identify which Sun StorEdge 3510 FC array is connected to which HBA card, and the WWN assigned to each card.

▼ To Determine the WWN in the Solaris Operating Environment

1. If you have a new HBA device installed on your computer, reboot the computer.
2. Type the following command:

```
# luxadm probe
```

3. Scroll down the listing to see the Fibre Channel devices and the related WWNs.



```
Terminal
Window Edit Options Help
falcon# luxadm probe
Found Fibre Channel device(s):
  Node WWN:200000c0ff100010 Device Type:Disk device
  Logical Path:/dev/rdisk/c6t220000C0FF100010d0s2
  Node WWN:201000c0ff000010 Device Type:Disk device
  Logical Path:/dev/rdisk/c6t221000C0FF000010d0s2
```

▼ To Determine the WWN for Linux, Windows NT, or Windows 2000

1. **Boot a specific host, and note the BIOS version and HBA card models connected to your host.**
2. **Access the HBA card's BIOS with the appropriate command (Alt-Q or Ctrl-A are commonly used).**

If the host has multiple HBA cards, select the card that is connected to the storage.

3. **Scan the card to look for devices attached to it (usually with the Scan Fibre Devices or the Fibre Disk Utility).**

The node name (or similar label) is the WWN. The following example shows the node name for a Qlogic card.

ID	Vendor	Product	Rev	Node Name	Port ID
0	Qlogic	QLA22xx Adapter	B	210000E08B02DE2F	0000EF

▼ To Determine the WWN for Other Supported Operating Systems

Refer to appendixes in the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the Sun StorEdge 3510 FC Array to see information about determining the WWN for servers running HP-UX or IBM AIX, and for more detailed information regarding WWNs on all supported platforms.

Creating Host Filter Entries

Host filter entries enable you to assign exclusive access to a logical drive to a specific host. Otherwise, in some configurations, multiple hosts would be able to view and access the same drive.

When you create a host filter entry you are also mapping each partition to a specific host, so the procedure described in this section is used instead of the procedure described in [“Using the Map Host LUN Option” on page 83](#).

Note – If you plan to map hundreds of LUNs, the process is easier if you use the Sun StorEdge Configuration Service program.

▼ To Create Host Filter Entries

1. After the completion of the steps in the procedure “[Mapping a Partition to a LUN](#)” on page 81, choose “Create Host Filter Entry → Add from current device list.”

< Main Menu >		LUN	LU/LD	DRU	Partition	Size(MB)	RAID
Quick installation		0	LD	0	0	150000	RAID5
view and edit Logical drives							
view and edit logical Volumes							
view and edit Host luns		1					
v CHL 0 ID 40 <Primary Contro		Map Host LUN					
v CHL 1 ID 42 <Secondary Cont		Create Host Filter Entry					
v CHL 4 ID 44 <Primary Contro		Add from current device list					
v CHL 5 ID 46 <Secondary Cont		Manual add host filter entry					
v Edit Host-ID/WWN Name List		4					
		5					
		6					
		7					

Note – This step automatically performs a discovery of the attached HBAs. Alternatively, you can add them manually by choosing “Manual add host filter entry” rather than “Add from current device list” and typing the Host-ID/WWN in the text area provided.

2. From the device list, select the server WWN number for which you are creating a filter.

< Main Menu >		LUN	LU/LD	DRU	Partition	Size(MB)	RAID
Quick installation		0	LD	0	0	150000	RAID5
view and edit Logical drives							
view and edit logical Volumes							
view and edit Host luns		1					
v CHL 0 ID 40 <Primary Contro		Map Host LUN					
v CHL 1 ID 42 <Secondary Cont		Create Host Filter Entry					
v CHL 4 ID 44 <Primary Contro		Host-ID/WWN					
v CHL 5 ID 46 <Secondary Cont		Host-ID/WWN:0x0000000000323542					
v Edit Host-ID/WWN Name List		5					
		6					
		7					

3. At the confirmation screen, choose Yes.

LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					

< Main Menu > Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns u u CHL 0 ID 40 <Primary Contro u CHL 1 ID 42 <Secondary Conto u CHL 4 ID 44 <Primary Contro s CHL 5 ID 46 <Secondary Cont v Edit Host-ID/WWN Name List	Map Host LUN Create Host Filter Entry
Host-ID/WWN: 0x0000000000323542 Yes No	

LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					
6					
7					

4. Review the filter configuration screen. Make any changes necessary by using the arrow keys to select an item.

LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					

< Main Menu > Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns u u CHL 0 ID 40 <Primary Contro u CHL 1 ID 42 <Secondary Conto u CHL 4 ID 44 <Primary Contro s CHL 5 ID 46 <Secondary Cont v Edit Host-ID/WWN Name List	Map Host LUN Create Host Filter Entry				
4	Logical Drive 0 Partition 0 Host-ID/WWN - 0x0000000000323542 Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFFF Filter Type - Include Access Mode - Read/Write Name - Not Set				
5					
6					
7					

Note – If you press Return to change the first item displayed, which is the logical drive and partition assignment, press Esc and repeat the “Create Host Filter Entry” procedure.

5. If you want to edit the WWN, choose “Host-ID/WWN.”

6. Make the desired changes and press Return.

LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					

< Main Menu > Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns u u CHL 0 ID 40 <Primary Contro u CHL 1 ID 42 <Secondary Conto u CHL 4 ID 44 <Primary Contro s CHL 5 ID 46 <Secondary Cont v Edit Host-ID/WWN Name List	Map Host LUN Create Host Filter Entry				
4	Logical Drive 0 Partition 0 Host-ID/WWN - 0x0000000000323542 H F A N Host-ID/WWN:				
5					
6					
7					



Caution – Be sure that you edit the WWN correctly. If the WWN is incorrect, the host will be unable to recognize the LUN.

- To edit the WWN Mask, use the arrow keys to select “Host-ID/WWN Mask” and press Return.

< Main Menu >						
Quick installation	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit Logical drives	0	LD	0	0	150000	RAID5
view and edit logical Volumes	1					
view and edit Host luns						
v						
v CHL 0 ID 40 <Primary Contro	Map Host LUN Create Host Filter Entry					
v CHL 1 ID 42 <Secondary Cont						
v CHL 4 ID 44 <Primary Contro						
s CHL 5 ID 46 <Secondary Cont						
v Edit Host-ID/WWN Name List						
	4				Logical Drive 0 Partition 0 Host-ID/WWN - 0x000000000323542 Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFF	
	5				Host-ID/WWN Mask: _	
	6					
	7					

- If you want to change the filter setting, which toggles between “Include” and “Exclude,” choose “Filter Type -.”

- At the confirmation screen, choose Yes.

< Main Menu >						
Quick installation	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit Logical drives	0	LD	0	0	150000	RAID5
view and edit logical Volumes	1					
view and edit Host luns						
v						
v CHL 0 ID 40 <Primary Contro	Map Host LUN Create Host Filter Entry					
v CHL 1 ID 42 <Secondary Cont						
v CHL 4 ID 44 <Primary Contro						
s CHL 5 ID 46 <Secondary Cont						
v Edit Host-ID/WWN Name List						
	4				Logical Drive 0 Partition 0 Host-ID/WWN - 0x000000000323542 Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFF Filter Type - include	
	5					
	6				Set Filter Type to Exclude ? Y N	
	7				<input checked="" type="checkbox"/> Yes No	

- If you want to change the access mode, which toggles between “Read/Write” and “Read Only,” use the arrow keys to choose “Access mode -.”

11. At the confirmation screen, choose Yes.

< Main Menu > Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns					
0	LD	0	0	150000	RAID5
1					
4					
5					
6					
7					

CHL 0 ID 40 <Primary Contro CHL 1 ID 42 <Secondary Cont CHL 4 ID 44 <Primary Contro CHL 5 ID 46 <Secondary Cont Edit Host-ID/WWN Name List	Map Host LUN Create Host Filter Entry
--	--

L	Set Access Mode to Read-Only ?	
H	Yes	No
H		
F		F

4					
5					
6					
7					

12. If you want to set a name, choose "Name -."

13. Type the name you want to use for this filter and press Return.

< Main Menu > Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns					
0	LD	0	0	150000	RAID5
1					
4					
5					
6					
7					

CHL 0 ID 40 <Primary Contro CHL 1 ID 42 <Secondary Cont CHL 4 ID 44 <Primary Contro CHL 5 ID 46 <Secondary Cont Edit Host-ID/WWN Name List	Map Host LUN Create Host Filter Entry
--	--

L	Name :mars	542
H		FFFFFF
H		
F		
A		

4					
5					
6					
7					

14. Verify all settings and press Esc to continue.

< Main Menu > Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns					
0	LD	0	0	150000	RAID5
1					
4					
5					
6					
7					

CHL 0 ID 40 <Primary Contro CHL 1 ID 42 <Secondary Cont CHL 4 ID 44 <Primary Contro CHL 5 ID 46 <Secondary Cont Edit Host-ID/WWN Name List	Map Host LUN Create Host Filter Entry
--	--

4	Logical Drive 0 Partition 0	
5	Host-ID/WWN - 0x0000000000323542	
6	Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFF	
7	Filter Type - Exclude	
	Access Mode - Read/Write	
	Name - mars	

15. At the confirmation screen, choose Yes.

< Main Menu >						
LUN	LU/LD	DRU	Partition	Size(MB)	RAID	
0	LD	0	0	150000	RAID5	
1						
Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns						
CHL 0 ID 40 <Primary Contro CHL 1 ID 42 <Secondary Cont CHL 4 ID 44 <Primary Contro CHL 5 ID 46 <Secondary Cont Edit Host-ID/WWN Name List						
Map Host LUN Create Host Filter Entry						
Add Host Filter Entry ? <input checked="" type="radio"/> Yes <input type="radio"/> No						
4						
5						
6						
7						

16. At the Host-ID/WWN list, repeat the previous steps to create additional filters, or press Esc to continue.

< Main Menu >						
LUN	LU/LD	DRU	Partition	Size(MB)	RAID	
0	LD	0	0	150000	RAID5	
1						
Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns						
CHL 0 ID 40 <Primary Contro CHL 1 ID 42 <Secondary Cont CHL 4 ID 44 <Primary Contro CHL 5 ID 46 <Secondary Cont Edit Host-ID/WWN Name List						
Map Host LUN Create Host Filter Entry						
Host-ID/WWN - 0x0000000000323542						
5						
6						
7						

Note – Unlike most firmware operations where you must complete each entry individually and repeat the procedure if you want to perform a similar operation, you can add multiple WWNs to your list before you actually create the host filter entry in [Step 17](#). Be sure to follow the instructions carefully.

17. At the confirmation screen, verify the settings and choose Yes to complete the host LUN filter entry.

< Main Menu >						
LUN	LU/LD	DRU	Partition	Size(MB)	RAID	
0	LD	0	0	150000	RAID5	
1						
Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns						
CHL 0 ID 40 <Primary Contro CHL 1 ID 42 <Secondary Cont CHL 4 ID 44 <Primary Contro CHL 5 ID 46 <Secondary Cont Edit Host-ID/WWN Name List						
Map Logical Drive: 0 Partition : 0 To Channel : 0 ID : 112 Lun : 1 ? <input checked="" type="radio"/> Yes <input type="radio"/> No						
6						
7						

A mapped LUN displays a number and a filtered LUN displays an “M” for masked LUN in the host LUN partition window.

<pre> < Main Menu > Quick installation view and edit Logical drives view and edit Logical Volumes view and edit Host luns </pre>		LUN	LU/LD	DRU	Partition	Size(MB)	RAID
		0	LD	0	0	150000	RAID5
		M 1	LD	0	0	150000	RAID5
	<pre> v CHL 0 ID 40 <Primary Contro v CHL 1 ID 42 <Secondary Cont v CHL 4 ID 44 <Primary Contro v CHL 5 ID 46 <Secondary Cont v Edit Host-ID/WWN Name List v </pre>	2					
		3					
		4					
		5					
		6					
		7					

Viewing and Modifying Host Filter Information

Once you have created host filter entries you can display detailed information about them.

Note – If you have host filter entries for two or more WWNs, an asterisk next to an entry indicates that more information is available than is currently displayed. To display the additional information, select the entry and press Enter.

▼ To View or Modify Host Filter Information

1. Choose “view and edit Host luns.”
2. Select the channel and ID that are mapped to the host LUN.
3. Select the filtered LUN.
4. Choose “View and edit host filtering.”
5. Select the Host-ID/WWN whose information you want to view or edit.
6. Choose “View Host Filter Information” to see detailed information about that filter.
7. Choose “Add Host Filter Entry” to add another filter. See [“To Create Host Filter Entries” on page 90](#) for the steps to follow once you choose this menu option.

8. Choose "Delete Filter Entry" to delete the current filter.
9. Choose "Add Host-ID/WWN Name List" to manually add a WWN. See ["To Create Host Filter Entries" on page 90](#) for the steps to follow once you choose this menu option.

Viewing and Editing SCSI Drives

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

- “Viewing the Physical Drive Status Table” on page 98
- “SCSI Drive IDs (SCSI Only)” on page 100
- “FC Drive IDs (Fibre Channel Only)” on page 102
- “Assigning a Local Spare Drive” on page 102
- “Assigning a Global Spare” on page 103
- “Viewing Drive Information” on page 104
- “Viewing the Connected Drives” on page 105
- “Scanning Drives” on page 106
- “Deleting a Spare Drive” on page 107
- “Setting Slot Numbers” on page 108
- “Adding or Deleting Drive Entries” on page 110
- “Fault Protection Measures” on page 111
 - “Cloning a Failing Drive” on page 111
 - “Replacing After Clone” on page 112
 - “To Replace After Clone” on page 112
 - “Perpetual Clone” on page 114
 - “To Enable Perpetual Clone” on page 114
 - “Viewing the Status of a Cloning Operation” on page 116
 - “To View the Status of a Cloning Operation” on page 116
 - “Understanding SMART Technology” on page 117
 - “How Sun StorEdge 3000 Family Arrays Use SMART” on page 117
 - “Enabling SMART From Firmware Menus” on page 118
 - “To Enable SMART” on page 118
 - “Detect Only” on page 120
 - “To Enable Detect Only” on page 120
 - “Detect and Perpetual Clone” on page 120
 - “To Detect and Perpetual Clone” on page 120
 - “Detect and Clone+Replace” on page 121
 - “To Detect and Clone+Replace” on page 121
- “Identifying a Failed Drive for Replacement” on page 121
- “SCSI Drive Utilities (Reserved)” on page 125

TABLE 6-1 Parameters Displayed in the Drive Status Window

Parameters	Description
Slot	Slot number of the SCSI drive.
Chl	Channel of the connected drive.
ID	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 drive is a drive member of logical drive <i>x</i> . If Status shows "STAND-BY," the drive is a local spare drive of logical drive <i>x</i> .
Status	GLOBAL The drive is a global spare drive. INITING The drive is initializing. ON-LINE The drive is in good condition. REBUILD The drive is rebuilding. 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 part of a logical drive from which it has been removed; it still contains data from this logical drive. FRMT DRV The drive has been formatted with reserved space allocated for controller-specific information. BAD Failed drive. ABSENT Drive slot is not occupied. MISSING Drive once existed, but is now missing. SB-MISS Spare drive missing. SES A SES device that monitors the array's environmental conditions.
Vendor and Product ID	Vendor and product model information of the drive.

A physical drive has a USED status when it was once part of a logical drive but no longer is. This can happen, for instance, when a drive in a RAID 5 array is replaced by a spare drive and the logical drive is rebuilt with the new drive. If the removed drive is later replaced in the array and scanned, the drive status is identified as USED since the drive still has data on it from a logical drive.

When the RAID set is deleted properly, this information is erased and the drive status is shown as FRMT rather than USED. A drive with FRMT status has been formatted with either 64 KB or 256 MB of reserved space for storing controller-specific information, but has no user data on it.

If you remove the reserved space, using the “View and Edit SCSI drives” menu, the drive status changes to NEW.

To handle BAD drives, see [“Controller Failover” on page 185](#). If two drives show BAD and MISSING status, see [“Recovering From Fatal Drive Failure” on page 190](#).


SCSI Drive IDs (SCSI Only)

Each SCSI 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 *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array.


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 six disk drive IDs to CH 0 and six disk drive IDs to CH 2 in the RAID array, and then typically adds an additional six 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



FC Drive IDs (Fibre Channel Only)

When an expansion unit is attached to a RAID array, unique, hard-assigned loop IDs are assigned to each expansion unit drive. A loop ID is the decimal version of an AL_PA. The lowest number loop ID is the lowest priority address on the loop.

On the left front side of an expansion unit, an ID switch is used to set the loop IDs for the disk drives to a different range of values so that the same IDs are not repeated on the same loop. Press the button to change the ID number.

By default, the ID switch on all RAID arrays and expansion units is set to 0, where the range of IDs is automatically 0 to 11 for 12 drives (the IDs 12–15 are ignored).

The ID switch offers eight ID ranges. Each set contains 16 IDs (the last four IDs in each range are ignored). These ranges are shown in [TABLE 6-2](#).

TABLE 6-2 ID Switch Settings for FC Expansion Units

ID Switch Setting	Range of IDs
0	0–15
1	16–31
2	32–47
3	48–63
4	64–79
5	80–95
6	96–111
7	112–125

Assigning a Local Spare Drive

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

Note – Logical drives configured with a nonredundant RAID level (RAID 0) do not support spare drive rebuild.

▼ To Assign a Local Spare Drive

1. Choose “view and edit scsi Drives.”
2. Select the drive you want to designate as a spare.
3. Choose “add Local spare drive.”
4. Select a logical drive to assign the local spare to that logical drive.
A confirmation prompt is displayed,
5. Choose Yes to confirm.

Assigning a Global Spare

A global spare drive automatically replaces any failed drive in the array.

▼ To Assign a Global Spare

1. Choose “view and edit scsi Drives.”
2. Select the drive you want to designate as a spare
3. Choose “add Global spare drive.”
A confirmation prompt is displayed.

Quick view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
view	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		Add Global Spare Drive ?					T DRV	SEAGATE ST336605LSUN36G
view		Yes No					T DRV	SEAGATE ST336605LSUN36G
view		disk Reserved space - unformatted					T DRV	SEAGATE ST336605LSUN36G
view		0	13	34732	160MB	NONE	NEW DRV	SEAGATE ST336605LSUN36G
view		0	14				SAF-TE	DotHill ERMLVD SCRUNCH

4. Choose Yes to confirm.

Viewing Drive Information

Displayed SCSI drive information includes the revision number, serial number, and disk capacity of individual drives. If you select an SES chip rather than a drive, the “View drive information” menu option also displays the Node name (WWN). From this information you can determine the serial number of the chassis, as described in the following procedure.

▼ To View SCSI Drive Information

1. From the Main Menu, choose “view and edit scsi Drives.”
2. Select the SCSI drive you want to view.
3. Choose “View drive information.”

The drive’s revision number, serial number and disk capacity (in blocks of 512K) of the drive are displayed.

The screenshot shows the Main Menu with the following options:

- Quick installation
- view and edit Logical drives
- view and edit logical Volumes
- view and edit Host luns
- view and edit scsi Drives
- View drive information
- add Local spare drive
- add Global spare drive
- Delete global/local spare drive
- Scan scsi drive
- set slot Number
- add drive Entry
- Clear drive status
- Identifying scsi drive
- clone Failing drive
- scsi drive Utilities

The View drive information screen displays the following information:

Status	Vendor and Product ID
0 ON-LINE	HITACHI DK31CJ-72FC
0 ON-LINE	HITACHI DK31CJ-72FC

The Revision Number, Serial Number, Disk Capacity <blocks>, Fibre Port Name, and Redundant Loop ID are also displayed:

Revision Number	G7A7
Serial Number	ID009863
Disk Capacity <blocks>	0144410879
Fibre Port Name	500E10015B73
Redundant Loop ID	16

If you selected the SES chip rather than a drive, the Node name (WWN) text area displays the worldwide name assigned to the enclosure.

< Main Menu >								
Quick installation								
view and edit Logical drives								
view and edit logical Volumes								
view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
view								
view	Revision Number					1000	36753FSUN36G	
view	Serial Number						36753FSUN36G	
view	Disk Capacity (blocks)					N/A	36753FSUN36G	
view	Node Name(WWNN)					20 40 00 C0 FF 00 2F 18	36753FSUN36G	
view	Redundant Loop ID					12	36753FSUN36G	
	View drive information					0	ON-LINE	SEAGATE ST336753FSUN36G
	Scan scsi drive					GLOBAL	STAND-BY	SEAGATE ST336753FSUN36G
	add drive Entry							
	2(3)	12				SES	SUN	StorEdge 3510F A

In the example above, the world wide name is shown in hexadecimal format:

Node Name(WWNN) 20 40 00 C0 FF 00 2F 18

The last six hexadecimal digits of the WWN indicate the serial number of the chassis, which is the same as the chassis Field-Replaceable Unit Identifier (FRU-ID). This number is sometimes shown as only the last four hexadecimal digits. In the example above, the FRU-ID is 002F18, or simply 2F18.

Viewing the Connected Drives

Prior to configuring disk drives into a logical drive, you must know the status of physical drives in your enclosure.

▼ To View the List of Available SCSI Drives

1. Choose “view and edit scsi Drives.”

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
view								

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

If a drive is installed but not listed, the drive might be defective or installed incorrectly.

3. If a physical drive was connected after the controller completed initialization, select the drive from the table and press Return. Use the “Scan scsi drive” menu option to recognize the newly added physical drive and to configure it as a member of a logical drive.



Caution – Scanning an existing drive removes its assignment to any logical drive. All data on that drive is lost.

Scanning Drives

When an array is powered on, the controller scans all physical drives that are connected through drive channels. If a physical drive is connected after the controller completes initialization, use the “Scan scsi drive” menu option

▼ To Scan a New SCSI Drive

1. Choose “view and edit scsi Drives” from the Main Menu.
2. Select a drive from the SCSI drive table.



Caution – Scanning an existing drive removes its assignment to any logical drive. All data on that drive is lost.

3. Choose "Scan scsi drive."

SCSI channels are displayed.

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

The menu options vary according to the drive status.

4. Select a SCSI channel.

A menu displays drive IDs.

5. Select the SCSI ID of the drive you want to scan.

A confirmation prompt is displayed.

6. Choose Yes to confirm.

Deleting a Spare Drive

This section describes how to delete a local spare or global spare drive.

▼ To Delete a Spare Drive

1. Choose "view and edit scsi Drives" from the Main Menu.

2. Select a local spare drive or a global spare drive.




FIGURE 6-1 Set Slot Number

4. Type in a value (0–15) that represents the slot number of the drive.

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.

Assigning a Slot Number to an Empty Slot

When there is an empty slot (or air management sled) that does not contain a drive, the corresponding SCSI channel/ID does 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.

Deleting a Slot Number

This section describes how to delete a slot number of a SCSI drive.

▼ To Delete the Slot Number of a SCSI Drive

1. Choose “view and edit scsi Drives” from the Main Menu.
2. Select the desired SCSI drive.
3. Choose “Set Slot Number → 0.”

Adding or Deleting Drive Entries

Use the “add drive Entry” menu option to add an additional record to the SCSI drive table prior to adding a SCSI drive.

Use the “Clear drive status” menu option if you want to remove an empty drive designation from the table later.




