



Sun StorEdge™ 3000 Family RAID Firmware 4.2x User's Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 817-3711-18
July 2009, Revision A

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

Copyright © 2002–2009 Dot Hill Systems Corporation and others, 2200 Faraday Avenue, Suite 100, Carlsbad, California 92008, 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 © 2002–2009 Dot Hill Systems Corporation et d’autres, 2200 Faraday Avenue, Suite 100, Carlsbad, Californie 92008, 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 Etats-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 Etats-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 Etats-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.

Contents

Preface	xxix
1. Before You Begin	1
Supported Hardware Platforms	2
Key Concepts	4
RAID Planning Considerations	4
Local and Global Spare Drives	5
Local Spare Drives	6
Global Spare Drives	6
Using Both Local and Global Spare Drives	7
Accessing the Firmware Application	8
Setting an IP Address	8
▼ To Set an Array's IP Address	9
2. Basic Firmware Components	11
Viewing the Initial Firmware Screen	11
Navigating Firmware Menus	15
Navigation Terminology and Conventions	17
Menu Options That Toggle Between Conditions	18
Progress Indicators	18

Device Capacities	19
3. Configuration Defaults and Guidelines	21
Default Configurations	21
Default Logical Drive Configuration	22
Default Channel Configurations	22
Maximum Drive Configurations per Array	24
Maximum Number of Disks and Maximum Usable Capacity per Logical Drive	25
Controller Operation Guidelines	28
Dual-Controller Guidelines	28
Single-Controller Guidelines	29
Cache Optimization Mode and Stripe Size Guidelines	30
Cache Write Policy Guidelines	32
Fibre Connection Protocol Guidelines	33
A Sample SAN Point-to-Point Configuration	35
▼ To Set Up a Typical Point-to-Point SAN Configuration	39