FIGURE 6-2 Add Drive Entry

▼ To Add a Drive Entry

1. Choose “view and edit scsi Drives” from the Main Menu.
2. Select an insertion spot within the SCSI drive table.
3. Choose “add drive Entry.”
A channel list is displayed.
4. Select a channel.
For installed SCSI drives, a table shows the available IDs.
5. Type the desired ID number in the provided text area and press Return.
A confirmation prompt is displayed.
6. Choose Yes to confirm.
A drive entry is created. Its status is displayed as ABSENT.

▼ To Remove an Empty Drive Entry

1. On the Main Menu, choose “view and edit scsi Drives.”
2. Select the desired SCSI drive.
3. Choose “set slot Number → 0 → Clear drive status.”

The empty drive entry is removed from the SCSI drive table.

Note – You cannot remove an empty drive entry if it has been assigned a slot number. Delete the slot number before removing the empty drive entry.

Fault Protection Measures

With the maturity of industry-standard technologies such as Self-Monitoring, Analysis and Reporting Technology (SMART), disk drive failures can sometimes be predicted before they happen. Encountering drive bad block reassignments is one common predictor of a drive that is about to fail.

System administrators can decide when to substitute a healthy drive for a drive showing symptoms of impending failure. This section discusses manual and automated procedures for averting disk failures.

Cloning a Failing Drive

To assist fault prevention, a system administrator can manually clone a disk drive that shows signs of failing, choosing a convenient time when system performance will not be adversely affected.

Note – The “Clone Failing Drive” menu option is not displayed with RAID 1 logical drives because its use is not supported in RAID 1 configurations.

Use the Clone Failing Drive menu option when:

- Replacing drives about to fail (indicated by the 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

These options are described in the following sections.

Replacing After Clone

Data on the source drive (the drive with the predicted error or any selected member drive) is cloned to a standby spare. The spare then becomes the new source drive. The status of the original source drive is redefined as a “USED DRIVE.” System administrators can 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 must add a new drive and configure it as a standby drive. If there is no standby drive, the “clone failing drive” option is not displayed.

▼ To Replace After Clone

1. Choose “view and edit scsi Drives.”
2. Select the member drive that you want to clone.
3. Choose “clone failing drive.”

This option is displayed only if there is a standby drive available.

4. Choose “Replace After Clone.”

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

	Slot	Ch1	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
Quick view		2	0	319	20MB	0	ON-LINE	
view						0	ON-LINE	
view						0	ON-LINE	
view						0	STAND-BY	
view						NONE	NEW DRV	
view							EW DRV	
view		2					EW DRV	
view		2					EW DRV	

5. A notification message is displayed.

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

6. 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="width: 28%; background-color: gray; height: 10px; margin-bottom: 5px;"></div> 28% Completed </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.

7. Select the drive indicated as "CLONING."

Note – To quit the status bar, press Esc to return to the table of SCSI drives.

Quick view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
view		2	0	319	20MB	0	ON-LINE	
view		2	1	319	20MB	0	ON-LINE	
view		2	2	319	20MB	0	ON-LINE	
view		2	3	319	20MB	0	CLONING	
view							<div style="border: 1px solid black; padding: 5px;"> V S t a t u s Source Drive: Channel 2 ID 0 View clone progress Abort clone Clone Failing drive </div>	
view							EW DRV	
view							EW DRV	
view							EW DRV	
view						NONE	NEW DRV	

8. Choose "clone Failing drive" again to view the current status.

Note – You can identify the source drive and choose "View clone progress," or choose "Abort clone" if you selected the wrong drive.

When the process is completed, the following message is displayed.

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

9. Press Esc to proceed.

5. Press Esc to view current progress on a status bar.

Quick view view view view view view view view view view	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
		2	0	319	20MB	0	ON-LINE	
		<div style="border: 1px solid black; padding: 5px; margin: 5px auto; width: 80%;"> <div style="background-color: #ccc; width: 35%;"></div> <p style="text-align: center;">35% Completed_</p> </div>						
		2	4	319	20MB	NONE	NEW DRV	
		2	5	319	20MB	NONE	NEW DRV	

6. To quit viewing the status bar, press Esc to return to the previous menu screen.

7. Select the drive indicated as "CLONING" by pressing Return.

8. Choose "clone Failing drive" again to view the progress.

Note – You can identify the source drive and choose "View clone progress." Choose "Abort clone" if you have selected the wrong drive.

Quick view view view view view view view view view view	Slot	Chl	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	
		<div style="border: 1px solid black; padding: 5px; margin: 5px auto; width: 80%;"> <div style="display: flex; justify-content: space-between;"> V EW DRV </div> <div style="display: flex; justify-content: space-between;"> S EW DRV </div> <div style="display: flex; justify-content: space-between;"> a EW DRV </div> <div style="display: flex; justify-content: space-between;"> t NEW DRV </div> </div>						
						NONE	NEW DRV	

A notification message informs you when the process is complete.

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

9. Press Esc to clear the notification message and to see the SCSI drives' status after the cloning process.

The source drive (Channel 1 ID 5) remains as a member of logical drive 0, and the standby drive (Channel 1 ID 2, the local or global spare drive) becomes a CLONE drive.

	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
Quick view		2	0	319	20MB	0	ON-LINE	
view		2	1	319	20MB	0	ON-LINE	
view		2	2	319	20MB	0	ON-LINE	
view		2	3	319	20MB	0	CLONE	
view		2	4	319	20MB	NONE	NEW DRV	
view		2	5	319	20MB	NONE	NEW DRV	
view		2	6	319	20MB	NONE	NEW DRV	
view		2	8	319	20MB	NONE	NEW DRV	

Viewing the Status of a Cloning Operation

While a cloning operation is underway you can examine its status, including the operation's progress and the identity of the target drive.

▼ To View the Status of a Cloning Operation

1. Choose "view and edit Scsi drives" from the Main Menu.
The SCSI drive table is displayed.
2. Select the target drive whose status is CLONING.
3. Choose "clone failing drive" to identify the source drive that is being cloned and to see a menu whose options enable you to display the progress of the cloning operation or abort the cloning operation.

Note – If you are viewing an active monitoring session with Sun StorEdge Configuration Service software, the progress of the cloning operation is displayed by the Controller Array Progress bar.

Understanding SMART Technology

SMART is an industry-standard technology that provides near-term failure prediction for disk drives. When SMART is enabled, the drive monitors predetermined drive attributes that are susceptible to degradation over time. If a failure is likely to occur, SMART makes a status report available so that the host can prompt the user to back up data on the failing drive.

Not all failures can be predicted, however. SMART predictability is limited to the attributes the drive can monitor that are selected by the device manufacturer, based on the attribute's ability to contribute to the prediction of degrading or fault conditions.

Although SMART attributes are drive-specific, a variety of typical characteristics can be identified:

- Head flying height
- Data throughput performance
- Spin-up time
- Reallocated sector count
- Seek error rate
- Seek time performance
- Spin try recount
- Drive calibration retry count

How Sun StorEdge 3000 Family Arrays Use SMART

Sun StorEdge 3000 family arrays implement the ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190 standard.

Sun StorEdge 3000 family array firmware supports four manual selections related to SMART function in firmware:

- **Disable.** SMART functions are not activated.
- **Detect Only.** SMART functions are enabled. The controller sends a command to enable all the drives' SMART functions. If a drive predicts a problem, the controller reports the predicted problem as an entry in the event log.
- **Perpetual Clone.** The controller sends a command to enable all the drives' SMART functions. If a drive predicts a problem, the controller reports the predicted problem as an entry in the event log. The controller then clones the drive whose failure has been predicted if a global or local spare drive is available. The clone drive still functions as a standby drive.

If the drive whose failure has been predicted does fail subsequently, the clone drive takes over immediately.

Note – If the drive whose failure has been predicted continues to work successfully and another drive in the same logical drive fails, the clone drive performs as a standby spare drive and start to rebuild the failed drive immediately. This helps prevent a fatal drive error if yet another drive fails.

- **Clone + Replace.** The controller sends a command to enable all the drives' SMART functions. If a drive predicts a problem, the controller reports the predicted problem as an entry in the event log. The controller then clones the drive whose failure has been predicted to a standby spare drive and takes the drive whose failure has been predicted offline as soon as the cloning process is completed.

Enabling SMART From Firmware Menus

This section describes how to configure the SMART settings you want.

▼ To Enable SMART

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → Periodic Drive Check Time.”
2. Choose a time interval.



A confirmation message is displayed.

3. Choose Yes to continue.
4. Choose “Drive-side SCSI Parameters → Drive Predictable Failure Mode (SMART).”
5. Choose one of the menu options from the “Drive Predictable Failure Mode (SMART)” menu.
 - “Disable”
 - “Detect Only”
 - “Detect and Perpetual Clone”

Detect Only

This section describes setting SMART to only detect future problems and not attempt to fix them.

▼ To Enable Detect Only

1. Choose “Drive-side SCSI Parameters → Drive Predictable Failure Mode (SMART) → Detect Only.”

Whenever a drive predicts symptoms of predictable drive failure, the controller writes an error message to the event log.

Detect and Perpetual Clone

This section describes setting SMART to detect future problems and clone a spare drive so that it is ready when failure does occur.

▼ To Detect and Perpetual Clone

1. Choose “Drive-side SCSI Parameters → Drive Predictable Failure Mode → Detect and Perpetual Clone.”
2. Assign at least one spare drive to the logical drive (either a local spare or global spare).

When a drive (logical drive member) detects a predictable drive failure, the controller clones the drive to a spare drive.

3. To see the status of the source drive or the cloning progress, or to cancel the cloning process, from the “View and Edit SCSI Drive” menu select the spare drive (either local or global) and choose the appropriate menu option.



Quick view view view view view view view view view view view	Slot	Chl	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	CLONING	

V	S	s	s	s	clone failing drive	NONE	NEW DRV
					Source Drive: Channel 2 ID 0		EW DRV
					View clone progress		EW DRV
					Abort clone		EW DRV

Note – The spare drive stays mirrored to the source drive (the drive whose failure has been predicted) but does not replace it until the source drive fails.

Note – While the spare drive is mirroring the source drive and no other spare drive is available, any drive failure forces the spare drive to give up the mirrored data and resume its original role. It becomes a spare drive again and rebuilds the failed drive.

Detect and Clone+Replace

This section describes setting SMART to detect future problems, clone a spare drive, and immediately use the cloned drive to replace the original drive whose failure is predicted.

▼ To Detect and Clone+Replace

1. Choose “Drive-side SCSI Parameters → Drive Predictable Failure Mode (SMART) → Detect and Clone+Replace.”
2. Assign at least one spare drive (either local or global) to the logical drive.

When a drive failure is predicted, the controller clones that drive to a spare drive. After the clone process is complete, the controller immediately replaces the source drive (the drive whose failure has been predicted). The status of the source drive then is changed to a used drive, and you can replace this drive with a new one.

Note – If you want to see the progress of cloning, press Esc to clear the notification message and see the status bar.

Identifying a Failed Drive for Replacement

If there is a failed drive in a RAID 5 logical drive, replace the failed drive with a new drive to keep the logical drive working.

4. Change the duration if you want and press Return.

A confirmation message is displayed.

Quick view	Slot	Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
view	2<3>		0	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G
view	View drive information						0	ON-LINE SEAGATE ST336752FSUN36G
view	Scan scsi drive						0	ON-LINE SEAGATE ST336752FSUN36G
view	set slot Number						0	ON-LINE SEAGATE ST336752FSUN36G
view	add drive Entry						0	ON-LINE SEAGATE ST336752FSUN36G
syst	Identify scsi drive						0	ON-LINE SEAGATE ST336752FSUN36G
view	Flash All drives						Flash Channel:2 ID:0 SCSI Drive ?	
	Flash Selected drive						Yes No	
							f	
	Flash Drive Time(Second) : 15						E SEAGATE ST336752FSUN36G	
							E SEAGATE ST336752FSUN36G	

5. Choose Yes to confirm.

The read/write LED of a failed hard drive does not light. The absence of a lit LED helps you locate and remove the failed drive.

In addition to flashing all drives, you can flash the read/write LED of only a selected drive or flash the LEDs of all drives except the selected drive, using steps similar to those outlined in this section. These three drive-flashing menu options are described in the remainder of this section.

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 6-3 Flashing the Drive LED of a Selected Drive

Flashing All SCSI Drives

The “Flash All SCSI Drives” menu option lights LEDs of all good drives but does not light LEDs for any defective drives.




FIGURE 6-4 Flashing All Drive LEDs to Detect a Defective Non-Flashing Drive

Flashing All But Selected Drive

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




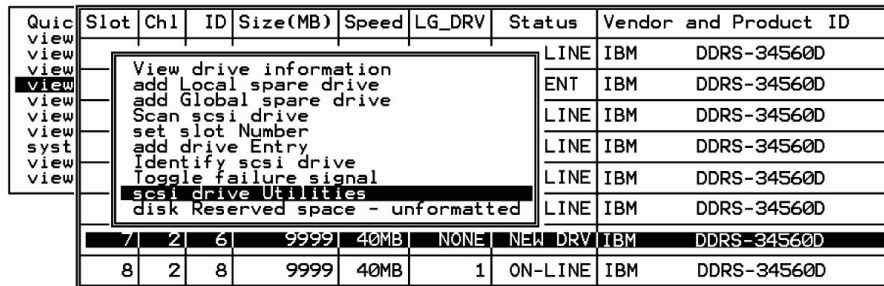
FIGURE 6-5 Flashing All Drive LEDs Except a Selected Drive LED

SCSI Drive Utilities (Reserved)

Do not use the “scsi drive Utilities” menu option. These utilities are reserved for specific troubleshooting methods and should be used only by qualified technicians.

▼ To Use the SCSI Drive Utilities (Reserved) Menu

1. Choose “view and edit scsi Drives” from the Main Menu.
2. Select the drive that the utility is to performed on.
3. Choose “scsi drive Utilities → SCSI Drive Low-level Format” or “scsi drive Utilities → Read/Write Test.”



The screenshot shows a terminal-style interface. On the left, a vertical menu lists options: "Quic", "view", "view", "view", "view", "view", "view", "view", "view", "view", "view". The "view" option is highlighted with a black bar. A large rectangular menu is overlaid on the table, listing the following options: "View drive information", "add Local spare drive", "add Global spare drive", "Scan scsi drive", "set slot Number", "add drive Entry", "Identify scsi drive", "Toggle failure signal", "scsi drive Utilities" (highlighted with a black bar), and "disk Reserved space - unformatted".

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 (Reserved)



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

The SCSI disk drive on which a low-level disk format will be performed cannot be a spare drive (local or global) nor a member drive of a logical drive. “Disk Reserved space” must be removed before this menu option appears. See [“Changing Disk Drive Reserved Space” on page 128](#) for more information.

The “SCSI Drive Low-level Format” menu option appears only if the drive status is a NEW or USED drive.

▼ To Low-Level Format a SCSI Drive (Reserved)

1. Choose “view and edit scsi Drives” from the Main Menu.
2. Select a new or used drive that the utility is to be performed on.
3. Choose “scsi drive Utilities → scsi Drive Low-level Format.”

A confirmation message is displayed.

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

View drive information	LINE	IBM	DDRS-34560D
add Local spare drive	ENT	IBM	DDRS-34560D
Scan			-34560D
add			-34560D
set			-34560D
iden			-34560D
logs			-34560D
scsi			-34560D
d			-34560D

SCSI Drive Low-Level Format		Yes	No
[WARNING!!] Channel 2 ID 6 SCSI Drive			
All data on the disk will be erased !			
!! Low-Level Format Disk !!			

4. Choose Yes.

Note – Do not switch the controller or SCSI disk drive power off during the SCSI Drive Low-level Format. If any power failure occurs during a drive low-level format, the formatting must be performed again when power resumes.

SCSI Drive Read/Write Test (Reserved)

This section describes how to perform a read/write test.

▼ To Perform the SCSI Drive Read/Write Test (Reserved)

1. Choose “view and edit scsi Drives” from the Main Menu.
2. Select a new or used drive on which the utility is to be performed.
3. Choose “scsi drive Utilities → Read/Write Test.”
4. Enable or disable the following options and press Return after each change.
 - Auto Reassign Bad Block
 - Abort When Error Occurs

- Drive Test for - Read Only/Read and Write

5. When configuration is complete, choose “Execute Drive Testing.”

Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	2	0	9999	40MB	0	ON-LINE	IBM DDRS-34560D
2	2	1	9999	40MB	0	ON-LINE	IBM DDRS-34560D
3	2	2	9999	40MB	0	ON-LINE	IBM DDRS-34560D
4							IBM DDRS-34560D
5							IBM DDRS-34560D
6							IBM DDRS-34560D
7	2	6	9999	40MB	NONE	NEW DRV	IBM DDRS-34560D
8	2	8	9999	40MB	0	ON-LINE	IBM DDRS-34560D

The Read/Write test progress is indicated by a status bar.

Note – At any time you can press Esc, choose “Read/Write Test,” and then choose “View Read/Write Testing Progress” or “List Current Bad Block Table.”

Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	2	0	9999	40MB	0	ON-LINE	IBM DDRS-34560D
2	2	1	9999	40MB	0	ON-LINE	IBM DDRS-34560D
3	2	2	9999	40MB	0	ON-LINE	IBM DDRS-34560D
4					0	ON-LINE	IBM DDRS-34560D
5							IBM DDRS-34560D
6							IBM DDRS-34560D
7	2	6	9999	40MB	NONE	NEW DRV	IBM DDRS-34560D
8	2	8	9999	40MB	0	ON-LINE	IBM DDRS-34560D

Address	Stat	Address	Stat	Address	Stat	Address	Stat	Address	Stat	Address	Stat

Note – If you want to stop testing the drive, choose “Abort Drive Testing.”

Changing Disk Drive Reserved Space

Before a disk can be included in a logical drive, the RAID controller needs to format an amount of space for storing controller-specific data separately from user data. The default reserved space is 256 MB, but older drives might have only 64 KB of reserved space. Use the default value unless you are using an older (non-upgraded) version of controller firmware that can only recognize 64 KB of reserved space.

Note – It is preferable to upgrade your controller firmware rather than make your drive reserved space backward-compatible.

Note – You can change disk drive reserved space only for drives that are not spares or currently part of a logical drive. Attempting to change reserved space on a drive that is a member of a logical drive displays an error message. Since reserved space is a feature of the physical drive rather than of a logical drive, it is not RAID-protected.

▼ To Change the Allocated Reserved Space on a Drive

1. **Make sure the drive is not part of a logical drive.**
1. **Choose “view and edit scsi Drives” from the Main Menu.**
2. **Select the drive whose reserve space you want to change.**
3. **Choose “disk Reserved space - ”**
If reserved space currently exists, a confirmation prompt asks if you want to remove the reserved space.
4. **Choose Yes to continue.**
The “disk Reserved space - ” menu option now indicates that the reserved space is unformatted.
5. **Choose “disk Reserved space - ” again if you want to allocate reserved space.**
6. **Choose either “256 MB” or “Backward-Compatible (64KB)” to allocate reserved space.**
A confirmation prompt asks if you want to format disk reserved space.
7. **Choose Yes to confirm.**

Viewing and Editing SCSI Channels

This chapter explains how to view and edit SCSI channels for Fibre Channel and SCSI arrays.

Note – Where procedures vary between FC and SCSI arrays, headings are marked “SCSI Only” or “FC Only.”

Topics covered in this chapter include:

- “SCSI Channel Status Table” on page 130
- “Configuring SCSI Channels as Host or Drive” on page 132
- “Creating Additional Host IDs” on page 135
- “Adding a Host ID” on page 136
- “Deleting a Host Channel SCSI ID” on page 137
- “Drive Channel SCSI IDs (Reserved)” on page 138
- “Setting SCSI Channel Termination (SCSI Only) (Reserved)” on page 139
- “Setting Transfer Clock Speed (SCSI Only)” on page 140
- “Setting the SCSI Transfer Width (SCSI Only)” on page 141
- “Viewing and Editing Drive Channel SCSI Targets (SCSI Only)” on page 141
 - “To View or Edit the SCSI Targets for a Selected Drive Channel (SCSI Only)” on page 142
 - “Providing a Slot Number (SCSI Only)” on page 142
 - “To Provide a Slot Number (SCSI Only)” on page 142
 - “Maximum Synchronous Transfer Clock (SCSI Only)” on page 143
 - “To Set Maximum Synchronous Transfer Clock (SCSI Only)” on page 143
 - “Maximum Transfer Width (SCSI Only)” on page 144
 - “To Set Maximum Transfer Width (SCSI Only)” on page 144
 - “Enabling Parity Checking (SCSI Only)” on page 144
 - “To Enable or Disable Parity Checking” on page 145
 - “Allowing Disconnect Support (SCSI Only)” on page 145
 - “To Allow or Disallow Disconnect Support (SCSI Only)” on page 145
 - “Maximum Tag Count (Reserved) (SCSI Only)” on page 146
 - “Restoring a SCSI Target to Default Settings” on page 146
 - “To Restore Default SCSI Target Settings” on page 146

- “Viewing Chip Information” on page 147
- “Viewing Channel Host-ID WWN Information (FC Only)” on page 148
- “Viewing Device Port Names (WWPN) (FC Only)” on page 149
- “Setting a Channel’s Data Rate (FC Only)” on page 150

SCSI Channel Status Table

The SCSI channel status table provides an easy way to view SCSI channel parameters.

▼ To Check and Configure SCSI Channels

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

The status of all SCSI channels for this controller is displayed.



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

Note – Each controller has a separate RS-232 port, as well as an Ethernet chip. This architecture ensures continuous communication should a controller fail. Since the connection is established to only one controller (even when the array is in redundant mode), the CurSyncClk and CurWid settings are displayed for that individual controller. Therefore, if a user maps one LUN to the primary controller, and another LUN to a secondary controller, only the established connection to that individual controller is displayed through the serial and Ethernet port menu.

Note – On SCSI arrays, a mapped SCSI host channel sometimes shows the current sync clock as “Async/Narrow” and correctly identifies 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.

2. Select a SCSI channel to view the additional menu options available for that channel.

Note – SCSI channel menu options differ substantially between Fibre Channel and SCSI arrays. In this chapter, menu options or procedures that apply to only one of these arrays are noted in the section headings as (FC Only) or (SCSI Only).

TABLE 7-1 Parameters Displayed in the Channel Window

Parameters	Description
Chl	Channel's ID
Mode	Channel mode:
	RCC or RCom 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 ID mapping:
	* Multiple IDs were applied (host channel mode only).
	x The ID for host LUNs mapped to this channel in host channel mode. ID for the primary controller in drive channel mode.
	NA No ID applied.
SID	Secondary controller's ID mapping:
	* Multiple IDs (host channel mode only).
	x The ID for host LUNs mapped to this channel in host channel mode. ID for the secondary controller in drive channel mode.
	NA No 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 bus width:
	Wide Channel is set to allow wide (16-bit) transfers.
	Narrow Channel is set to allow narrow (8-bit) transfers.
	Serial Fibre Channel loops do not use narrow or wide bus widths.
S	Signal:
	S Single-ended
	L LVD

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

Parameters	Description
	F Fibre
Term	Terminator status:
	On Termination is enabled.
	Off Termination is disabled.
	NA For a redundant controller communications channel (RCCOM).
CurSynClk	Current SCSI bus synchronous clock:
	xx.x MHz The current speed at which the channel is communicating.
	Async. The channel is communicating asynchronously or no device is detected.
	(Empty) The default SCSI bus synchronous clock has changed. Reset the controller for changes to take effect.
CurWid	Current SCSI bus width:
	Wide The channel is currently servicing wide 16-bit transfers.
	Narrow The channel is currently servicing wide 8-bit transfers.
	Serial Fibre Channel loops do not use narrow or wide bus widths.
	(Empty) The default SCSI bus width has changed. Reset the controller for the changes to take effect.

Configuring SCSI Channels as Host or Drive

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

Fibre Channel Default Channel Settings (FC Only)

Sun StorEdge 3510 FC RAID arrays are preconfigured when they arrive from the factory. Default channel settings and rules are specified as follows.

- Default channel settings are:
 - CH 0, CH 1, CH 4, and CH 5 = Host channels
 - CH 2 and CH 3 = Drive channels

- Channels 2 and 3 (CH 2 and 3) must be drive channels.
- Channels 0, 1, 4, and 5 (CH 0, 1, 4, and 5) can be drive or host channels.

SCSI Default Channel Settings (SCSI Only)

- Channel 0 (CH 0) must be a drive channel.
- Channels 1, 2, and 3 (CH 1, CH 2, CH 3) can 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 need only one host channel.

▼ To Change the SCSI Channel Assignments (FC and SCSI)

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

Note – For SCSI arrays, the Mode column for at least one channel must include the RCC or RCCOM abbreviation for redundant controller communications.

2. Select the channel that you want to modify.
3. Choose “channel Mode.”

For SCSI arrays, you can only toggle between Host and Drive modes.
4. For FC arrays, select the channel mode menu option to which you want to change the selected channel from these available menu options:
 - Host
 - Drive
 - RCCOM

■ Drive+RCCOM

A confirmation prompt is displayed:

Main Menu									
Quick installation									
view and edit Logical drives									
view and edit logical Volumes									
Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	channel Mode			er scsi id	erial	F	NA	1 GHz	Serial
1	Host			ller scsi id	erial	F	NA	2 GHz	Serial
2	Drive			ation	erial	F	NA	2 GHz	Serial
Change Channel Mode ?									
Yes No									
4	H				Serial	F	NA	1 GHz	Serial
5	Host	NA	46	AUTO	Serial	F	NA	2 GHz	Serial

5. Choose Yes if you want to change the host or drive assignment.

A confirmation message asks if you want to reset the controller so that the change takes effect now.

Main Menu									
Quick installation									
view and edit Logical drives									
view and edit logical Volumes									
Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	channel								
1	Host								
2	RCC								
NOTICE: Change made to this setting will NOT take effect until the controller is RESET. Prior to resetting the controller, operation may not proceed normally.									
Do you want to reset the controller now ?									
Yes No									
4	H								
5	H				Serial	F	NA	2 GHz	Serial

6. Choose Yes to reset the controller.



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 made to the secondary controller.

Creating Additional Host IDs

This section describes how to create host IDs for Fibre Channel and SCSI arrays.

Note – To map 1024 partitions on a Fibre Channel array, you must add additional host IDs so that 32 IDs are mapped to the array’s channels. Several such configurations are possible, such as eight IDs mapped to each of the four host channels or sixteen IDs mapped to two channels and none to the other two. For more information, see [“Planning for 1024 LUNs on FC Arrays” on page 79](#).

Note – To map 128 partitions into 128 LUNs on a SCSI array, you must add additional host IDs. A minimum of four host IDs is required; a maximum of six host IDs is possible. For details on mapping 128 LUNs, see [“Planning for 128 LUNs on SCSI Arrays” on page 78](#).

Default Fibre Channel Host IDs (FC Only)

In loop mode, up to 16 IDs can be assigned to each Fibre Channel. In point-to-point mode, only one ID should be assigned to each channel. The array can have a maximum of 32 IDs.

Typically, host IDs are distributed between the primary and secondary controllers to load-balance the I/O in the most effective manner for the network. For more information, refer to the *Sun StorEdge 3000 Family Best Practices Manual* for more information.

Default host channel IDs are displayed in [TABLE 7-2](#).

TABLE 7-2 Default Host Channel IDs

Channel	Primary Controller ID (PID)	Secondary Controller ID (SID)
Channel 0	40	N/A
Channel 1	N/A	42
Channel 4	44	N/A
Channel 5	N/A	46

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

- Edit each host ID number to change the target number of each controller host channel that is seen by the host
- Add other host ID numbers

Note – To map 1024 partitions on a Fibre Channel array, you must add other host IDs. A minimum of eight host IDs per channel are required. For more information, see [“Planning for 1024 LUNs on FC Arrays” on page 79](#).

Default SCSI Host IDs (SCSI Only)

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 other 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).

Adding a Host ID

This section describes how to add a host channel SCSI ID to Fibre Channel and SCSI arrays.

▼ To Select a Unique ID Number for a Host Channel

1. Choose “view and edit Scsi channels” from the Main Menu.
2. Select the host channel on which you want to edit the Primary or Secondary ID.

3. Choose “view and edit scsi Id.”

The existing primary and secondary IDs (if both exist) are displayed.

4. Select the ID you want to edit.

5. Choose “Add Channel SCSI ID.”

Note – For a Fibre Channel array, by default channel 0 has a primary ID (PID) and no secondary ID (SID), while channel 1 has a SID but no PID.

6. Select the primary or secondary controller to which you want to assign the ID.

7. Select an ID number for that controller.

A confirmation prompt is displayed.

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	Cur	ID
0	Host	40	NA	2 GHz	Serial	F	NA	2 GHz	Ser	32
1	ID 40 <Primary Controller>				Serial	F	NA	2 GHz	Ser	33
2<3;D	Add Channel SCSI ID				Serial	F	NA	2 GHz	Ser	34
3<2;D	Primary Controller Secondary Controller				Serial	F	NA	2 GHz	Ser	35
4	H				Serial	F	NA	2 GHz	Ser	36
5	Host	NA	46	2 GHz	Serial	F	NA	2 GHz	Ser	37
										38
										39
										40
										41
										42
										43
										44
										45
										46
										47
										48
										49
										50
										51
										52
										53
										54
										55
										56
										57
										58
										59
										60
										61
										62
										63
										64
										65
										66
										67
										68
										69
										70
										71
										72
										73
										74
										75
										76
										77
										78
										79
										80
										81
										82
										83
										84
										85
										86
										87
										88
										89
										90
										91
										92
										93
										94
										95
										96
										97
										98
										99
										100

Note – On a SCSI array, 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.

8. Choose Yes to confirm.

A confirmation prompt informs you that the controller must be reset for the change to take effect.

9. Choose Yes to reset the controller.

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

Deleting a Host Channel SCSI ID

This section describes how to delete a host channel SCSI ID.

▼ To Delete a Host Channel SCSI ID

1. Choose “view and edit Scsi channels” from the Main Menu.
2. Select the host channel with the SCSI ID you want to delete.
3. Choose “view and edit scsi Id.”
The existing primary and secondary IDs (if both exist) are displayed.
4. Select the ID you want to delete.
5. Choose “Delete Channel SCSI ID.”

A “Delete Primary/Secondary Controller SCSI ID?” confirmation message 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
```

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 ? ync Narrow
ID 1 Yes No 0MHz Wide
A Delete Channel SCSI ID Wide L On Async Narrow

6. Choose Yes.
7. Choose Yes to reset the controller.

A confirmation prompt informs you that the controller must be reset for the change to take effect.

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

Drive Channel SCSI IDs (Reserved)

When a drive channel or DRV+RCC channel is selected, the “view and edit Scsi channel” menu provides two choices used to change the default drive SCSI IDs:

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

Do not use these menu options. They are reserved for specific troubleshooting methods and should be used only by qualified technicians.



Caution – If you edit a drive channel SCSI ID, you might create conflicts with the controller communication channels and create confusion tracking the status of drive IDs.

Setting SCSI Channel Termination (SCSI Only) (Reserved)


Do not use this menu option. It is reserved for specific troubleshooting methods and should be used only by qualified technicians.

Typically the default setting is not changed.

▼ To Enable or Disable SCSI Channel Termination (SCSI Only)

1. Choose “view and edit Scsi channels” from the Main Menu.
2. Select the channel for which you want the terminator enabled or disabled.
3. Choose “scsi Terminator.”

A confirmation message is displayed.



4. Choose Yes.

The controller is reset and the channel configuration is updated.

Setting Transfer Clock Speed (SCSI Only)

Typically, the default setting for “sync transfer clock” is not changed for the host or drive channel.

▼ To View the Options for Host or Drive Transfer Clock Speed (SCSI Only)

Typically the default setting for “Host Channel Clock Speed” is not changed for the host channel.

1. Choose “view and edit Scsi channels” from the Main Menu.
2. Select the drive channel or host channel.
3. Choose “sync transfer Clock.”

A menu of sync transfer clock speeds is displayed.

4. Select a clock speed.

A confirmation prompt is displayed.

The screenshot shows a menu with the following options: Quick installation, view and edit Logical drives, view channel Mode, view and edit scsi Id, view scsi terminator, view sync transfer Clock (highlighted), view wide transfer, parity check - Enabled, and view chip inFormation. Below this is a table with columns: Ch, Host, ID, NA, Clk, DefWid, S, and a list of options. The 'Clk' column shows 40.0MHz for all channels. A 'Change Sync Transfer Clock?' dialog box is open with '40.0MHz' selected and 'Yes' highlighted. A list of clock speed options is also visible: 8.0MHz, 6.7MHz, 5.0MHz, 4.0MHz, 3.3MHz, 2.8MHz, 2.5MHz, and Async.

Ch	Host	ID	NA	Clk	DefWid	S	
2	Drive	7	6	40.0MHz	Wide	L	On Async Nar
3	Drive	7	NA	40.0MHz	Wide	L	On Async Nar

5. Choose Yes.

A confirmation prompt informs you that you must reset the controller for the change to take effect.

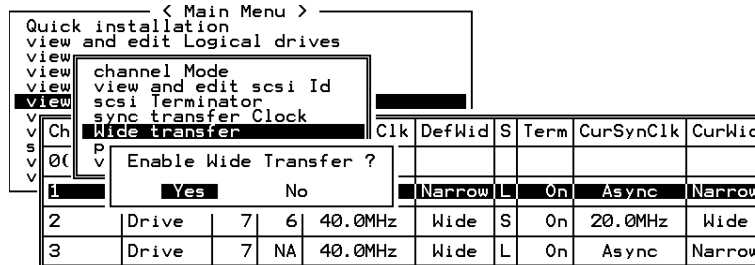
6. Choose Yes to reset the controller.

Setting the SCSI Transfer Width (SCSI Only)

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

▼ To View the Transfer Speed Options (SCSI Only)

1. Choose “view and edit Scsi channels” from the Main Menu.
2. Select the drive channel or host channel.
3. If you want to change the transfer width, choose “Wide transfer.”
A confirmation message asks if you want to enable or disable wide transfer.
4. Choose Yes to confirm.



A confirmation prompt informs you that you must reset the controller for the change to take effect.

5. Choose Yes to reset the controller.

Viewing and Editing Drive Channel SCSI Targets (SCSI Only)

This section describes how to view and edit drive channel SCSI targets.

▼ To View or Edit the SCSI Targets for a Selected Drive Channel (SCSI Only)

1. Choose “view and edit Scsi channels” from the Main Menu.
2. Select the appropriate drive channel.
3. Choose “View and edit scsi target.”
A list of all the SCSI targets and their current settings is displayed.
4. Select a SCSI target to display the “View and edit scsi target” menu.

Slot	Ch1	ID	SyncClk	XfrWid	ParityChk	Disconnect	TagCount
2	0	10	Wide	Enabled	Enabled	Def (32)	
Slot number				bled	Enabled	Def (32)	
maximum sync. xfer Clock				bled	Enabled	Def (32)	
maximum xfer Width				bled	Enabled	Def (32)	
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)	
1	2	5	10	Wide	Enabled	Enabled	Def (32)
2	2	8	10	Wide	Enabled	Enabled	Def (32)
3	2	9	10	Wide	Enabled	Enabled	Def (32)

Procedures for changing these parameters are described in the following sections.

Note – Alter the SCSI target settings only when adjustments need to be made to specific devices on a drive channel. You can change the SCSI parameters for specific drives when mixing different drives or connecting other SCSI devices such as a CD-ROM on a drive channel. Neither mixing drives nor connecting a CD-ROM is recommended for the controller.

Providing a Slot Number (SCSI Only)

This section describes how to provide a slot number for a SCSI target.

▼ To Provide a Slot Number (SCSI Only)

1. Choose “view and edit Scsi channels” from the Main Menu.
2. Select the appropriate drive channel.
3. Choose “View and edit scsi target.”

A list of all the SCSI targets and their current settings is displayed.

4. Select a SCSI target.

5. Choose "Slot number."

A text area is displayed.

6. Type a slot number in the text area and press Return.

Maximum Synchronous Transfer Clock (SCSI Only)

This section describes how to set a maximum synchronous transfer clock period.

▼ To Set Maximum Synchronous Transfer Clock (SCSI Only)

1. Choose "view and edit Scsi channels" from the Main Menu.

2. Select the appropriate drive channel.

3. Choose "View and edit scsi target."

A list of all the SCSI targets and their current settings is displayed.

4. Select a SCSI target.

5. Choose "maximum sync. xfer Clock."

A text area is displayed.

Slot	Chl	ID	SyncClk	XfrWid	ParityChk	Disconnect	TagCount	
1	0	9	Wide	Enabled	Enabled	Def (32)		
Slot number		maximum sync. xfer Clock		bled	Enabled	Def (32)		
Synchronous Transfer Period Factor								2)
Maximum Sync. Xfer Clock: 9_								2)
1	5	9	Wide	Enabled	Enabled	Def (32)		
2	6	9	Wide	Enabled	Enabled	Def (32)		
3	8	9	Wide	Enabled	Enabled	Def (32)		

6. Type a clock transfer period factor and press Return.

Maximum Transfer Width (SCSI Only)

This section describes how to set maximum transfer width. You can toggle between Wide and Narrow maximum transfer widths.

▼ To Set Maximum Transfer Width (SCSI Only)

1. Choose “view and edit Scsi channels” from the Main Menu.
2. Select the appropriate drive channel.
3. Choose “View and edit scsi target.”

A list of all the SCSI targets and their current settings is displayed.

4. Select a SCSI target.
5. Choose “maximum xfer Width.”

A confirmation prompt is displayed.

Slot	Chl	ID	SyncClk	XfrWid	ParityChk	Disconnect	TagCount
1	1	0	9	Wide	Enabled	Enabled	Def(32)
2	1	1	9	Narrow	Enabled	Enabled	Def(32)
Slot number				bled	Enabled	Def(32)	
maximum sync. xfer Clock				bled	Enabled	Def(32)	
maximum xfer Width				bled	Enabled	Def(32)	
Set SCSI Target Maximum Xfer Wide Supported ?							2)
Yes							2)
No							2)
1	6		9	Wide	Enabled	Enabled	Def(32)
1	8		9	Wide	Enabled	Enabled	Def(32)

6. Choose Yes to change the setting from Wide to Narrow or from Narrow to Wide.

Enabling Parity Checking (SCSI Only)

The integrity of redundant data on fault-tolerant arrays (RAID 3 and 5) is verified with a parity check. 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 is reported and the new, correct parity is substituted for the stored parity.

You can toggle between enabling and disabling parity checking.

4. Select a SCSI target.
5. Choose "Disconnect support."

A confirmation prompt is displayed.

Slot	Chl	ID	SyncClk	XfrWid	ParityChk	Disconnect	TagCount	
1	2	0	9	Wide	Enabled	Enabled	Def (32)	
2	2	1	9	Wide	Enabled	Enabled	Def (32)	
Slot number						bled	Enabled	Def (32)
maximum sync. xfer Clock						bled	Enabled	Def (32)
maximum xfer Width						bled	Enabled	Def (32)
Parity check						bled	Enabled	Def (32)
Disconnect support						bled	Enabled	Def (32)
Disallow target disconnect ?							Enabled	Def (32)
Yes							Enabled	Def (32)
No							Enabled	Def (32)
7	2	8	9	Wide	Enabled	Enabled	Def (32)	

6. Choose Yes to allow disconnect support if it is currently disallowed, or to disallow disconnect support if it is currently allowed.

Maximum Tag Count (Reserved) (SCSI Only)

Do not use this menu option. To set maximum tag count, use the Maximum Tag Count menu option available from the "view and edit Configuration parameters" menu. See ["Maximum Tag Count \(Tag Command Queuing\)"](#) on page 173 for more information.

Restoring a SCSI Target to Default Settings

Use this menu option to restore the default settings for this SCSI target.

▼ To Restore Default SCSI Target Settings

1. Choose "view and edit Scsi channels" from the Main Menu.

2. Select the appropriate drive channel.
3. Choose “View and edit scsi target.”

A list of all the SCSI targets and their current settings is displayed.

4. Select a SCSI target.
5. Choose “Restore to default setting.”
6. Choose Yes to confirm.

The SCSI target settings are reset to default values.

Viewing Chip Information

Each controller has multiple channels (I/O paths) and each channel is powered by an I/O processor. The “view chip inFormation” menu option provides information about the host or drive channel’s chip type and revision level, and the firmware identifier, which might include version information.

▼ To View Chip Information

1. Choose “view and edit Scsi channels.”
2. Select a host or drive channel.
3. Choose “view chip inFormation.”

The channel’s chip information is displayed.

< Main Menu >														
Quick installation														
view and edit Logical drives														
view and edit logical Volumes														
U	Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid				
U	0	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Serial				
U	1 channel Mode									1	F	NA	2 GHz	Serial
U	view and edit scsi Id									1	F	NA	2 GHz	Serial
U	view chip inFormation									1	F	NA	2 GHz	Serial
U	3<	Chip Type	ISP2312		wwpn>		1	F	NA	2 GHz	Serial			
U	4	Chip Rev. ID	2				1	F	NA	1 GHz	Serial			
U	4	Chip FW Rev. ID	3.01.18				1	F	NA	1 GHz	Serial			
U	5	Host	NA	46	AUTO	Serial	F	NA	2 GHz	Serial				

Viewing Channel Host-ID WWN Information (FC Only)

Each controller has multiple channels (I/O paths) and each channel is powered by an I/O processor. Use the “view channel host-id/Wwn” menu option to view the I/O processor’s node name and port name for a selected host channel:

- worldwide node name (WWNN)
- worldwide port nName (WWPN)

Some host-based management software requires these names in order to address storage devices.

▼ To View a Channel’s Host-ID/WWN (FC Only)

1. Choose “view and edit Scsi channels.”
2. Select a host channel.
3. Choose “view channel host-id/Wwn.”

The channel’s worldwide node name and worldwide port name are displayed.

The screenshot shows a menu titled "< Main Menu >" with options: "Quick installation", "view and edit Logical drives", and "view and edit logical Volumes". A table of channels is displayed with columns: Ch1, Mode, PID, SID, DefSynClk, DefWid, S, Term, CurSynClk, and CurWid. Channel 0 is selected, and a sub-menu is shown with options: "channel Mode", "view and edit scsi Id", "view chip inFormation", "view channel host-id/wpn", and "D". The "view channel host-id/wpn" option is selected, displaying "WWNN: 0x206000C0FF002F18" and "WWPN: 0x266000C0FFE02F18".

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Serial
1						1 F	NA	2 GHz	Serial
2						1 F	NA	2 GHz	Serial
3						1 F	NA	2 GHz	Serial
4						1 F	NA	1 GHz	Serial
5						Serial F	NA	2 GHz	Serial

Viewing Device Port Names (WWPN) (FC Only)

The “View device port name list(wwpn)” menu option displays device port names for the host bus adapter (HBA) detected on a host loop. Device port names on the loop are displayed, with the exception of the controller’s I/O processor itself.

Once you have displayed an HBA port name here, you can add it to the WWN list displayed when you choose the “Host-ID WWN name list” menu option from the “View and Edit Host LUN” menu. Adding port names to this list can speed the Host LUN mapping process.

You can also assign each port on the Host-ID WWN name list a name of your own, for ease of identification. This becomes especially useful when you want to set up multiple filtering entries so that you can grant or deny specific hosts access to logical drives. See “[Creating Host Filter Entries](#)” on [page 89](#) for more information.

▼ To View a Channel’s Device Port Name List

1. Choose “view and edit Scsi channels.”
2. Select a host channel.
3. Choose “View device port name list(wwpn).”

A list of the device port names on the host loop are displayed.

```

      < Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
  
```

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Serial
1				0	Serial	F	NA	2 GHz	Serial
2<3;C>	DRU+RCC	14	15	AUTO	Serial	F	NA	2 GHz	Serial
3<2;C>	DRU+RCC	14	15	AUTO	Serial	F	NA	2 GHz	Serial
4	Host	44	NA	AUTO	Serial	F	NA	1 GHz	Serial
5	Host	NA	46	AUTO	Serial	F	NA	2 GHz	Serial

Setting a Channel's Data Rate (FC Only)

An FC channel communicates at a rate of either 1 GHz or 2 GHz. You can set this data rate manually or use the default "Auto" setting to autosense the rate of communications.

▼ To Set a Channel's Data Rate

1. Choose "view and edit Scsi channels."
2. Select a host or drive channel.
3. Choose "Data rate."

A menu of data rate choices is displayed:

- Auto
- 1 GHz
- 2 GHz

4. Choose a data rate from the menu.

A confirmation message is displayed.

The screenshot shows a terminal window with a menu at the top: "Quick installation", "view and edit Logical drives", and "view and edit logical Volumes". Below this is a table of channels. Channel 2 is selected, and a dialog box is open asking "Set Fibre Channel Data Rate ?" with options "Auto", "1", and "2". A "Data" menu is also visible with "Yes" and "No" options.

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
2	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Serial
1						1 F	NA	2 GHz	Serial
2<							NA	2 GHz	Serial
3<							NA	2 GHz	Serial
4							NA	1 GHz	Serial
5	Host	NA	46	AUTO	Serial	F	NA	2 GHz	Serial

5. Choose Yes to continue.

A confirmation message asks if you wish to reset the controller so that your change takes effect.

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
  
```

Chl	Mod								
0	Hos	NOTICE: Change made to this setting will NOT take effect until the controller is RESET. Prior to resetting the controller, operation may not proceed normally.							
		Do you want to reset the controller now ?							
		Yes				No			
1									
2									
3									
4									
5	Host	NA	46	AUTO	Serial	F	NA	2 GHz	Serial

6. Choose Yes to continue.

Viewing and Editing Configuration Parameters

This chapter describes viewing and editing configuration parameters. Topics covered include:

- [“Communication Parameters” on page 154](#)
 - [“Setting an IP Address” on page 154](#)
 - [“To Set an Array’s IP Address” on page 155](#)
- [“Caching Parameters” on page 156](#)
 - [“Enabling and Disabling Write-Back Cache” on page 156](#)
 - [“Optimization Modes” on page 157](#)
 - [“Optimization for Random I/O \(32 Kbyte block size\)” on page 158](#)
 - [“Optimization for Sequential I/O \(128 Kbyte block size\)” on page 158](#)
- [“Host-side SCSI Parameters Menu Options” on page 160](#)
 - [“Maximum Queued I/O Count” on page 161](#)
 - [“To Set the Maximum Queued I/O Count” on page 161](#)
 - [“LUNs Per Host SCSI ID” on page 162](#)
 - [“To Change the Number of LUNs Per Host SCSI ID” on page 162](#)
 - [“Maximum Number of Concurrent Host-LUN Connections” on page 163](#)
 - [“To Change the Maximum Number of Concurrent Host-LUN Connections” on page 163](#)
 - [“Number of Tags Reserved for Each Host LUN Connection” on page 164](#)
 - [“To Modify the Tag Command Queuing on the Host-LUN Connection” on page 164](#)
 - [“Host Cylinder/Head/Sector Mapping Configuration” on page 165](#)
 - [“To Configure Sector Ranges, Head Ranges, and Cylinder Ranges” on page 166](#)
 - [“Preparing for Logical Drives Larger Than 253 Gbytes on Solaris Systems” on page 166](#)
 - [“To Confirm or Change the Fibre Connection for the Array” on page 169](#)
- [“Drive-side SCSI Parameters Menu” on page 170](#)
 - [“SCSI Motor Spin-Up \(Reserved\)” on page 171](#)
 - [“SCSI Reset at Power-Up \(Reserved\)” on page 172](#)
 - [“Disk Access Delay Time” on page 172](#)
 - [“SCSI I/O Timeout” on page 173](#)

- “Maximum Tag Count (Tag Command Queuing)” on page 173
- “Periodic SAF-TE and SES Device Check Time” on page 175
- “Periodic Drive Check Time” on page 175
- “Periodic SAF-TE and SES Device Check Time” on page 175
- “Auto-Assign Global Spare Drive” on page 177
- “Disk Array Parameters Menu” on page 178
 - “Setting Rebuild Priority” on page 178
 - “Verification on Writes” on page 179
- “Redundant Controller Parameters Menu (Reserved)” on page 181
- “Preparing for Logical Drives Larger Than 253 Gbytes on Solaris Systems” on page 166
- “Controller Parameters” on page 181
 - “Controller Name” on page 181
 - “To View and Display the Controller Name” on page 182
- “Controller Failover” on page 185
- “Rebuilding Logical Drives” on page 185
 - “Automatic Logical Drive Rebuild” on page 185
 - “Manual Rebuild” on page 188
 - “Concurrent Rebuild in RAID 1+0” on page 189
- “Recovering From Fatal Drive Failure” on page 190
 - “To Recover From Fatal Drive Failure” on page 190
 - “LCD Title Display - Controller Logo (Reserved)” on page 182
 - “Password Validation Timeout” on page 182
 - “Controller Unique Identifier (Reserved)” on page 184
 - “SDRAM ECC Function (Reserved)” on page 184

Communication Parameters

The “Communication parameters” menu option enables you to view communication settings. Most of these parameters are reserved and should not be changed. Use the “Internet Protocol (TCP/IP) menu option to set or change your array’s IP address.

Setting an IP Address

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

- The Sun StorEdge Configuration Service application. Refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide* for details.
- The firmware application you access when you use the `telnet` command to connect to the IP address of the controller.

To access the array using the Ethernet port, you must set up an IP address for the controller. You can set the IP address by typing in values for the IP address itself, the subnet mask, and the IP address of the gateway manually. If your network is using a Reverse Address Resolution Protocol (RARP) server or a Dynamic Host Configuration Protocol (DHCP) server to automatically configure IP information for devices on the network, you can specify the appropriate protocol instead of typing in the information manually.



Caution – If you assign an IP address to an array in order to manage it out-of-band, for security reasons make sure that the IP address is on a private network rather than a publicly routable network.

▼ To Set an Array's IP Address

To set the IP address, subnet mask, and gateway addresses of the RAID controller, perform the following steps:

1. **Access the array through the COM port on the controller module of the array.**
2. **Choose “view and edit Configuration parameter → Communication Parameters → Internet Protocol (TCP/IP).”**
3. **Select the chip hardware address.**
4. **Choose “Set IP Address.”**
5. **Type in the desired IP address, subnet mask, and gateway address, choosing each menu option in turn.**

If your network sets IP addresses using a RARP server, type RARP rather than an IP address and do not type in a subnet mask or gateway address. If your network sets IP addresses using a DHCP server, type DHCP rather than an IP address and do not type in a subnet mask or gateway address.

6. **Press Esc to continue.**

A confirmation prompt is displayed.

Change/Set IP Address ?

7. **Choose Yes to continue.**

Note – You need to reset the controller for the configuration to take effect.

You are prompted to reset the controller.

8. Choose Yes to reset the controller.

Caching Parameters

Caching parameters include the write-back cache and optimization modes described below.


Enabling and Disabling Write-Back Cache

The write-back cache function significantly enhances controller performance. When it is disabled, a write-through strategy replaces it. A 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

1. Choose “view and edit Configuration parameters → Caching Parameters → Write-Back Cache.”

A confirmation message asks if you want to change the write-back cache setting.



2. Choose Yes to confirm.

Optimization Modes

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
- Optimization for sequential I/O

The random I/O optimization mode reads and writes small blocks of data, and sequential optimization mode reads and writes large blocks of data. [TABLE 8-1](#) specifies the stripe size for each RAID level according to the optimization mode.

TABLE 8-1 Stripe Size Per Optimization Mode

RAID Level	Sequential I/O	Random I/O
0, 1, 5	128	32
3	16	4

The types of applications appropriate for random and sequential optimization are described in [“Database and Transaction-Based Applications” on page 158](#) and [“Video Recording, Playback, and Imaging Applications” on page 158](#).

Optimization Limitations

There are two limitations that apply to the optimization modes:

- One optimization mode must be applied to all logical units in a RAID array.
- Once the optimization mode is chosen 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.

This limitation results from the redundant configuration of controllers. Data inconsistency can occur when a controller configured with one optimization mode is used to replace a failed controller with a different mode.

Note – The maximum allowable size of a logical drive optimized for sequential I/O is 2 Tbyte. The maximum allowable size of a logical drive optimized for random I/O is 512 Gbyte. When creating a logical drive that is greater than these limits, an error message is displayed.

Database and Transaction-Based Applications

Typical transaction sizes for transactions and database updates range from 2 Kbyte to 4 Kbyte. These applications keep each transaction small so that I/O transfers are not clogged.

Transaction-based applications do not read or write data in a sequential order. Instead, access to data occurs randomly. Transaction-based performance is usually measured in I/O operations per second (IOPS).

Video Recording, Playback, and Imaging Applications

Video playback, multimedia post-production editing, and similar applications read and write large files to and from storage in sequential order. The size of each I/O operation typically ranges from 28 Kbytes up to 1 MByte or higher. Performance is measured in MBytes per second.

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

Optimization for Random I/O (32 Kbyte block size)

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

Optimization for Sequential I/O (128 Kbyte 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 parameters are also changed to optimize for sequential I/O. These changes take effect after the controller resets.

The logical drive, cache memory, and other controller parameters are adjusted for the use of video and imaging applications.

Maximum Number of Disks and Maximum Usable Capacity for Random and Sequential Optimization

Your choice of random or sequential optimization affects the maximum number of disks you can include in an array and the maximum usable capacity of a logical drive. The following tables contain the maximum number of disks per logical drive and the maximum usable capacity of a logical drive.

Note – You can have a maximum of eight logical drives and 36 disks, using one array and two expansion units.

TABLE 8-2 Maximum Number of Disks per Logical Drive for a 2U Array

Disk Capacity (GByte)	RAID 5 Random	RAID 5 Sequential	RAID 3 Random	RAID 3 Sequential	RAID 1 Random	RAID 1 Sequential	RAID 0 Random	RAID 0 Sequential
36.2	14	31	14	31	28	36	14	36
73.4	7	28	7	28	12	30	6	27
146.8	4	14	4	14	6	26	3	13

TABLE 8-3 Maximum Usable Capacity (Gbyte) per Logical Drive for a 2U Array

Disk Capacity	RAID 5 Random	RAID 5 Sequential	RAID 3 Random	RAID 3 Sequential	RAID 1 Random	RAID 1 Sequential	RAID 0 Random	RAID 0 Sequential
36.2	471	1086	471	1086	507	543	507	1122
73.4	440	1982	440	1982	440	1101	440	1982
146.8	440	1908	440	1908	440	1908	440	1908


Note – You might not be able to use all disks for data when using 36 146-Gbyte disks. Any remaining disks can be used as spares.

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

▼ To Choose the Optimization Mode for All Drives

1. Choose “view and edit Configuration parameters → Caching Parameters → Optimization for Random I/O, or choose “view and edit Configuration parameters → Caching Parameters → Optimization for Sequential I/O.”

A confirmation message is displayed.



2. Choose Yes to confirm.

Host-side SCSI Parameters Menu Options

The Host-side SCSI Parameters menu options are discussed in the following sections:

- [Maximum Queued I/O Count](#)
- [LUNs Per Host SCSI ID](#)
- [Maximum Number of Concurrent Host-LUN Connections](#)
- [Number of Tags Reserved for Each Host LUN Connection](#)
- [Host Cylinder/Head/Sector Mapping Configuration](#)
- [Preparing for Logical Drives Larger Than 253 Gbytes on Solaris Systems](#)
- [Peripheral Device Type Parameters \(Reserved\)](#)
- [Fibre Connection Options \(FC Only\)](#)

Maximum Queued I/O Count

The “Maximum Queued I/O Count” menu option enables you to configure the maximum number of I/O operations per logical drive that can be accepted from servers. The predefined range is from 1 to 1024 I/O operations per logical drive, or you can choose the “Auto” (automatically configured) setting. The default value is 256 I/O operations per logical drive.

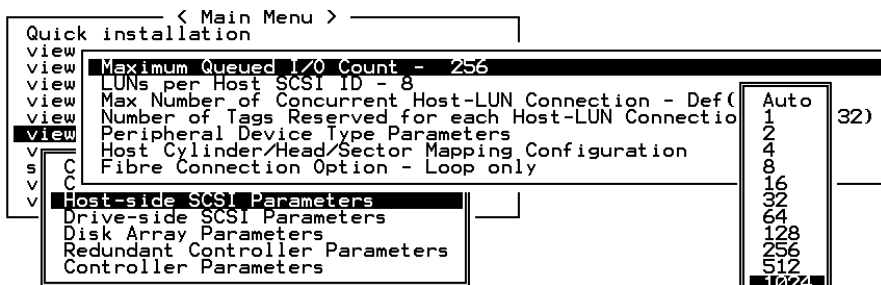
The maximum number of queued I/O operations is 4096. (A maximum of eight logical drives, multiplied by the maximum of 1024 queued I/O operations per logical drive, equals a total of 4096 queued I/O operations).

The appropriate “Maximum Queued I/O Count” setting depends on how many I/O operations attached servers are performing. This can vary according to the amount of host memory present as well as the number of drives and their size. If you increase the amount of host memory, add more drives, or replace drives with larger drives, you might want to increase the maximum I/O count. But usually optimum performance results from using the “Auto” or “256” settings.

▼ To Set the Maximum Queued I/O Count

1. Choose “view and edit Configuration parameters → Host-side SCSI Parameters → Maximum Queued I/O Count.”

A list of I/O count values is displayed.



2. Choose a value.
A confirmation message is displayed.
3. Choose Yes to confirm.

LUNs Per Host SCSI ID


This function is used to change the number of LUNs per host SCSI ID. Each time a host channel ID is added, it uses the number of LUNs allocated in this setting, no matter how many LUNs are actually mapped to it. The default setting is 32 LUNs, with a predefined range of 1 to 32 LUNs per logical drive available.

Note – For the Sun StorEdge 3310 SCSI array, the maximum number of LUN assignments is 128. If you use the default setting of 32 LUNs per host ID, you can only add four host channel IDs (4 x 32 = 128). If you want to allocate more than four host channel IDs, you need to set the LUNs per Host SCSI ID parameter to a value less than 32.

▼ To Change the Number of LUNs Per Host SCSI ID

1. Choose “view and edit Configuration parameters → Host-side SCSI Parameters → LUNs per Host SCSI ID.”

A list of values is displayed.



2. Choose a value.

A confirmation message is displayed.

3. Choose Yes to confirm.

Maximum Number of Concurrent Host-LUN Connections

The “Max Number of Concurrent Host-LUN Connection” menu option is used to set the maximum number of concurrent host-LUN connections. Change this menu option setting only if you have more than four logical drives or partitions. Increasing this number might 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 (IDs 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).

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

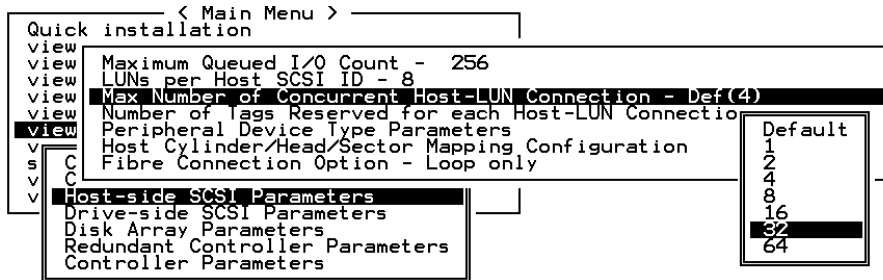
If there is I/O in the cache with four different nexus, and another host I/O comes with a nexus different than the four in the cache (for example, host A accesses ID 3), the controller returns busy. This occurs with the concurrent active nexus; if the cache is cleared, it accepts four different nexus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

▼ To Change the Maximum Number of Concurrent Host-LUN Connections

1. Choose “view and edit Configuration parameters → Host-side SCSI Parameters → Max Number of Concurrent Host-LUN Connection.”

A list of values is displayed. For FC arrays, this list ranges from 1 to 1024. For SCSI arrays, the list ranges from 1 to 128.

2. Choose a value.



A confirmation message is displayed.

3. Choose Yes to confirm.

Number of Tags Reserved for Each Host LUN Connection

This menu option 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. The default factory setting should not be changed unless necessary.

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

▼ To Modify the Tag Command Queuing on the Host-LUN Connection

1. From the Main Menu, choose “view and edit Configuration parameters → Host-side SCSI Parameters → Number of Tags Reserved for each Host-LUN Connection.”

A list of values is displayed.

```
Quick installation
view
view Maximum Queued I/O Count - 256
view LUNs per Host SCSI ID - 8
view Max Number of Concurrent Host-LUN Connection - Def(4)
view Number of Tags Reserved for each Host-LUN Connection - Def(32)
view Peripheral Device Type Parameters
view Host Cylinder/Head/Sector Mapping Configuration
view Fibre Connection Option - Loop only
v
s
c
v
Host-side SCSI Parameters
Drive-side SCSI Parameters
Disk Array Parameters
Redundant Controller Parameters
Controller Parameters
```

Default
1
2
4
8
16
32
64
128
256

2. Choose a value.

A confirmation message is displayed.

3. Choose Yes to confirm.

Host Cylinder/Head/Sector Mapping Configuration

SCSI drive capacity is determined by the host computer according to the number of blocks. Some host operating systems read the capacity of the array based on the cylinder/head/sector count of the drives. The RAID controller firmware enables you to either specify the appropriate number of cylinders, heads, and sectors, or to use the Variable menu option for one or more of these settings. When you use the Variable menu option, the firmware calculates the settings appropriately.


Leaving the cylinder, head, and sector - settings at "Variable" ensures that all three values are calculated automatically. If you choose a specific value for one of these settings and leave the other two set to "Variable," the firmware calculates the other two settings. If you set two, the firmware automatically calculates the third.

For the Solaris operating environment, the number of cylinders cannot exceed 65,535, so you can choose "< 65536 Cylinders" and "255 Heads" to cover all logical drives over 253 GB and under the maximum limit. The controller automatically adjusts the sector count, and then the operating environment can read the correct drive capacity.

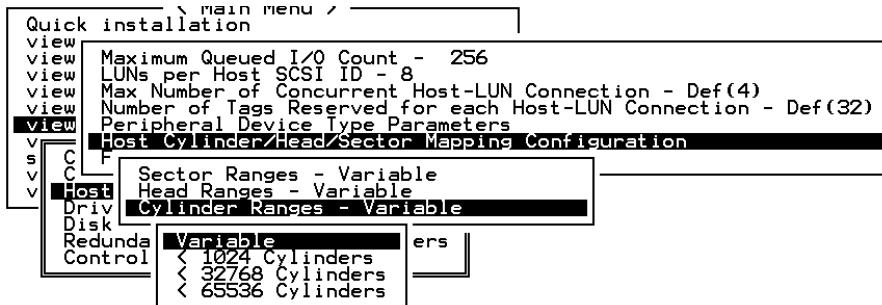
After changing the size of a disk in the Solaris operating environment, run the format utility and choose the 0, autoconfigure option from the menu. This enables the host to reconfigure the size of the disk appropriately and relabel the disk with the current firmware revision level.

▼ To Configure Sector Ranges, Head Ranges, and Cylinder Ranges

1. Choose “view and edit Configuration parameters → Host-Side SCSI Parameters → Host Cylinder/Head/Sector Mapping Configuration → Sector Ranges.”
2. Choose a value.
3. Choose “Head Ranges” and choose a value.



4. Choose “Cylinder Ranges” and choose a value.



Preparing for Logical Drives Larger Than 253 Gbytes on Solaris Systems

The Solaris operating environment requires drive geometry for various operations, including `newfs`. For the appropriate drive geometry to be presented to the Solaris operating environment for logical drives larger than 253 Gbyte, change the default settings to “< 65536 Cylinders” and “255 Heads” to cover all logical drives over 253 GB. The controller automatically adjusts the sector count, and then the operating environment can read the correct drive capacity.

For Solaris operating environment configurations, use the values in the following table.

TABLE 8-4 Cylinder and Head Mapping for the Solaris Operating Environment

Logical Drive Capacity	Cylinder	Head	Sector
< 253 GB	variable (default)	variable (default)	variable (default)
253 GB - 1 TB	< 65536 Cylinders *	255 *	variable (default)

* These settings are also valid for all logical drives under 253 GBytes.

Note – Earlier versions of the Solaris operating environment do not support drive capacities larger than 1 terabyte.

▼ To Prepare Logical Drives Larger than 253 Gbytes

1. Choose “view and edit Configuration parameters → Host-Side SCSI Parameters → Host Cylinder/Head/Sector Mapping Configuration → Sector Ranges - Variable → 255 Sectors → Head Ranges - Variable.”
2. Specify “255 Heads.”
3. Choose “Cylinder Ranges - Variable → < 65536 Cylinders.”

Peripheral Device Type Parameters (Reserved)

Do not use this menu option to change the Peripheral Device Type setting from “Enclosure Services Device.”

The “Peripheral Device Type Parameters” menu option is used only when attempting to configure an array through an in-band connection before a logical drive has been created and mapped to a host LUN. If you follow the instructions for creating a logical drive using a `tip` or telnet session, using the “Peripheral Device Type Parameters” menu option is unnecessary.



Caution – Changing this setting might cause unexpected results.

Note – Do not change the Peripheral Device Qualifier setting from “Connected.”

Fibre Connection Options (FC Only)

Your FC array contains internal circuits called port bypass circuits (PBCs). These PBCs are controlled by firmware configuration settings. Choose the “Loop only” menu option from the Fibre Connection Option menu to make a FC loop configuration possible. Choose the “Point to point only” menu option to make point-to-point connections possible.

Note – It is important that you choose the correct one of these two options for your configuration.



Caution – An additional menu option defaults to a loop configurations but, upon failure to connect at boot time, switches to a point-to-point configuration. Do not use this option unless directed to use it by technical support personnel.

Refer to *Sun StorEdge 3000 Family Best Practices Manual* for the 3510 FC Array and *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the 3510 FC Array for more information about point-to-point and loop configurations.

It is important for point-to-point configurations to also specify only a primary Fibre Channel Host ID (PID) or a secondary ID (SID) for each host channel. For loop configurations with failover, it is important to specify both a PID and SID. See [“Default Fibre Channel Host IDs \(FC Only\)” on page 135](#) for more information about creating host IDs.

Note – The following steps show you how to change a loop configuration to a point-to-point configuration.

▼ To Confirm or Change the Fibre Connection for the Array

1. Choose “view and edit Configuration parameters → Host-side SCSI Parameters → Fibre Connection Option.”
2. Choose a connection type.

```
< Main Menu >
quick installation
view
view Maximum Queued I/O Count - 256
view LUNs per Host SCSI ID - 32
view Max Number of Concurrent Host-LUN Connection - 1024
view Number of Tags Reserved for each Host-LUN Connection - Def(32)
view Peripheral Device Type Parameters
view Host Cylinder/Head/Sector Mapping Configuration
view Fibre Connection Option - Loop only
  C
  C
  v
  s
  v
  v
  Host Loop only
  Driv Point to point only
  Disk Loop preferred,otherwise point to point
  Redu Point to point preferred,otherwise Loop
  Cont
  DMEP Parameters
```



Caution – Do not use the “Loop preferred, otherwise point to point” menu option. This menu option is reserved for special use and should be used only if directed by technical support personnel.

A confirmation message is displayed.

3. Choose Yes to continue.

```
< Main Menu >
quick installation
view
view Maximum Queued I/O Count - 256
view LUNs per Host SCSI ID - 32
view Max Number of Concurrent Host-LUN Connection - 1024
view Number of Tags Reserved for each Host-LUN Connection - Def(32)
view Peripheral Device Type Parameters
v Host Cylinder/Head/Sector Mapping Configuration
s C Fibre Connection Option - Loop only
v C
v Host Loop only
v Driv Point to point only
Disk Lo
Redu Po
Cont Po
DMEP Pa
Set Fibre Channel Connection Option ?
Yes No
```

Note – You need to reset your controller for this configuration change to take effect.

To reset your controller, perform these steps.

4. Choose “system Functions → Reset controller.”

Drive-side SCSI Parameters Menu


The Drive-side SCSI Parameters menu options include:

- SCSI Motor Spin-Up (Reserved)
- SCSI Reset at Power-Up (Reserved)
- Disk Access Delay Time
- SCSI I/O Timeout
- Maximum Tag Count (Tag Command Queuing)
- Periodic Drive Check Time
- Periodic SAF-TE and SES Device Check Time
- Periodic SAF-TE and SES Device Check Time
- Drive Predictable Failure Mode (SMART)
- Auto-Assign Global Spare Drive

These parameters are user-configurable. However they should not be changed from their preset values without good reason, and without an understanding of potential impacts on performance or reliability.

▼ To Access the Drive-side Parameter Menu

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters.”
The Drive-side SCSI parameters menu is displayed.



SCSI Motor Spin-Up (Reserved)



Caution – Do not use the SCSI Motor Spin-Up menu option. It is reserved for specific troubleshooting methods and should be used only by qualified technicians.

The SCSI spin-up decides how the SCSI drives in a disk array are started. When the power supply is unable to provide sufficient current for the hard drives and controllers that are powered on at the same time, spinning-up the hard drives serially uses less current.

If the drives are configured as “Delay Motor Spin-up” or “Motor Spin-up” in Random Sequence, some of these drives might not be ready for the controller to access when the array powers up. Increase the disk access delay time so that the controller will wait a longer time for the drive to be ready.

By default, all hard drives spin up when powered on. These hard drives can be configured so that they do not all spin up at the same time.

▼ To Spin-Up SCSI Hard Drives (Reserved)

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → SCSI Motor Spin-Up.”
A confirmation message is displayed.
2. Choose Yes.

SCSI Reset at Power-Up (Reserved)



Caution – Do not use the SCSI Reset at Power-Up menu option. It is reserved for specific troubleshooting methods and should be used only by qualified technicians.

By default, when the controller is powered on, it sends a SCSI bus reset command to the internal SCSI bus. When disabled, it does not send a SCSI bus reset command when powered on.

▼ To Reset SCSI at Power-Up (Reserved)

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → SCSI Reset at Power-Up.”

A confirmation message is displayed.

2. Choose Yes.
3. Power off all hard drives and controller.
4. Power the hard drives and controller on again.

The controller spins up the hard drives in sequence, separated by a four-second interval.

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 from no delay to 75 seconds.

▼ To Set Disk Access Delay Time

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → Disk Access Delay Time.”

A list of selections is displayed.

2. Choose the desired delay time.
A confirmation message is displayed.
3. Choose Yes.

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.



Caution – The correct setting for “SCSI I/O Timeout” is 30 seconds for Fiber Channel arrays and 15 seconds for SCSI arrays. Do not change this setting. Setting the timeout to a lower value, or to Default, causes 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 causes the controller to keep waiting for a drive, and it can sometimes cause a host timeout.

When the drive detects a media error while reading from the drive platter, it retries the previous reading or recalibrates 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 brands and models of drives.

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

▼ To Choose SCSI I/O Timeout

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → SCSI I/O Timeout –.”

A list of selections is displayed.

2. Select a timeout.

A confirmation message is displayed.

3. Choose Yes.

Maximum Tag Count (Tag Command Queuing)

The maximum tag count 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, enabling the drive to finish the requests more quickly.

The cache size and maximum number of tags varies between brands and models of drive. Use the default setting of 32.

Note – Changing the maximum tag count to “Disable” causes the Write-Back cache in the hard drive to not be 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.

It is possible to configure command tag queuing with a maximum tag count of 128 (SCSI) and 256 (FC).

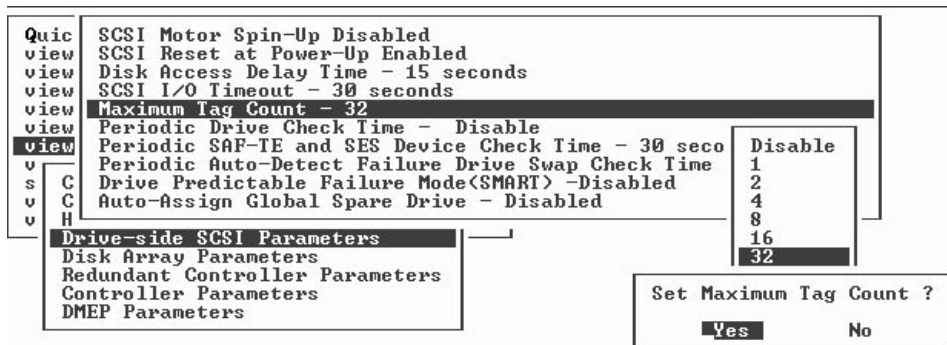
▼ To Change the Default Maximum Tag Count Setting

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → Maximum Tag Count.”

A menu of available tag count values is displayed.

2. Select a tag count number.

A confirmation message is displayed.



Choose Yes to confirm.



Caution – Disabling the maximum tag count disables the drive’s internal cache.

3. Reset the controller to have the change take effect.

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. The default value is Disabled, which means that if a drive is removed from the bus, the controller does 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 all of the drives that are listed under “view and edit scsi Drives” at the specified interval. If any drive is then removed, the controller recognizes the removal even if a host does not access that drive.

▼ To Set the Periodic Drive Check Time

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → Periodic Drive Check Time -.”

A list of intervals is displayed.

2. Choose the interval you want

A confirmation message is displayed.

3. Choose Yes.

Periodic SAF-TE and SES Device Check Time

If there are remote devices within your RAID enclosure monitored by SAF-TE or SES, use this function to decide at what interval the controller checks the status of these devices.



Caution – Do not set this interval for less than one second. Setting it to less than one second can adversely impact reliability.

▼ To Set the Periodic SAF-TE and SES Device Check Time

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → Periodic SAF-TE and SES Device Check Time.”

A list of intervals is displayed.

2. Choose the interval you want.

A confirmation message is displayed.

3. Choose Yes.

Periodic Auto-Detect Failure Drive Swap Check Time

This menu option periodically polls the unit to detect a replacement of a bad drive. If no spare drive is present within the array, the logical drive begins an automatic rebuild of a degraded RAID set if the firmware detects a bad drive replacement.

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 detects the failed drive (at the specified time interval). Once the failed drive has been swapped with a drive that has adequate capacity to rebuild the logical drive, the rebuild begins automatically.

The default setting is "Disabled," meaning that the controller does not auto-detect the swap of a failed drive. When "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 detects drive removal only when a host attempts to access the data on the drive.

Note – This feature requires system resources and can impact performance.

▼ To Set the Auto-Detect Failure Drive Swap Check Time

1. Choose "view and edit Configuration parameters → Drive-side SCSI Parameters → Periodic Auto-Detect Failure Drive Swap Check Time."

A list of intervals is displayed.

2. Choose the interval you want.

A confirmation message is displayed.

3. Choose Yes 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.

Drive Predictable Failure Mode (SMART)

Use this menu option to enable or disable SMART functionality. See [“Understanding SMART Technology” on page 117](#) for a detailed description of this functionality. See [“Enabling SMART From Firmware Menus” on page 118](#) for information about how to configure your Drive Predictable Failure Mode setting.

Auto-Assign Global Spare Drive

This feature is Disabled by default.

When you enable the “Auto-Assign Global Spare Drive” menu option, the system automatically assigns a global spare to the minimum drive ID in unused drives. This enables the Fibre Channel array to rebuild automatically without user intervention if a drive is replaced.

▼ To Automatically Assign Replacements to Faulty Drives

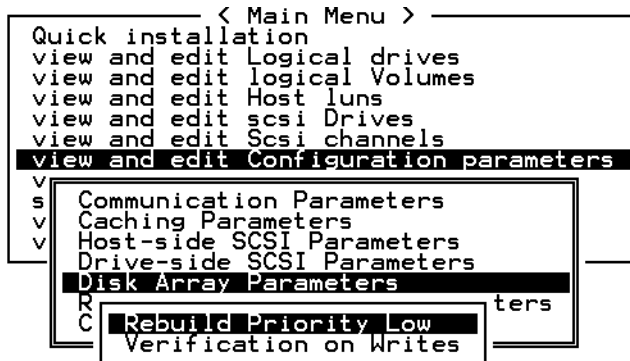
1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → Auto-Assign Global Spare Drive.”
2. When the prompt *Enable Auto-Assign Global Spare?* appears, select Yes.
As soon as a faulty drive is replaced, the replacement drive is identified as a global spare drive.

Disk Array Parameters Menu

The menu options on the Disk Array Parameters menu are described in this section.

▼ To Display the Disk Array Parameters Menu

1. Choose “view and edit Configuration parameters → Disk Array Parameters.”



Setting 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 depends largely on the total capacity of the logical drive being rebuilt. Additionally, the rebuilding process is totally transparent to the host computer and its operating system.

▼ To Set the Rebuild Priority


1. Choose “view and edit Configuration parameters → Disk Array Parameters → Rebuild Priority.”

A list of the priority selections is displayed:

- Low (the default that 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)

2. Choose the desired setting.



Verification on Writes

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

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

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

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

▼ To Enable and Disable Verification Methods

1. Choose “view and edit Configuration parameters → Disk Array Parameters → Verification on Writes.”

The verification methods available are 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
v
s
v
v
v
Host
Driv
Disk Verification on LD Initialization Writes Disabled
R Verification on LD Rebuild Writes Disabled
C Verification on Normal Drive Writes Disabled
R
C Verification on Writes
```

2. Choose the method you want to enable or disable.

A confirmation message 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
v
s
v
v
v
Host
Driv
Disk Verification on LD Initialization Writes Disabled
R V Enable Initialize RAID with Verify Data ?
C Yes No
R
C Veri
```

3. Choose Yes.

Note – Follow the same procedure to enable or disable each method.

Redundant Controller Parameters Menu (Reserved)

The Redundant Controller Parameters menu options are:

- “Secondary Controller RS-232C”
- “Remote Redundant Controller”

Do not use these menu options. They are reserved for specific troubleshooting procedures and should be used by only qualified technicians.

▼ To Display the Redundant Controller Parameters Menu (Reserved)

1. Choose “view and edit Configuration parameters → Redundant Controller Parameters.”

The “Redundant Controller Parameters” menu options are displayed.

Controller Parameters

Procedures for viewing and displaying controller parameters are described in this section.

Controller Name

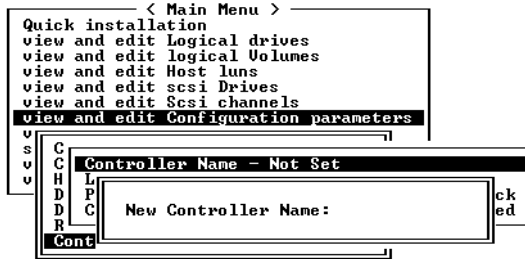
The controller name is displayed only in the firmware program and is used to identify separate controllers.

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.

▼ To View and Display the Controller Name

1. Choose “view and edit Configuration parameters → Controller Parameters → Controller Name.”

A text area is displayed where you can type in a controller name. Depending on the controller’s current settings, you are prompted to either enter a new name or modify the existing name for the designated controller.



2. Type a name for the controller and press Return.

LCD Title Display - Controller Logo (Reserved)

This function is not applicable to this product.

Password Validation Timeout

This function sets a timeout when a password.

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 “Always Check” default value 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 can continue to retry a password until the preset timeout expires, unless the default “Always Check” value is chosen. 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 provides immediate access to firmware menu options, whether or not a password has been established.

Controller Unique Identifier (Reserved)

The controller unique identifier is automatically set by the SAF-TE or SES device. The controller unique identifier is used to create Ethernet addresses and WWNs, and to identify the unit for some network configurations.



Caution – Do not change the controller unique identifier unless instructed to do so by qualified service personnel.



Caution – 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 might become inaccessible.

▼ To Set the Controller Unique Identifier

1. Choose “view and edit Configuration parameters → Controller Parameters → Controller Unique Identifier <hex>.”
2. 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.



Caution – A nonzero 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.

3. To implement the revised parameter value, choose “system Functions → Reset Controller.”

SDRAM ECC Function (Reserved)

The default is set to enabled. Do not use this setting. It is reserved for specific troubleshooting methods and should be used only by qualified technicians.

Controller Failover

Some 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 “Bus Reset Issued” warning message is displayed for each of the channels. In addition, a “Redundant Controller Failure Detected” alert message is displayed.

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 assumes the role of the secondary controller.

The failover and failback processes are completely transparent to the host.

Controllers are hot-swappable if you are using a redundant configuration. Replacing a failed unit takes only a few minutes. Since the I/O connections are on the controllers, you might experience some unavailability between the time when the failed controller is removed and the time when a new one is installed in its place.

To maintain your redundant controller configuration, replace the failed controller as soon as possible. For details, refer to *Sun StorEdge 3000 Family FRU Installation Guide*.

Rebuilding Logical Drives

This section describes automatic and manual procedures for rebuilding logical drives.

Automatic Logical Drive Rebuild

Rebuild with Spare. When a member drive in a logical drive fails, the controller first examines whether there is a local spare drive assigned to this logical drive. If yes, it automatically starts to rebuild the data of the failed disk to it.

If there is no local spare available, the controller searches for a global spare. If there is a global spare, it automatically uses 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 does not attempt to rebuild unless you apply a forced-manual rebuild.

To enable this feature, go to the Main Menu choose “view and edit Configuration parameters → Drive-side SCSI Parameters → 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 chosen), the controller detects whether the failed drive has been swapped by checking the failed drive's channel/ID. Once the failed drive has been swapped, the rebuild begins immediately.

Note – This feature requires system resources and can impact performance.

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

For a flowchart of automatic rebuild, see [FIGURE 8-1](#).

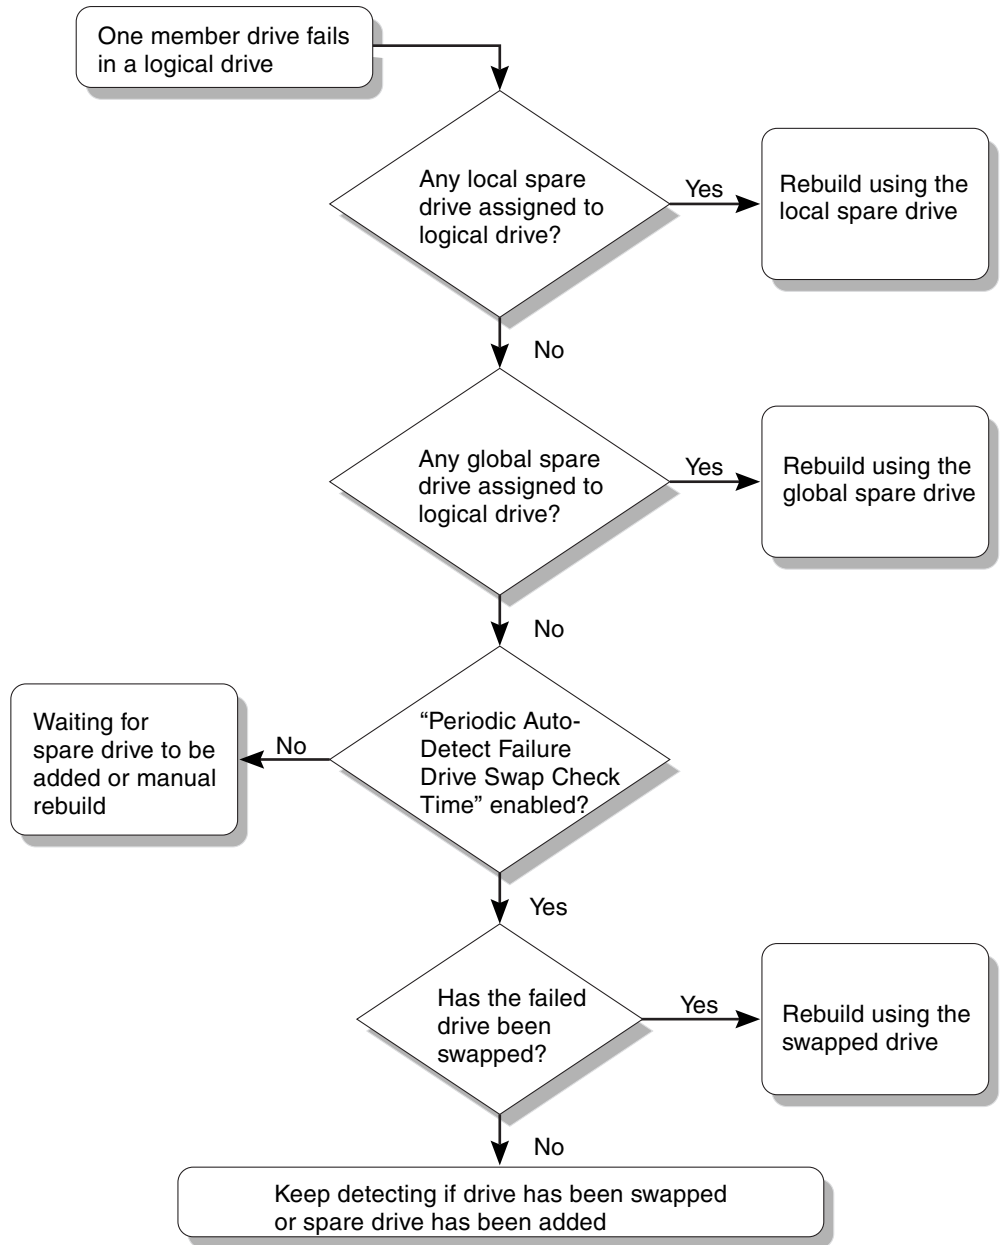


FIGURE 8-1 Automatic Rebuild

Manual Rebuild

When a user applies forced-manual rebuild, the controller first examines whether there is any local spare assigned to the logical drive. If a local spare is available, it automatically starts to rebuild.

If no local spare is available, the controller searches for a global spare. If there is a global spare, the logical drive rebuild begins. See [FIGURE 8-2](#).

If neither local spare nor global spare is available, the controller examines the 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 does not attempt to rebuild until the user applies another forced-manual rebuild.




FIGURE 8-2 Manual Rebuild

Concurrent Rebuild in RAID 1+0

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

Recovering From Fatal Drive Failure

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

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

It is rare for two or more drives to fail at the same time.

▼ To Recover From Fatal Drive Failure

1. **Discontinue all I/O activity immediately.**
2. **To cancel the beeping alarm, choose “system Functions → Mute beeper.”**
3. **Physically check whether all the drives are firmly seated in the array and make sure that none have been partially or completely removed.**
4. **Choose “view and edit Logical drives” from the Main Menu and look for:**
Status: FAILED DRV (one failed drive) or
Status: FATAL FAIL (two or more failed drives)
5. **Select the logical drive.**
6. **Choose “view scsi drives.”**
If two physical drives have a problem, one drive will have a BAD status and one drive will have a MISSING status. The MISSING status is a reminder that one of the drives might be a “false” failure. The status does not tell you which drive might be a false failure.
7. **Perform one of the following steps:**
 - Choose “system Functions → Reset controller.”
 - Power off the array, wait five seconds, and power on the array.
8. **Repeat Steps 4 and 5 to check the logical and drive status.**
After resetting the controller, if there is a false bad drive, the array automatically starts 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 188).
 - If the status is still “FATAL FAIL,” you have lost all data on the logical drive and must recreate the logical drive. Proceed with the following procedures.
- a. **Replace the failed drive. Refer to *Sun StorEdge 3000 Family FRU Installation Guide* for more information.**
 - b. **Delete the logical drive. See “[Deleting a Logical Drive](#)” on page 51 for more information.**
 - c. **Create a new logical drive. See “[Assigning a Logical Drive Name](#)” on page 54 for more information.**

For additional troubleshooting tips, refer to the *Sun StorEdge 3510 FC Array Release Notes* located at:

http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3510

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 193
- “Viewing SES Status (FC Only)” on page 194
 - “To Check the Status of SES Components (FC Only)” on page 195
 - “SES Temperature Sensor Locations (FC Only)” on page 199
 - “Identifying Fans” on page 197
 - “To View the Status of Each Fan” on page 198
- “Viewing Peripheral Device SAF-TE Status (SCSI Only)” on page 201
 - “SAF-TE Temperature Sensor Locations (SCSI Only)” on page 203
- “SAF-TE Temperature Sensor Locations (SCSI Only)” on page 203
 - “Redundant Controller Mode” on page 204
 - “Enabling UPS Status” on page 204
- “Setting the UPS Power Fail Signal” on page 205
 - “To Set the UPS Power Fail Signal (Reserved)” on page 205
 - “Setting the Temperature Exceeds Threshold Trigger” on page 205
 - “To Configure Over-Temperature Controller Shutdown” on page 205
- “Viewing Controller Voltage and Temperature Status” on page 206
- “Fibre Channel Error Statistics” on page 209

Viewing Peripheral Device Controller Status

To view the status of each controller, choose “view and edit Peripheral devices → View Peripheral Device Status.”

A table displays the status of the available peripheral devices.

```

    < 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
  S
  v
  v
  S
  D
  A
  C
  View Peripheral Device Status
  ITEM          STATUS          LOCATION
  Redundant Controller Enabled          Primary
  SAF-TE Device Operational      Channel 0 ID 14

```

FIGURE 9-1 View Peripheral Device Status

Viewing SES Status (FC Only)

A Fibre Channel array's SCSI Enclosure Services (SES) processor, located on the I/O module, monitors environmental conditions such as temperature sensors, cooling fans, the beeper speaker, power supplies, and slot status. The SES processor is supported by Sun StorEdge Configuration Service and the command-line interface.

For Sun StorEdge 3510 FC JBOD arrays only, both Sun StorEdge Configuration Service and the CLI access the SES processor using device files in `/dev/es`, such as `/dev/es/ses0`, as shown in the following example:

```

# sccli

Available devices:

1. /dev/rdisk/c4t0d0s2 [SUN StorEdge 3310 SN#000280] (Primary)

2. /dev/es/ses0 [SUN StorEdge 3510F D SN#00227B] (Enclosure)

```


▼ To Check the Status of SES Components (FC Only)

1. Choose “view and edit Peripheral devices → View Peripheral Device Status.”

A table displays the status of the available peripheral devices.

2. Choose “SES Device.”

```

< 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
  
```

```

S
D
A
C
F
  
```

ITEM	STATUS	LOCATION
Redundant Controller	Degraded	Primary
SES Device	Enclosure Device	Channel 2 ID 12
SES Device	Enclosure Device	Channel 2 ID 76

A list is displayed of environmental sensors and other hardware components of that SES device.

```

< Main Menu >
Quick installation
view and edit Logical drives
view and e
view an
view an En
view an De
view an Co
view an Te
view an Uo
  
```

```

S
D
A
C
F
  
```

Element Descriptor	STATUS	LOCATION
Overall Status	Degraded	Primary
Element 0		
Element 1		
Element 2		
Element 3		
Element 4		
Element 5		
Element 6		
Element 7		
Element 8		
Element 9		
SES Element 10	Enclosure Device	Channel 2 ID 12
SES Element 11	Enclosure Device	Channel 2 ID 76

3. Select an item from the list and press Return to display information about it or see a submenu of its component attributes.

```

< Main Menu >
Quick installation
-----
Status      :OK
Actual Speed:Fan at speed 4
-----
view an Device
view an Cooling element
view an Temperature
-----
s view an Uo Element Descriptor
v view an Po Overall Status
v view an Au Element 0
v view an D No Element 1
v view an A SC Element 2
v view an C R PB Element 3
v view an F
-----
SES Device Enclosure Device Channel 2 ID 12
SES Device Enclosure Device Channel 2 ID 76
-----

```

STATUS		LOCATION
Degraded		Primary
SES Device	Enclosure Device	Channel 2 ID 12
SES Device	Enclosure Device	Channel 2 ID 76

Choosing Overall Status in the illustration above displays the status of the SES device and its operating temperature.

Overall status of an SES device is reported independently from the status of the individual components of that device. An SES device showing an overall status in the menu has its own sensors that report its overall status and its overall temperature.

4. Select other attributes in which you are interested and press Return to learn more about the SES device.

Selecting the Element Descriptor in the following illustration displays the descriptive name of the element.

Quick installation		< Main Menu >	
view an	view and edit Logical drives		
view an	view and e	Element Descriptor	
view an	En Overall Status	Element 0	
view an	De Element	Element 1	Parameters
view an	Co Element	Element 2	ES
view an	Te Element	Element 3	
s	Uo Element	Element 4	us
u	Po Element	Element 5	
v	Au Element	Element 6	
S	No Element	Element 7	
D	SC Element	Element 8	
A	RPB Element	Element 9	
C	R Element	Element 10	
F	SES Element	Element 11	
	SES		
	SES		
		STATUS	LOCATION
		Degraded	Primary
		Enclosure Device	Channel 2 ID 12
		Enclosure Device	Channel 2 ID 76

In this case the descriptor is "Disk Drives."

Quick installation		< Main Menu >	
view an	view and edit Logical drives		
view an	view and e	Disk Drives	
view an	En Overall Status	Element 0	
view an	De Element	Element 1	Parameters
view an	Co Element	Element 2	ES
view an	Te Element	Element 3	
s	Uo Element	Element 4	us
u	Po Element	Element 5	
v	Au Element	Element 6	
S	No Element	Element 7	
D	SC Element	Element 8	
A	RPB Element	Element 9	
C	R Element	Element 10	
F	SES Element	Element 11	
	SES		
	SES		
		STATUS	LOCATION
		Degraded	Primary
		Enclosure Device	Channel 2 ID 12
		Enclosure Device	Channel 2 ID 76

Identifying Fans

You can view the status of SES components, including the pair of fans located in each fan and power supply module. A fan is identified by the SES Device menus as a cooling element.

Fan Status	Fan RPM
5 Fan at speed 5	5143 - 5428
6 Fan at speed at intermediate speed	5429 - 5713
7 Fan at highest speed	5714 to maximum

If a fan fails and the Status field does not display the OK value, you must replace the fan and power supply module.

Cooling elements in the status table can be identified for replacement as shown:

TABLE 2 Relationship Between Cooling Elements, Fans, and Power Supply Modules

Cooling Element #	Fan # and Power Supply Module #
Cooling element 0	FAN 0, PS 0
Cooling element 1	FAN 1, PS 0
Cooling element 2	FAN 2, PS 1
Cooling element 3	FAN 3, PS 1




FIGURE 9-3 Cooling Fan Locations

SES Temperature Sensor Locations (FC Only)

Monitoring temperature at different points within the array is one of the most important SES functions. High temperatures can cause significant damage if they go unnoticed. There are a number of different sensors at key points within the enclosure. The following table shows the location of each of those sensors. The

element ID corresponds to the identifier shown when you choose “view and edit Peripheral devices → View Peripheral Device Status → SES Device → Temperature Sensors.”

TABLE 9-1 Sun StorEdge 3510 Temperature Sensor Locations

Element ID	Description
0	Drive Midplane Left Temperature Sensor #1
1	Drive Midplane Left Temperature Sensor #2
2	Drive Midplane Center Temperature Sensor #3
3	Drive Midplane Center Temperature Sensor #4
4	Drive Midplane Right Temperature Sensor #5
5	Drive Midplane Right Temperature Sensor #6
6	Upper IOM Left Temperature Sensor #7
7	Upper IOM Left Temperature Sensor #8
8	Lower IOM Temperature Sensor #9
9	Lower IOM Temperature Sensor #10
10	Left PSU Temperature Sensor #11
11	Right PSU Temperature Sensor #12

Viewing Peripheral Device SAF-TE Status (SCSI Only)

A SCSI array's SAF-TE controller is located on the SCSI I/O module. It controls environmental monitoring SAF-TE devices such as temperature sensors, cooling fans, the beeper speaker, power supplies, and slot status.

▼ To Check the Status of SAF-TE Components (SCSI Only)

1. Choose “view and edit Peripheral devices → View Peripheral Device Status → SAF-TE Device.”

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 9-4](#).
- **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 9-5](#).

To check whether you have all slots filled in a dual-bus configuration, see “[Viewing the Physical Drive Status Table](#)” on [page 98](#) and check the column labeled “Chl 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 9-4 Example of SAF-TE Device Status Window for a Single-Bus Configuration

In a dual-bus configuration example, the following SAF-TE window displays “No Device Inserted” for six drives that are actually inserted into slots. The SAF-TE protocol does not support a dual-bus configuration and recognizes only 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 9-5 Example of SAF-TE Device Status Window for a Dual-Bus Configuration

SAF-TE Temperature Sensor Locations (SCSI Only)

Monitoring temperature at different points within the array is one of the most important SAF-TE functions. High temperatures can cause significant damage if they go unnoticed. There are a number of different sensors at key points within the enclosure. The following table shows the location of each of those sensors. The Element ID corresponds to the identifier shown when you choose “view and edit Peripheral devices → View Peripheral Device Status → SAF-TE Device.”

TABLE 9-2 Sun StorEdge 3310 Temperature Sensor Locations

Temp Sensor ID	Description
0	Port A Drive Midplane Temperature #1
1	Port A Drive Midplane Temperature #2
2	Port A Power Supply Temperature #1 (PS 0)
3	Port B EMU Temperature #1 (left module as seen from back)
4	Port B EMU Temperature #2 (right module as seen from back)
5	Port B Drive Midplane Temperature #3
6	Port B Power Supply Temperature #2 (PS 1)
CPU Temperature	CPU on Controller
Board1 Temperature	Controller
Board2 Temperature	Controller

Setting Peripheral Device Entry

The Set Peripheral Device Entry menu options include the following:

- Redundant Controller
- UPS Status
- Event Trigger Operations

Redundant Controller Mode

The redundant controller mode is automatically enabled. Do not change this setting.

For more information about redundant controller operation, refer to [“Controller Defaults and Limitations”](#) on page 18.

Note – Dual independent controllers are sometimes used in high-performance situations where data integrity is not crucial and nonredundant is appropriate.



Caution – In single-controller configurations, do not disable the Redundant Controller setting and do not set the controller as a secondary controller. The primary controller controls all firmware operations and must be the assignment of the single controller. If you disable the Redundant Controller function and reconfigure the controller with the Autoconfigure option or as a secondary controller, the controller module will become inoperable and will need to be replaced.


Enabling 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.”

▼ To Enable the UPS Status

1. Choose “UPS Status.”

A confirmation prompt is displayed.



2. Choose **Yes to confirm.**

Setting the UPS Power Fail Signal

The “UPS Power Fail Signal” menu option sets the priority of the alert notification level if power to your UPS device should fail for any reason. Do not change the default High priority.

▼ To Set the UPS Power Fail Signal (Reserved)

1. Choose “**view and edit Peripheral devices → Define Peripheral Device Active Signal → UPS Power Fail Signal.**”

A confirmation prompt is displayed.

Choose Yes to change the setting.

Setting the Temperature Exceeds Threshold Trigger

Use the “Temperature exceeds threshold” menu option to force a controller shutdown if a temperature is detected that exceeds system threshold limits. You can adjust this setting to shut down the controller as soon as the temperature limit is exceeded, or after a configurable delay.

▼ To Configure Over-Temperature Controller Shutdown

1. Choose “**view and edit Peripheral devices → Set Peripheral Device Entry → Event Trigger Operations → Temperature exceeds threshold.**”

A menu of options and shutdown delay intervals is displayed.

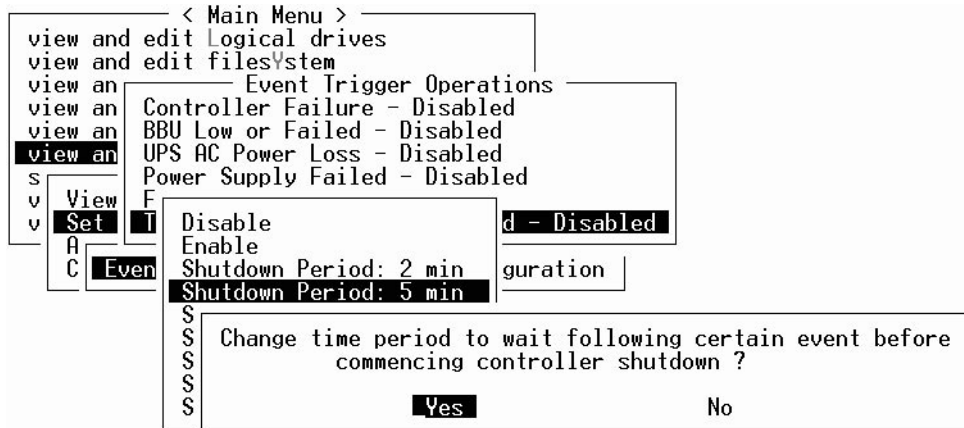
2. Choose **one of the following options.**

- **Disable:** Disable controller shutdown if the temperature exceeds limits.

- **Enable:** Immediately shut the controller down if the temperature exceeds limits.

- **Shutdown Period:** Shut the controller down after the specified delay interval if the temperature exceeds limits.

A confirmation prompt is displayed.



3. Choose Yes.

Viewing Controller Voltage and Temperature Status

This section describes how to see whether a controller's voltage and temperature are within normal ranges.

▼ To Check the Controller Voltage and Temperature Status

1. Choose "view and edit Peripheral devices → Controller Peripheral Device Configuration → View Peripheral Device Status."

- Repeat this step as many times as necessary to “drill down” to the threshold ranges and triggering events.

```

< 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
S
U
V
View
Set
Defi
Adj
Cont
F
U
Trigger Thresholds for +3.30 Events
Upper Threshold for +3.30 Event - Default<3.60>
Lower Threshold for +3.30 Event - Default<2.90>
Voltage and Temperature Parameters

```

```

< 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
S
U
V
View
Set
Defi
Adj
Cont
F
U
Voltage Range from 3.40 to 3.90
Default Trigger Event: "default"
Disable Trigger Event: "disable"
Input Voltage Trigger Threshold : default
Trigger
Upper Threshold for +3.30 Event - Default<3.60>
Lower Threshold for +3.30 Event - Default<2.90>
Voltage and Temperature Parameters

```

- If you want to edit a trigger or other editable value, backspace over the existing information and replace it, as shown in the following sequence.

```

< 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
s
v
v
View      Trig
Set      Trig
Defi     Trig
Adju     Trig
Cont     Trig
F
U
Uolt     Upper Threshold for CPU Temperature Event - 95 <C>
Lower Threshold for CPU Temperature Event - Default< 0<C> >

```

```

< 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
s
v
v
View      Trig
Set      Trig
Defi     Trig
Adju     Trig
Cont     Trig
F
U
Uolt     Upper Threshold for CPU Temperature Event - 95 <C>
Lower Threshold for CPU Temperature Event - Default< 0<C> >

```

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit
view and edit
view and edit
view and edit
view and edit
view and edit
s
v
v
View      Trig
Set      Trig
Defi     Trig
Adju     Trig
Cont     Trig
F
U
Uolt     Upper Threshold for Board Temperature Event - 85 <C>
Lower Threshold for Board Temperature Event - Default< 0<C> >
Trig Therholds for Board Temperature Events
Uolt     Voltage and Temperature Parameters

```

Fibre Channel Error Statistics

You can view FC error statistics that indicate the status of loopback operation on local channels and on drives.

The statistics are provided under the following headings:

- CH/ID – Channel number of the Fibre Channel Port from which error information is obtained.

- LIP – Total number of loop initializations that have occurred on the channel.
- LinkFail – Total number of instances of link failures. This hardware counter is a sum of the remainder of the following counters.
- LossOfSync – Total number of instances of loss of sync. This is the number of times the Qlogic Fibre Channel chip failed to receive the proper comma character within a primitive three times.
- LossOfSignal – Total number of instances of loss of signal.
- PrimErr – Total number of instances of primitive sequence protocol errors.
- InvalTXWord – Total number of instances of invalid transmission words. This error indicates either an invalid transmit word or disparity error.
- InvalCRC – Total number of instances of invalid CRC, or the number of times a frame was received and the CRC was not as expected.

To check the Fibre Channel Error statistics, perform the following steps.

1. Choose “view and edit Peripheral devices → Fibre Channel Error Statistics → Local Channel Statistics.”

Local channel statistics are displayed.

CH/ID	LIP	LinkFail	LossOfSy	LossOfSi	PrimErr	InvalTxW	InvalCRC
0/28	2	0	0	0	0	4	0
1/2A	2	0	0	0	0	0	0
2/7C	4	0	0	0	0	0	0
3/7C	4	0	0	0	0	0	0
4/2C	1	0	0	0	0	3	0
5/2E	2	0	0	0	0	0	0

2. Choose “Drive Side Device Statistics.”

Drive-side device statistics are displayed.

CH/ID	LIP	LinkFail	LossOfSy	LossOfSi	PrimErr	InvalTxW	InvalCRC
2/7C	4	0	0	0	0	0	0
2/0B	4	0	1E	0	0	722	0
2/06	4	0	6	0	0	2A2	0
2/0A	4	0	1	0	0	2FA	0
2/07	4	0	4	0	0	2E1	0
2/0C	4	0	0	0	0	0	0
2/41	4	0	1D	0	0	A2A	0
2/44	4	0	9	0	0	3B4	0
2/42	4	0	1	0	0	2FE	1
2/43	4	0	3	0	0	2E2	0
2/45	4	0	1	0	0	202	0
2/40	4	0	1	0	0	2A2	0
2/48	4	0	1	0	0	20F	1
2/49	4	0	9	0	0	2F5	0
2/4D	4	0	4	0	0	307	0
2/46	4	0	1	0	0	1A8	1

System Functions and Event Logs

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

- [“Muting the Beeper” on page 211](#)
 - [“To Change the Beeper Setting” on page 212](#)
- [“Setting and Changing the Controller Password” on page 212](#)
 - [“To Set a New Password” on page 213](#)
 - [“To Disable the Password” on page 214](#)
- [“Resetting the Controller” on page 214](#)
- [“Shutting Down the Controller” on page 215](#)
 - [“To Shut Down a Controller” on page 216](#)
- [“Saving Your Configuration \(NVRAM\) to a File” on page 216](#)
 - [“To Save Your Configuration \(NVRAM\) to a File” on page 219](#)
- [“Restoring Your Configuration \(NVRAM\) From a File” on page 219](#)
- [“Restoring Your Configuration \(NVRAM\) From a File” on page 219](#)
 - [“To View the Event Logs” on page 221](#)

Muting the Beeper

An audible alarm indicates that either a component in the array has failed or a specific controller event has occurred. Error conditions and controller events are reported by event messages and event logs. Component failures are also indicated by LED activity on the array.

For information about failed component alarms, refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array. For information on controller events, see [“Event Messages” on page 241](#).

Note – It is important to know the cause of the error condition because how you silence the alarm depends on the cause of the alarm.


If the alarm is caused by a controller event, use the “Mute beeper” menu option to disable the system alarm until another error event occurs. If the alarm is caused by a failed component, push the Reset button on the right ear of the array.

Note – If the alarm is caused by a failed component, using the “Mute beeper” menu option has no effect.

▼ To Change the Beeper Setting

1. Choose “system Functions → Mute beeper.”

A confirmation message is displayed.



2. Choose Yes to turn the beeper off until the next event occurs.

Setting and Changing the Controller Password

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

Note – The controller is able to verify the password when entering the Main Menu from the initial screen or making configuration changes. If the controller will be left unattended, “Password Validation Timeout” can be set to “Always Check.” Setting validation timeout to “Always Check” protects the controller configuration from any unauthorized change.

Note – 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.

▼ To Set a New Password

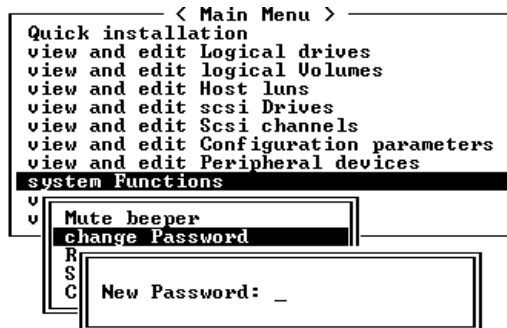
1. Choose “system Functions → change Password.”

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

2. Type the old password (if any) and press Return.

If the password is incorrect, the controller will not allow you to change the password. Instead, it will display a “Password incorrect!” error message and return to the previous menu.

If the password is correct, or there is no preset password, you are prompted for the new password.



Note – The controller password is case-sensitive.

3. **Type the desired password in the column and press Return.**

You are prompted to type the new password again.

4. **Type the password again to confirm and press Return.**

The new password now becomes the controller's password.

▼ To Disable the Password

To disable or delete the password, press Return without typing anything in the text area used for entering a new password. The existing password is deleted and no password checking occurs when entering the Main Menu from the initial screen.

Resetting the Controller

It is sometimes necessary after changing controller parameters to reset the controller before the parameter changes can take effect. However, there are two ways of resetting a controller from the firmware application, the "Reset Controller" menu option and the "Shutdown Controller" menu option. It is important that you distinguish between the results of these two menu options.

Use the "Reset Controller" menu option to reset the controller without saving the contents of the controller's cache to disk. This can be desirable if you believe that a software crash or hardware fault might have corrupted the cached data.

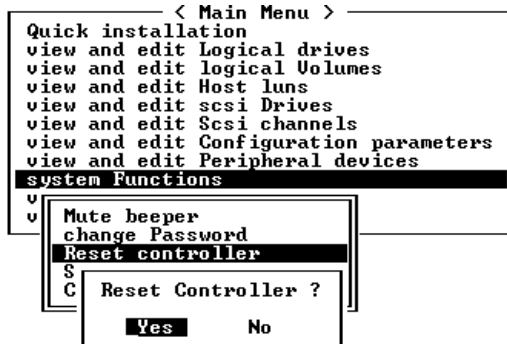


Caution – If you want to write the cache contents to disk, do not use the "Reset Controller" menu option while the controller is connected to host systems. Instead, use the "Shutdown Controller" menu option and choose Yes when you see the "Reset Controller?" prompt. See ["Shutting Down the Controller" on page 215](#) for more information.

▼ To Reset the Controller Without Saving Cache Contents

1. Choose “system Functions → Reset Controller.”

A confirmation prompt is displayed.



2. Choose Yes.

The controller is now reset.



Caution – Resetting the controller does not preserve the contents of the cache or write them to disk. When the controller is reset, all cache contents are lost. See [“Shutting Down the Controller” on page 215](#) for information about writing the contents of cache to disk before resetting the controller.

Shutting Down the Controller

The “Shutdown Controller” menu option first halts all I/O activity. It should be used only when all I/O activity from hosts to and from the controller has already been halted.

The “Shutdown Controller” menu option writes the contents of the cache to the drives.

Note – If you want to restart the controller after shutdown, confirm when prompted that you want to reset the controller after the shutdown function has completed.

▼ To Shut Down a Controller

1. Choose “system Functions → Shutdown Controller.”

A confirmation prompt is displayed.

2. Choose Yes.

A status and confirmation message tells you that the controller shutdown is complete and asks if you want to reset the controller.

```
< 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
view
view
svst
v
v M
v C
R
Shutdown controller
Controller maintenance
```

```
**** Shutdown Controller Completed ****
Power off Controller or Reset Controller

Reset Controller ?
Yes          No
```

3. Choose Yes to reset the controller.

Note – If you choose No, you lose access to the controller and must power it up manually or use the CLI to restart it.

Saving Your Configuration (NVRAM) to a File

Back up your controller-dependent configuration information. Use the “save nvram to disks” function to save configuration information whenever a configuration change is made, and record the configuration information.

When you save your configuration, it is stored in a logical drive.

Note – A logical drive must exist for the controller to write NVRAM content onto it.

Saving your NVRAM controller configuration to a file provides a backup of the controller-dependent configuration information such as channel settings, host IDs, FC protocol, and cache configuration. It does not save LUN mapping information. The NVRAM configuration file can restore all configuration settings but does not rebuild logical drives.

The firmware menu options whose parameter settings are saved when you save your NVRAM controller configuration to disk include:

- logical drive Assignments
- logical volume Assignments
- view and edit Host luns
- view and edit Scsi channels
- Baud-rate 38,400
- Data Routing Direct to Port
- Terminal Emulation Enabled
- Internet Protocol (TCP/IP)
- Write-Back Cache
- Optimization for Sequential or Random I/O
- Maximum Queued I/O Count
- Luns per Host SCSI ID
- Max Number of Concurrent Host-LUN Connections
- Peripheral Device Type
- Peripheral Device Qualifier
- Device Supports Removable Media
- LUN Applicability
- Host Cylinder/Head/Sector Mapping Configuration
- Head Ranges
- Cylinder Ranges
- Fibre Connection Option
- SCSI Motor Spin-Up
- SCSI Reset at Power-Up
- Disk Access Delay Time
- SCSI I/O Timeout
- Maximum Tag Count
- Periodic Drive Check Time
- Periodic SAF-TE and SES Device Check Time
- Periodic Auto-Detect Failure Drive Swap Check Time
- Auto-Assign Global Spare Drive
- Rebuild Priority
- Verification on LD Initialization Writes
- Remote Redundant Controller
- Controller Name
- LCD Title Display
- Password Validation Timeout
- SDRAM ECC
- change Password

The firmware menu options whose parameter settings are not saved when you save your NVRAM controller configuration to disk include:

- Delete logical drive
- Partition logical drive
- logical drive Name
- Delete logical volume
- Partition logical volume
- Edit Host-ID/WWN Name List
- disk Reserved space
- Global spare
- PPP Configuration
- Modem Operation
- SNMP Configuration
- Controller Unique Identifier (Hex)
- UPS Status
- UPS Power Fail Signal Active
- View Peripheral Device Status
- Trigger Thresholds for +3.3V Events
- Upper Threshold for +3.3V Event
- Lower Threshold for +3.3V Event
- Trigger Thresholds for +5V Events
- Upper Threshold for +5V Event
- Lower Threshold for +5V Event
- Trigger Thresholds for +12V Events
- Upper Threshold for +12V Event
- Lower Threshold for +12V Event
- Trigger Thresholds for CPU Temperature Events
- Upper Threshold for CPU Temperature Event
- Lower Threshold for CPU Temperature Event
- Trigger Thresholds for Board Temperature Events
- Upper Threshold for Board Temperature Event
- Lower Threshold for Board Temperature Event

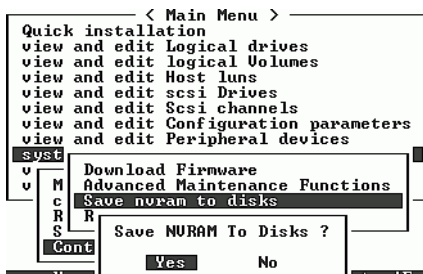
If you prefer to save and restore all configuration data, including LUN mapping information, use Sun StorEdge Configuration Service and CLI software in addition to saving your NVRAM controller configuration to disk. The information saved this way can be used to rebuild all logical drives and therefore can be used to completely duplicate an array configuration to another array.

Refer to the *Sun StorEdge 3000 Family Configuration Service User's Guide* for information about the "save configuration" and "load configuration" features. Refer to the `sccli` man page for information about the "reset nvram" and "download controller-configuration" commands.

▼ To Save Your Configuration (NVRAM) to a File

1. Choose “system Functions → Controller maintenance → Save nvram to disks.”

A confirmation prompt is displayed.



2. Choose Yes to confirm.

A message confirms that the NVRAM information has been successfully saved.

To restore the configuration, refer to [“Restoring Your Configuration \(NVRAM\) From a File” on page 219.](#)

Restoring Your Configuration (NVRAM) From a File

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

The NVRAM configuration file restores all configuration settings (channel settings, host IDs, etc.) but does not rebuild logical drives. See [“Saving Your Configuration \(NVRAM\) to a File” on page 216](#) for information about how to save a configuration file, including advice about saving controller-dependent configuration whenever a configuration change is made.



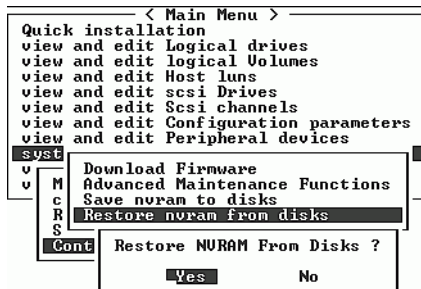
Caution – Before restoring a configuration file, be certain that the configuration file you apply matches the array to which you apply it. If host IDs, logical drive controller assignments, or other controller-dependent configuration information has changed since the configuration file was saved, you might lose access to mismatched channels or drives. You will have to change cabling or host or drive channel IDs to correct this mismatch and restore the access you have lost. The address of the RAID controller channel must also match what is described in `/etc/vfstab`.

Note – In the Sun StorEdge 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 the logical drives. Only perform this operation if no data has been stored or all data has been transferred to another array.

▼ To Restore Configuration Settings from a Saved NVRAM File

1. Choose “system Functions → Controller maintenance → Restore nvram from disks.”

A confirmation message is displayed.



2. Press Yes to confirm.

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

Viewing Event Logs

The controller event log records an event or alarm that occurs after the system is powered on. The controller can store up to 1,000 event logs. An event log records a configuration or operation event as well as an error message or alarm event. The SES log in each FC array and Event Monitoring Units in each SCSI array send messages to the controller log that report problems and status of the fans, temperature, and voltage.



Caution – If you power off or reset the controller, all recorded event logs are deleted.

▼ To View the Event Logs

1. Choose “view and edit Event logs” from the Main Menu.

A log of recent events is displayed.

Event Logs	
[1106] CHL:2 ID:74 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Wed Dec 25 16:25:49 2002	P
[1106] CHL:2 ID:0 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Thu Dec 26 17:47:26 2002	P
[1106] CHL:2 ID:4 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Thu Dec 26 21:33:00 2002	P
[1106] CHL:2 ID:2 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Thu Dec 26 22:54:17 2002	P
[1106] CHL:2 ID:0 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Fri Dec 27 12:00:35 2002	P
[1106] CHL:2 ID:72 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Sat Dec 28 14:56:19 2002	P
[1106] CHL:2 ID:2 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Sat Dec 28 16:11:32 2002	P
[1106] CHL:2 ID:70 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Sat Dec 28 17:52:19 2002	P

2. Use your arrow keys to scroll up and down through the list.

- To clear the events from the log once you've read them, use your arrow keys to move down to the last event you want to clear and press Return.

A "Clear Above xx Event Logs?" confirmation message is displayed.

Event Logs		
[1106] CHL:2 ID:74 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
[1]	Clear Above 14 Event Logs ?	ERT: SCSI Parity/CRC Error Detected
[1]	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	ERT: SCSI Parity/CRC Error Detected
[1106] CHL:2 ID:2 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Thu Dec 26 22:54:17 2002		
[1106] CHL:2 ID:0 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Fri Dec 27 12:00:35 2002		
[1106] CHL:2 ID:72 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Sat Dec 28 14:56:19 2002		
[1106] CHL:2 ID:2 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Sat Dec 28 16:11:32 2002		
[1106] CHL:2 ID:70 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Sat Dec 28 17:52:19 2002		

- Choose Yes to clear the recorded event logs.

Note – Resetting the controller clears the recorded event logs. If you want event logs to persist after controller resets, refer to *Sun StorEdge 3000 Family Configuration Service User's Guide* for information about how to install and configure the Sun StorEdge Configuration Service software.

Firmware Specifications

This appendix contains the firmware specifications in the following tables:

- “Basic RAID Management” on page 223
- “Advanced Features” on page 224
- “Caching Operation” on page 225
- “RAID Expansion” on page 226
- “Redundant Controller” on page 226
- “Data Safety” on page 227
- “Security” on page 227
- “Environment Management” on page 227
- “User Interface” on page 228

Note – Refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your SCSI or Fibre Channel array for up-to-date information about specific firmware features and descriptions that apply to your array.

TABLE A-1 Basic RAID Management

Feature	Description
RAID Levels	0, 1, 1+0, 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 coexist in an array.
Maximum Drive Number of Each Logical Drive	31 (RAID 3 or 5); 45 (RAID 0); 44 (RAID 1).
Logical Drive Identification	Unique, controller-generated logical drive ID; logical drive name is user-configurable.
Maximum Partitions per Array	128 for SCSI arrays, 1028 for FC arrays.

TABLE A-1 Basic RAID Management (Continued)

Feature	Description
Maximum Number of Logical Drives in a Logical Volume	8
Maximum Number of Luns 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 is available for all logical drives.
Auto-Rebuild Onto Spare Drive	Supported.
Auto-Scan of Replacement Drive Upon Manually Initiated Rebuild	Supported.
One-Step Rebuild Onto Replacement Drive	
Auto-Rebuild Onto Failed Drive Replacement	Supported. With no spare drive assigned, the controller auto-scans the failed drive and starts to rebuild automatically once the failed drive has been replaced.
Auto Recovery From Logical Drive Failure	Supported. When a user accidentally removes the wrong drive to cause the second drive failure of a one-drive-failed RAID 5 or RAID 3 logical drive, switches off the controller, puts the drive back and powers on the controller. The logical drive is restored to one-drive-failed status.

TABLE A-2 Advanced Features

Feature	Description
Drive Low-Level Format	Supported.
Drive Identification	Supported. Forces the drive to light the activity indicator so the user can recognize the correct drive.
Drive Information Listing	Supported.
Drive Read/Write Testing	Supported.
Configuration on Disk	Supported. The logical drive information is recorded on drive media.

TABLE A-2 Advanced Features (*Continued*)

Feature	Description
Save and Restore NVRAM to and From Disks	Supported. Saves all the settings stored in the controller NVRAM to the logical drive members.
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 sends spin-up (start unit) command to each drive at four second 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.
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-3 Caching Operation

Feature	Description
Write-Back and Write-Through Cache	Supported.
Supported Memory Type	SDRAM memory for enhanced performance. Fast Page Memory with Parity for enhanced data security.
Scatter / Gather	Supported.
I/O Sorting	Supported. Optimized I/O sorting for enhanced performance.
Variable Stripe Size	RAID 5: Optimization for random I/O (32 Kbyte), optimization for sequential I/O (128 Kbyte), user selectable. RAID 3: Optimization for random I/O (4 Kbyte), optimization for sequential I/O (16 Kbyte), user selectable.

TABLE A-4 RAID Expansion

Feature	Description
On-Line RAID Expansion	Supported.
RAID Expansion - Add Drive	Supported. Multiple drives can be added concurrently.
RAID Expansion – Copy and Replace Drives	Supported. Replace members with drives of larger capacity.

TABLE A-5 Redundant Controller

Feature	Description
Active-Active Redundant Controller	Supported.
Synchronized Cache for Both Controllers	Supported.
Write-back Cache Enabled in Redundant Controller Mode	Yes; with synchronized cache connection between controllers.
Hot-Swappable Controller	Supported.
No Single-Point-of-Failure	Supported.
Dynamic Cache Memory Allocation	Yes. Cache memory is dynamically allocated, not fixed.
Cache Battery Backup	Supported.
Load-Sharing	Supported. Workload can be flexibly divided between different controllers by assigning logical drives to different controllers.
User-Configurable Channel Mode	Supported. Channel modes are configurable 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.
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 autosynchronized later.

TABLE A-6 Data Safety

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 blocks.
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 SMART Support	Supported. Default: Disabled.
Clone Failing Drive	Users can choose to clone data from a failing drive to a backup drive manually.

TABLE A-7 Security

Feature	Description
Password Protection	Supported.
User-Configurable Password Validation Timeout	Supported. After a specific period of time without any user interaction, the password is requested again. This prevents unauthorized operation when the user is away.

TABLE A-8 Environment Management

Feature	Description
SAF-TE and SES Support	Supported.
SAF-TE and SES Polling Period	User-configurable (50 ms, 100 ms, 200 ms, 500 ms, 1~60 sec).
SAF-TE and SES Temperature Value Display	Supported. Displays the temperature value provided by enclosure SAF-TE module (if available).

TABLE A-8 Environment Management *(Continued)*

Feature	Description
On-Board Controller Voltage Monitors	Supported. Monitors the 3.3V, 5V, and 12V voltage status. Event trigger threshold user-configurable.
On-Board Controller Temperature Sensors	Supported. Monitors the CPU and board temperature status. Event trigger threshold user-configurable.
Enclosure Monitoring of Redundant Power Supply Status, Fan Status, UPS Status and Temperature Status	Supported. Fault-Bus, SAF-TE, SES, ISEMS.

TABLE A-9 User Interface

Feature	Description
RS-232C Terminal	Supports terminal modes: ANSI, VT-100, ANSI Color. Provides menu-driven, user-friendly, text-based interface.
Beeper Alarm	Warns user when any failure or critical event occurs.

Parameter Summary Tables

This appendix lists the firmware parameter settings for both the FC and SCSI arrays. You can fine-tune your array by making changes to these settings. It also lists the parameter defaults that you should not change unless advised to do so by technical support.

Topics covered include:

- [“Introducing Default Parameters” on page 229](#)
- [“Basic Default Parameters” on page 230](#)
- [“Default Configuration Parameters” on page 231](#)
- [“Default Peripheral Device Parameters” on page 238](#)
- [“Default System Functions” on page 239](#)
- [“Keeping Specific Parameter Defaults” on page 240](#)

Introducing Default Parameters

Although the factory defaults provide optimized controller operation, you might want to fine-tune your array through minor modifications to the following parameters listed.

Some parameters can be changed only at the initial array configuration; other parameters can be changed at any time. Note the listing of parameters defaults that should not be changed unless directed by technical support.

Note – For customer convenience, these tables apply to both the Sun StorEdge 3310 SCSI array, using firmware version 3.25, and the Sun StorEdge 3510 FC array, using firmware version 3.27.

Basic Default Parameters

These parameters are the primary settings for each array.

TABLE B-1 Logical Drive Parameters (View and Edit Logical Drives)

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Create Logical Drives	<ul style="list-style-type: none"> • 1–2 per array. • 1–2 spares per array. • 1–8 drives per array. 		
Change a Logical Drive Controller Assignment	Primary.		Secondary.

TABLE B-2 Logical Volume Parameters (View and Edit Logical Volumes)

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Create a Logical Volume	Primary controller.		Secondary.
Auto-Assign Global Spare Drive	Disabled.	Disabled.	Enabled. Disabled.

TABLE B-3 Host LUN Parameters (View and Edit Host LUNs)

User-Defined Parameter	FC Default Setting	SCSI Default Settings
Host LUN IDs	16 IDs per channel maximum in loop mode; 1 ID per channel in point-to-point mode. Channel 0 ID 40 - primary. Channel 1 ID 42 - secondary. Channel 4 ID 44 - primary. Channel 5 ID 46 - secondary.	2 IDs per channel maximum. Channel 1 ID 0 - primary. Channel 1 ID NA- secondary. Channel 3 ID NA- primary. Channel 3 ID 1- secondary.

TABLE B-4 SCSI Drive Parameters (View and Edit SCSI Drives)

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
FC Drive ID Switch Settings	0		0–7

TABLE B-5 SCSI Channel Parameters (View and Edit SCSI Channels)

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Host Channel Settings	0, 1, 4, 5 Host Channels.	1, 3 Host Channels.	Vary by product.
Drive Channel Settings	2 and 3	0 and 2	Vary by product.
Sync Transfer Clock	80 MHz	80 MHz	Do not change this parameter. 2.5 MHz to 80 MHz and Async.
Wide Transfer	Enabled.	Enabled.	Do not change this parameter. Enabled. Disabled.
Parity Check	Enabled.	Enabled.	Do not change this parameter. Disabled. Enabled.

Default Configuration Parameters

The most important parameter to review are the caching parameters, which impact the block size and optimization performance. Many parameters are optional or unused, depending on the applicable product.

The parameters in the following tables can be set using the View and Edit Configuration Parameters menu:

- [“Communication Parameters” on page 232](#)
- [“Caching Parameters” on page 233](#)
- [“Peripheral Device Type Parameters” on page 234](#)
- [“Host-Side and Drive-Side Parameters” on page 234](#)
- [“Other Configuration Parameters” on page 236](#)

TABLE B-6 Communication Parameters

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Communication Parameters > RS-232 Port Configuration			
Baud Rate	38,400	38,400	9,600, 19,200, 4,800, 2,400, 38,400
Data Routing	Direct to Port.	Direct to Port.	Point to Point (PPP). Direct to Port.
Terminal Emulation	Enabled.	Enabled.	Disabled. Enabled.
Communication Parameters > PPP Configuration			
Access Name	Not Set.	Not Set.	Type an access name. Not Set.
Access Password	Not Set.	Not Set.	Type an access password. Not Set.
Communication Parameters > Modem Operation > Modem Setup			
Configure Modem Port	Not configured.	Not configured.	Type a port.
Communication Parameters > Modem Operation > Modem Setup > Dial-out Function			
Dial-out Modem command	Not set.	Not Set.	Type a command.
Auto Dial-out on Initialization	Disabled.	Disabled.	Enabled. Disabled.
Dial-out Timeout (Seconds)	None.	None.	Type the number of seconds.
Dial-out Retry Count	2	2	Type a number.
Dial-out Retry Interval (Minutes)	5 minutes.	5 minutes.	Type a number.

TABLE B-6 Communication Parameters *(Continued)*

Dial-out Event Condition	Disabled.	Disabled.	Critical Events only. Critical Events and Warnings. All Events, Warnings, and Notifications. Disabled.
--------------------------	-----------	-----------	---

Communication Parameters > SNMP Configuration

SNMP Agent Community Name	Not Set.	Not Set.	Type a name.
SNMP Trap Destination Community Name (Parameters 1 through 4)	Not Set.	Not Set.	Type a community name for each of the parameters that you want to use.
SNMP Trap Destination IP Address (Parameters 1 through 4)	Not Set.	Not Set.	Type an IP address for each of the parameters that you want to use.
Internet Protocol	Not Set.	Not Set.	Type the Internet Protocol for the SNMP configuration.

TABLE B-7 Caching Parameters

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Write-Back Cache	Enabled.	Enabled.	Disabled. Enabled.
Optimization for Random/Sequential	Sequential (cannot be changed after the creation of a logical drive).	Sequential (cannot be changed after the creation of a logical drive).	Sequential or Random.

TABLE B-8 Peripheral Device Type Parameters

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Peripheral Device Type	Enclosure Services Device.	Enclosure Services Device.	No Device Present. Direct-Access Device. Sequential-Access Device. Processor Device. CD-ROM Device. Scanner Device. MO Device. Storage Array Controller Device. Enclosure Services Device. Unknown Device.
Peripheral Device Type Qualifier	Connected.	Connected.	Disconnected. Connected.
Device Supports Removable Media	Disabled.	Disabled.	Enabled. Disabled.
LUN Applicability	Undefined LUN-0s Only.		Disabled.

TABLE B-9 Host-Side and Drive-Side Parameters

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Host-side SCSI Parameters			
Maximum Queued I/O Count	1024 bytes	256 bytes	Auto, or 1-1024 bytes
LUNs per Host SCSI ID	32	32	1 to 32
Maximum Number of Concurrent Host-LUN Connections	1024	128	1 to 1024
Number of Tags Reserved for Each Host-LUN Connection	1024	32	1 to 1024 (FC) 1 to 256 (SCSI)

TABLE B-9 Host-Side and Drive-Side Parameters (*Continued*)

Host-side SCSI Parameters > Host Cylinder/ Head/ Sector Mapping Parameters

Sector Ranges	Variable.	Variable.	32, 64, 127, 255, Variable sectors.
Head Ranges	Variable.	Variable.	64, 127, 255, Variable heads.
Cylinder Ranges	Variable.	Variable.	1024, 32768, 65536, Variable cylinders.

Host-side SCSI Parameters > Fibre Connection Option

Loop only	Not applicable	Point to point only. Loop only. Loop preferred, otherwise point to point.
-----------	----------------	---

Drive-side SCSI Parameters

SCSI Motor Spin-Up	Disabled.	Disabled.	Do not change this parameter. Enabled. Disabled.
SCSI Reset at Power-Up	Enabled.	Enabled.	Do not change this parameter. Disabled. Enabled.
Disk Access Delay Time	15 seconds.	15 seconds.	Do not change this parameter. None to 75 seconds.
SCSI I/O Timeout	30 seconds.	7 seconds.	500 milliseconds to 30 seconds.
Maximum Tag Count	32	32	1-256 or Disabled.
Periodic Drive Check Time	Disabled.	Disabled.	1/2 to 30 seconds, Disabled.
Periodic SAF-TE/SES Check Time	30 seconds.	30 seconds.	Disabled to 60 seconds.

TABLE B-9 Host-Side and Drive-Side Parameters (*Continued*)

Periodic Auto-Detect Failure Drive Swap Check Time	Disabled.	Disabled.	5 to 60 seconds. Disabled.
Drive Predictable Failure Mode (SMART)	Disabled.	Disabled.	Do not change this parameter. Detect Only. Detect and Perpetual Clone. Detect and Clone + Replace. Disabled.
Auto-Assign Global Spare Drive	Disabled.	Disabled.	Enabled. Disabled.

TABLE B-10 Other Configuration Parameters

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Disk Array Parameters			
Rebuild Priority	Low.	Low.	Normal, Improved, High, Low.
Verification on Writes	Disabled.	Disabled.	On LD Initialization Writes Disabled. On LD Rebuild Writes Disabled. On Normal Drive Writes Disabled.
Redundant Controller Parameters			
Secondary Controller RS-232	Disabled.	Disabled.	Enabled. Disabled.
Remote Redundant Controller	Disabled.	Disabled.	Enabled. Disabled.
Controller Parameters			
Controller Name	Not Set.	Not Set.	Type a name.
LCD Title Display	Controller Logo.	Controller Logo.	Do Not Change this parameter.
Password Validation Timeout	Always Check.	Always Check.	Disable, 1, 2, or 5 minutes.
Controller Unique Identifier	Automatically set by the SAF-TE or SES device.	Automatically set by the SAF-TE or SES device.	Type a value.

TABLE B-10 Other Configuration Parameters *(Continued)*

SDRAM ECC	Enabled.	Enabled.	Do not change this parameter.
<hr/>			
DMEP Controller Parameters			
Total Usable Memory for DMEP	Disabled.	Disabled.	Type a percentage based on the RAM available from the controller.
<hr/>			

Default Peripheral Device Parameters

The following peripheral device parameters are available.

TABLE B-11 Peripheral Device Type Parameters (View and Edit Peripheral Devices)

User-Defined Parameter	FC Default Setting	SCSI Default Settings	Range of Values
Set Peripheral Device Entry			
Redundant Controller	Primary.	Primary.	Force primary controller failure. Force secondary controller failure.
UPS Status	Disabled.	Disabled.	Enabled. Disabled.
Set Peripheral Device Entry > Event Trigger Operations			
Temperature exceeds threshold	Not applicable in FC.	Disabled.	Disabled. Enabled.
Define Peripheral Device Active Signal	Low.	Low.	High. Low.
Controller Peripheral Device Configuration > Voltage and Temperature Parameters			
Upper Trigger Threshold for +3.3V Event	Default (3.6V).	Default (3.6V).	Disable, 3.4V - 3.9V.
Lower Trigger Threshold for +3.3V Event	Default (2.9V).	Default (2.9V).	Disable, 2.6V - 3.2V.
Upper Trigger Threshold for +5V Event	Default (5.5V).	Default (5.5V).	Disable, 5.2V - 6.0V.
Lower Trigger Threshold for +5V Event	Default (4.5V).	Default (4.5V).	Disable, 4.0V - 4.8V.
Upper Trigger Threshold for +12V Event	Default (13.2V).	Default (13.2V).	Disable, 12.5V - 14.4V.
Lower Trigger Threshold for +12V Event	Default (10.8V).	Default (10.8V).	Disable, 9.6V - 11.5V.
Upper Trigger Threshold for CPU Temperature Events	95 (C).	95 (C).	Disable, 50-100C.

TABLE B-11 Peripheral Device Type Parameters (View and Edit Peripheral Devices) (*Continued*)

Lower Trigger Threshold for CPU Temperature Events	Default 0 (C).	Default 0 (C).	Disable, 0-20C.
Upper Trigger Threshold for Board Temperature Events	85 (C).	85 (C).	Disable, 50-100C.
Lower Trigger Threshold for Board Temperature Events	Default 0 (C).	Default 0 (C).	Disable, 0-20C.

Default System Functions

The following system function parameters are available.

TABLE B-12 System Parameters (System Functions)

User-Defined Parameter	FC Default Setting	SCSI Default Setting	Range of Values
Mute Beeper	No.	No.	Yes. No.
Change Password	No.	No.	Type a password. No.
Reset Controller	No.	No.	Yes. No.
Shutdown controller (reserved)	No.	No.	Yes. No.
Controller Maintenance			
Restore NVRAM from disks	No.	No.	Yes. No.
Save NVRAM to disks	No.	No.	Yes. No.
Controller Maintenance > Advanced Maintenance Functions			
Download Boot Record with Firmware	No.	No.	Yes. No.

Keeping Specific Parameter Defaults

Default Parameter	Do Not Change these Default Parameters (unless directed by Technical Support)
Fault Management	
Drive Predictable Failure Mode (SMART)	Disabled on drives.
SDRAM ECC	Enabled.
SCSI Parameters	
Data Transfer Rate (sync transfer clock)	80 MHz.
Wide Transfer	Enabled.
Parity Check	Enabled.
Spin-Up Parameters	
SCSI Motor Spin-Up	Disabled.
SCSI Reset at Power-Up	Enabled.
Disk Access Delay Time	15 (none to 75 seconds).

Event Messages

This appendix lists the following event messages:

- “Controller Events” on page 242
- “SCSI Drive Events” on page 243
- “SCSI Channel Events” on page 244
- “Logical Drive Events” on page 245
- “General Target Alerts” on page 246
 - “SAF-TE Device” on page 246
 - “Controller On-Board” on page 247
 - “I²C Device” on page 247
 - “SES Devices” on page 248
 - “General Peripheral Device” on page 249

There are three categories of events.

TABLE C-1 Categories of Event Messages

Category	Description
Alert	Errors that need to be attended to immediately; might require reconnecting cables, replacing a component, or rebuilding a drive.
Warning	Errors that might indicate a temporary condition, a possible component problem, or the need to adjust controller parameters. Press Esc to clear message.
Notification	Informational message sent from controller firmware; press Esc to clear message.

Controller Events

The controller records all array events during power on; it records up to one thousand events.



Caution – Powering off or resetting the controller causes an automatic deletion of all recorded event logs.

Alerts

- [0104] Controller ALERT: SDRAM ECC Single-bit 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

SCSI Drive Events

SCSI drive event messages include:

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 (_).

Note – The three-digit code in parentheses for warning 1115 provides additional information about the drive error. The first of these three digits represents the SCSI Sense Key. The remaining two digits represent the Additional Sense Code (ASC). For more information about SCSI sense codes, refer to:

http://sunsolve.sun.com/handbook_pub/Systems/Sun4/TrDISK_SCSI_Sense_Codes.html

[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

SCSI Channel Events

SCSI channel event messages include:

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

Logical Drive Events

Logical drive event messages include:

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

General Target Alerts

General target alert messages include:

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)

[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 and logical drives.

- 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 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.
- ANSI** American National Standards Institute.
- ARP** Address Resolution Protocol.
- automatic rebuild** A process in which data is automatically reconstructed after a drive failure and written to a standby (spare) drive. An automatic rebuild also occurs when a new drive is installed manually in place of a failed drive. If the rebuild process is interrupted by a reset, use the Manual Rebuild command from the NAS appliance firmware to restart the rebuilding process.
- block striping** See striping.
- block striping with dedicated parity** (RAID 3) This technique breaks data into logical blocks, the size of a disk block, and then stripes these blocks across several drives. One drive is dedicated to parity. In the event that a disk fails, the original data can be reconstructed using the parity information and the information on the remaining drives.
- caching** Allows data to be stored in a predesignated 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.

capacity	The total number of physical drives available for data storage in a RAID array (logical drive). For example, if the capacity is N-1 and the total number of disk drives in a logical drives is six 36-Mbyte drives, the disk space available for storage is equal to five disk drives (5 x 36-Mbyte or 180 Mbytes).
CH	Channel.
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.
CIFS	Common Internet File System.
CISPR	International Special Committee on Radio Interference.
DHCP	Dynamic Host Configuration Protocol.
Directory	Storage in a NAS appliance file system is typically organized into directories for sharing.
disk mirroring	See mirroring.
EMC	Electromagnetic compatibility.
EMU	Event monitoring unit.
Fabric	Fibre Channel network built around one or more switches.
Fabric switch	A Fabric switch functions as a routing engine that 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 a data path of up to 100 mByte/sec 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 RAID 5 (striping with parity), or RAID 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, 1+0, 3, or 5.
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.

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 100 MByte/sec 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 rearranged in a star-like configuration because 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.
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.
global spare	A spare drive that is available to all logical drives in an array. Spare drives can be part of automatic logical drive rebuild.
Group (NAS)	A predefined grouping of users that require common directory access and privileges. After a group is created, it can be granted access to one or more share. Then any user that is added to the group has the share access that has been defined for that group.
Group (SSCS)	A data object that allows multiple servers to be contained under a single category. Groups are similar in concept to domains, and allow you to organize servers within Sun StorEdge Configuration Service.
HBA	Host bus adapter.
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 NAS appliance remains powered on and operational.
ID	Identifier number.
IEC	International Electrotechnical Commission.
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 are executed correctly.
JBOD	(Just a Bunch of Disks) A Sun StorEdge unit with drives and no controllers.

LAN	Local area network.
LD	Logical drive.
Load Balancing	A feature that assigns NAS clients to specific Ethernet data ports when both data ports are in the same subnet to distribute the communication load across both data ports.
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 might be located on one or more physical drives. Each NAS appliance controller can manage one to eight logical drives.
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.
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.
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.).
Management port	The NAS appliance 10/100 Ethernet port (LAN2). The management port is the only port that can be used to configure a NAS appliance.
MB	(megabyte) 1,000,000 bytes or characters of data.
mirroring (RAID 1)	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 percent 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. Disk mirroring provides 100 percent redundancy but is expensive because each drive in the array is duplicated.

**multiple-block striping
with distributed**

parity A RAID technique (RAID 5) that offers redundancy with the parity information distributed across all disks in the logical drive. Data and its parity are never stored on the same disk. In the event that a disk fails, the original data can be reconstructed using the parity information and the information on the remaining disks.

NAS Network attached storage.

NDMP Network Data Management Protocol.

NFS Network file system.

NVRAM (non-volatile random access memory) A memory unit equipped with a battery so that the data stays intact even after main power is switched off.

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 3 and 5) 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 is reported and the new correct parity is 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 A physical array is a group of physical drives in Sun StorEdge 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, Sun StorEdge Configuration Service requires that the same physical array of drives be used together in any remaining logical drives that are created with the remaining capacity.

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.

PID Primary controller identifier number.

RAID (Redundant Array of Independent Disks) An arrangement of two or more disk drives combined into a single virtual drive to provide more disk storage space, better performance and reliability, and redundant backup of data. Various combinations of these features are described by defined RAID levels. Sun StorEdge arrays can support RAID 0, 1, 1+0, 3, and 5.

RAID Level	Various techniques using combinations of mirroring, striping, duplexing, and parity to implement a RAID array are called RAID levels. Each technique uses a distinct algorithm to offer a mix of performance, reliability and cost.
RARP	Reverse Address Resolution Protocol.
read policy	A storage device parameter that determines whether the storage device holding data in cache before storing it to disk. The ability to hold data in cache while it is being written to disk can increase storage device speed during sequential reads.
rebuild	The process of reconstructing the data that was on a disk before it failed. Rebuilding can be done only in arrays with data redundancy, such as RAID levels 1, 1+0, 3, and 5.
rebuild priority	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 percent, the array administration activities have a higher priority than any other array activity. At 0 percent, the array administration activity is done only if there is no other activity on the array controller.
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	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).
Share	A share functions as a portal into a directory on a NAS appliance. Each share includes a set of permissions that apply to any group or user that accesses a directory through that share. Users must have access to a share to access files in a NAS directory.
SID	Secondary controller identifier number.
SMART	(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 email messages between servers and from mail clients to mail servers. The messages can then be retrieved with an email 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.
spanning	Making 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.
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, 1+0, 3, and 5.
state	The current operational status of a disk drive, a logical drive, or controller. The NAS appliance stores the states of drives, logical drives, and the controller in its nonvolatile memory. This information is retained across 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 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. This method of writing data increases the disk array throughput because multiple drives are working simultaneously, retrieving and storing. RAID 0, 1+0, 3, and 5 and 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.
UPS	Uninterruptible Power Supply.
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 policy	A cache-writing strategy used to control write operations. The write policy options are CIFS 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.
- WWN** (worldwide name) A globally unique, hard-coded and embedded number assigned by the manufacturer and registered under IEEE that is used to identify hardware.

Index

SYMBOLS

(, 87

NUMERICS

64 LUNs in redundant point-to-point configurations, 80

A

- active-active configuration, 18
- active-standby configuration, 18
- add drive Entry command, 110
- Add Global spare drive command, 103
- add Local spare drive command, 103
- add SCSI drives command, 59
- adding
 - drive entry, 110
 - global spare drive, 103
 - local spare drive, 103
 - SCSI drives, 59
- Adding a Host ID, 136
- Adding or Deleting Drive Entries, 110
- alarms silencing, 211
- Assigning a Local Spare Drive, 102
- automatic rebuild, 185

B

- BAD drive status, 99
- battery
 - period of usage, 20
 - shelf life, 20
- Battery Operation, 20
- battery support, 22
- beeper muting, 211
- Before You Begin, 2
- bus configurations, 100

C

- cabling
 - single- and dual-bus configurations, 100
- cache parameters, 158
- cache status, 28
- Changing a Logical Drive Controller
 - Assignment, 43
- channel IDs
 - permanent, 100
- channel status, 130
- channels
 - configuring as host or drive, 132
- check time
 - periodic auto-detection of failed drive swap, 176
- Clear drive status command, 110
- clone failing drive, 111
 - perpetual clone, 114
- Cloning a Failing Drive, 111

- COM port
 - connecting, 25
- concurrent rebuild, 189
- configuration
 - minimum requirements, 23
 - restoring from a file, 219
 - saving to a file, 216, 219
- Configuring SCSI Channels as Host or Drive, 132
- controller
 - muting beeper, 211
 - name, 181, 182
 - naming, 181, 182
 - optimization mode, 22
 - parameters, 230
 - password validation timeout, 182
 - SDRAM ECC, 184
 - unique identifier, 184
 - resetting, 214
 - restoring your configuration, 219
 - saving your configuration, 216, 219
 - shutting down, 215
- controller assignment, 43
- Controller Failover, 185
- controller failure, 185
- controller limitations, 18
- controller NAME, 55
- Controller Parameters, 181
- Controller Parameters command, 182
- Controller Voltage and Temperature Status, 206
- copy and replace
 - logical drive, 62
- copy and replace drive command, 63
- create
 - logical drive, 39
- creating
 - logical volume, 71
- Creating a Global Spare, 103
- Creating a Logical Volume, 70
- Creating Additional Host IDs, 135
- Creating Host Filter Entries, 89
- Creating Logical Drives, 37
- cylinder/head/sector mapping, 165

D

- Default Fibre Channel Host IDs, 135
- Default SCSI Host IDs, 136
- defaults
 - drive-swap check time, 176
 - enable UPS status, 204
 - host LUN connections, 163
 - logical drives, 34
 - number of LUNs per host SCSI ID, 162
 - password validation, 182
 - RAID levels, 34
 - SDRAM ECC, 184
 - UPS alert notification level, 205
- Define Peripheral Device Active Signal
 - command, 205
- Delete global/local spare drive command, 108
- deleting
 - drive entries, 110
 - host channel SCSI ID, 137
 - partition of logical drive
 - caution, 69
 - SCSI drive table
 - slot number, 109
 - SCSI ID, 137
 - spare drive, 107
- Deleting a Host LUN Mapping, 86
- deleting a logical drive, 51
- Deleting a Logical Drive Partition, 53
- Deleting a Spare Drive, 107
- Determining the Host Worldwide Name, 88
- DHCP, 155
- disconnecting support, 145
 - SCSI target/drive channel, 145
- disk
 - access delay time, 172
 - array parameters
 - advanced config, 178
- drive
 - cloning a failing, 111
 - entries
 - adding, 110
 - deleting, 110
 - failed, 15
 - checking, 176
 - partitioning a logical, 44
 - spare, 15
 - spares

- global, 16
- local, 16
- drive channel
 - viewing and editing targets, 141
- drive entries
 - adding or deleting, 110
- drive failure
 - recovering from fatal, 190
- drive information
 - viewing, 104
- drive status, 98
- drives
 - creating logical, 37
- drive-side parameters
 - advanced config, 170
- DRV FAILED status, 36
- dual-bus configurations, 100
- Dynamic Host Configuration Protocol, 155

E

- ECC SDRAM function, 184
- editing logical drives, 33
- Enabling and Disabling Write-Back and Write-Through Cache, 156
- error statistics, FC, 209
- Execute Drive Testing command, 127
- Expand logical drive command, 65
- Expand logical volume command, 72
- expanding
 - limitations, 58
 - logical drive, 64
 - copy and replace, 62
 - logical volume, 69, 72
- expansion
 - limitations, 58

F

- failed drive
 - identifying a, 121
- failing drive
 - cloning, 111
- failover

- controller, 185
- fatal drive failure
 - recovering from, 190
- fault management, controller parameters, 230
- fault prevention, 111
- FC arrays
 - planning for 1024 LUNs on, 79
- FC channel error statistics, 209
- fibre channel
 - controller parameter settings, 230
- Fibre Connection Option, 168
- filter
 - host, 87
- filter entries
 - creating host, 89
- filtering
 - LUN, 87
- firmware
 - automatic rebuild, 187
 - initial window, 28
 - logical drive status, 35, 116
 - main menu, 31
 - manual rebuild, 188
 - navigation keys, 30
 - quick installation, 31
 - SCSI drive status, 98
 - upgrading, 31
 - view event logs, 221
- firmware download
 - considerations, 2
- firmware window
 - initial screen, 28
 - viewing, 28
- Flash All Drives command, 122
- flash drive time command, 122
- flashing
 - all SCSI drives, 124
- format
 - low-level, 125

G

- gauge range, 28
- global spare drive, 15, 16
 - creating, 103

- deleting, 107
- explained, 16

H

- head
 - mapping, 165
- host application, 22
- host channel SCSI ID
 - adding, 135
 - deleting, 137
- Host Cylinder/Head/Sector Mapping
 - Configuration command, 166
- Host Filter Entries, 87
- host filter entries
 - creating, 89
- host ID
 - adding, 136
- host IDs
 - creating, 135
 - default Fibre Channel, 135
 - default SCSI, 136
- host LUN
 - connection
 - number of tags reserved, 164
 - example mappings, 85
- host LUN connections, 163
- host LUN mapping
 - deleting a, 86
- host LUNs
 - mapping logical drive partitions to, 75
- host worldwide name
 - determining, 88
- host-side parameters, 160

I

- I/O
 - maximum queued count, 161
 - SCSI timeout, 173
- ID
 - SCSI, deleting, 137
- Identifying a Failed Drive for Replacement, 121
- identifying drives, 15, 121
- Identifying SCSI drive command, 122

- INCOMPLETE status, 36
- initial firmware window, 28
- initial screen
 - main menu, 30
- INITING status, 36
- INVALID status, 36

K

- keys to navigate, 30

L

- LCD title display controller name
 - not applicable, 182
- limitations
 - controller, 18
 - expanding, 58
 - logical volume, 68
- local spare assignment, 41
- local spare drive, 15
 - assigning, 102
 - definition, 4
- local spare drives
 - deleting, 107
 - explained, 16
- logical drive
 - 128 LUN limit, 34
 - 128-LUN limit, 34
 - adding a SCSI drive, 58
 - assign local spare, 41
 - assigning a name, 54
 - change assignment, 43
 - commands, 33
 - view and edit logical drives, 35, 116
 - copying and replacing, 62
 - creating, 34, 37, 39
 - defaults, 34
 - deleting, 51
 - drive allocation, 5
 - expanding, 64
 - explained, 3
 - ID, 36
 - larger than 253 Gbyte, 166, 167
 - LG number, 36
 - maximum number of disks per logical drive, 159

- maximum physical capacity, 41
- maximum usable capacity per logical drive, 159
- minimum requirements, 23
- parity
 - checking, 56, 57
- partitioning, 44
- RAID level, 36
- RAID level selected, 40
- RAID levels, 34
- rebuilding, 55
- size, 36
- status, 35
- viewing and editing, 33
- logical drive name
 - changing, 54
 - creating, 54
- logical drive partition
 - deleting, 53
- logical drive partitions
 - mapping to host LUNs, 75
- logical drives, 178
 - creating, 37
 - rebuilding, 185
- logical volume, 68, 69, 70
 - avoiding failure, 69
 - creating, 70, 71
 - expanding, 72
 - using RAID expansion, 69
 - explained, 4, 68
 - limitations, 68
 - maximum partitions, 4
 - partitioning, 69
 - spare drives, 70
- logical volume status table, 73
- loop only, 168
- loop preferred, 168
- loopback error statistics, 209
- low-level format, 125
- LUN
 - described, 76
 - mapping a partition to, 81
 - per host SCSI ID, 162
 - changing, 162
- LUN filtering, 87
- LUN mapping
 - deleting a, 86
- LUNs

- mapping logical drive partitions to, 75
- maximum in point-to-point configurations, 80
- planning for 1024 on FC arrays, 79
- planning for 128 on SCSI arrays, 78

M

- Main Menu, 30
- manual rebuild, 188
- mapping
 - cylinder/head/sector, 165
- Mapping a Partition to a LUN, 81
- Mapping Logical Drive Partitions to Host LUNs, 75
- maximum
 - queued I/O count, 161
 - tag count (tag command queuing), 173
 - transfer width, 144
- maximum concurrent host LUN connections, 163
- maximum drive capacity, 41
- Maximum Queued I/O Count command, 161
- maximum sync. xfer Clock command, 143
- Maximum Tag Count (tag command queuing) command, 174
- maximum xfer Width command, 144
- MISSING drive status, 99
- motor spin-up, 171
 - SCSI, 171
- muting alarms, 211
- muting beeper, 211

N

- NAME (controller), 55
- naming
 - controller, 181, 182
- NEW DRV drive status, 99
- nexus (SCSI), 163
- non-volatile RAM
 - restoring your configuration from a file, 219
 - saving your configuration to a file, 216, 219
- Number of Tags Reserved for each Host-LUN Connection Command, 165
- NVRAM

restoring your configuration from a file, 219
saving your configuration to a file, 216, 219

O

optimization
 random I/O
 maximum size, 157
 sequential I/O
 maximum size, 157
optimization mode, 22
 limitations, 157
 random or sequential, 160

P

parameters
 cache, 158
 controller, 181
 drive-side, 170
 host-side
 advanced config, 160
 physical drives, 98
parity
 logical drive
 checking, 56, 57
parity check, 144, 145, 146
 SCSI target/drive channel, 144
partition
 deleting a logical drive, 53
 mapping to a LUN, 81
partitioning
 logical volume, 69
Partitioning a Logical Drive, 44
partitions
 logical volume
 maximum, 4
 mapping logical drive to host LUNs, 75
password
 disabling, 214
 setting a new, 213
 setting new, 212
 validation timeout, 182
Password Validation Timeout command, 183
PBCs, 168

periodic drive check time, 175
periodic drive swap auto check, 176
peripheral device
 setting, 203
 viewing SAF-TE status, 201
 viewing status, 193
Permanent SCSI Drive Channel IDs, 100
perpetual clone
 clone failing drive, 114
Physical Drive Status Table, 78
physical drives
 parameters, 98
physical drives listed, 39, 105
Planning for 1024 LUNs on FC Arrays, 79
Planning for 128 LUNs on SCSI Arrays, 78
Planning for 64 LUNs in Redundant FC Point-to-Point Configurations, 80
point to point only, 168
point to point preferred, 168
point-to-point configurations
 maximum LUNs, 80
port bypass circuit, 168
preparing for logical drives larger than 253 Gbyte, 166, 167
primary/secondary controller ID, 135

Q

quick installation, 31
 warning, 31

R

RAID
 advantages, 7
 controller, 178
 multilevel, 68
 planning considerations, 22
RAID (5+0), 70
RAID (5+1), 70
RAID (5+5), 70
RAID 1+0, 70
 concurrent rebuild in, 189
RAID Levels, 7

- RAID levels
 - controller optimization mode, 22
 - described, 9
 - planning, 22
 - RAID 0, 9
 - RAID 1, 10
 - RAID 1+0, 11
 - RAID 3, 12
 - RAID 3+0, 14
 - RAID 5, 13
 - RAID 5+0, 14
 - range of disks/logical drives supported, 8
 - selected, 40
- RAID Terminology Overview, 3
- random I/O optimization
 - maximum size, 157
- RARP, 155
- Read/Write Test, 126
- rebuild
 - automatic, 185
 - concurrent in RAID 1+0, 189
 - manual, 188
- Rebuild logical drive command, 55
- Rebuild Priority command, 178
- rebuilding, 178
 - logical drive, 55
- Rebuilding Logical Drives, 185
- Recovering From Fatal Drive Failure, 190
- redundant controller
 - explained, 185
- Regenerate Parity command, 58
- replace after clone
 - clone failing drive, 112
- Replace After Clone command, 112
- reset button
 - silencing failed component alarms, 212
- Restoring Your Configuration (NVRAM) From a File, 219
- RS-232
 - connecting to, 26

S

- SAF-TE status
 - viewing peripheral device, 201

- Saving Your Configuration (NVRAM) to a File, 216, 219
- SB-MISS drive status, 99
- scanning
 - new SCSI drive, 106
- scanning a new SCSI drive, 106
- Scanning Drives, 106
- SCSI
 - channel
 - status, 130
 - terminator, 139
 - I/O timeout, 173
 - ID
 - deleting, 137
 - motor spin-up, 171
 - reset at power-up, 172
 - transfer clock speed
 - setting, 140
- SCSI channel
 - setting termination, 139
- SCSI Channel Status Table, 130
- SCSI drive
 - adding to a logical drive, 58
 - disk capacity, 104
 - fibre port name, 104
 - global or local spare, 99
 - low-level format, 125
 - read/write test, 126
 - redundant loop ID, 104
 - revision number, 104
 - scanning new, 106
 - serial number, 104
 - setting slot numbers, 108
 - slot number
 - deleting, 109
- STANDBY mode, 99
- table
 - adding entries, 110
 - clear drive status, 110
 - deleting drive entries, 110
- USED DRV mode, 99
- utilities, 125, 128
- vendor ID, 99
- viewing information, 104

- scsi Drive Low-Level Format command, 126
- SCSI drive status, 98
- SCSI Drive Utilities, 125

- SCSI Motor Spin-Up command, 171
- SCSI target drive channel
 - viewing and editing, 141
- SCSI target/drive channel
 - parity check, 144
 - slot number, 142
- scsi Terminator command, 139
- SDRAM ECC, 184
 - default, 184
- sector
 - mapping, 165
- sequential I/O optimization
 - maximum size, 157
- serial port
 - connecting, 25
 - connection and setup, 25
- serial port parameters, 26
- SES status
 - viewing, 194
- set slot Number command, 111
- setting a new password, 213
- Setting a SCSI Channel Termination, 139
- Setting Peripheral Device Entry, 203
- Setting Slot Numbers, 108
- Setting the SCSI Transfer Width, 141
- Setting the UPS Power Fail Signal, 205
- Setting Transfer Clock Speed, 140
- Setting Up the Serial Port Connection, 25
- silencing alarms, 211
- single-bus configurations, 100
- slot number
 - assigning to empty slot, 109
 - deleting, 109
 - SCSI target/drive, 142
 - setting, 108
- slot numbers
 - setting, 108
- spare (local for logical drive), 41
- spare drive
 - assigning a local, 102
 - creating a global, 103
 - deleting, 107
- spare drives, 23, 189
 - deleting, 107
 - global, 15

- creating, 103
 - local, 15
 - explained, 102
 - logical volume, 70
- STAND-BY drive status, 99
- status
 - drive, 98
 - logical drive, 35
 - peripheral device, 193
 - UPS, 204
- status table
 - logical volume, 73
- sync transfer clock, 140
- system functions
 - controller
 - resetting, 214
 - shutting down, 215
 - password
 - disabling, 214
 - setting new password, 212

T

- tag count
 - maximum, 173
- temperature status
 - controller, 206
- terminator
 - SCSI channel, 139
- transfer clock
 - maximum synchronous, 143
- transfer clock speed
 - setting, 140
- transfer rate indicator, 28
- transfer width
 - maximum, 144
 - setting, 141

U

- uninterruptible power supply
 - enable status, 204
- Upgrading Firmware, 31
- UPS
 - enable status, 204

- fail signal, 205
 - status, 204
- UPS power fail signal, 205
 - setting, 205
- USED DRV drive status, 99
- utilities
 - SCSI drive, 125
- disabling, 156
- enabling, 156

V

- Verification on LD Initialization Writes Disabled
 - command, 180
- verification on writes, 179
- view
 - connected drives, 39, 105
- view and edit host LUNs command, 86
- view and edit logical drives command, 35, 116
- view and edit logical Volumes command, 71
- view and edit scsi drives command, 103
- view peripheral device status command, 194
- Viewing and Editing Drive Channel SCSI Targets, 141
- Viewing Drive Information, 104
- viewing logical drives, 33
- Viewing Peripheral Device Controller Status, 193
- Viewing Peripheral Device SAF-TE Status, 201
- Viewing SES Status, 194
- Viewing the Connected Drives, 105
- Viewing the Initial Firmware Windows, 28
- voltage status
 - controller, 206
- VT-100
 - connection and setup, 25

W

- warnings
 - quick installation, 31
- write error
 - avoiding, 179
- write-back cache
 - disabling, 156
 - enabling, 156
- write-through

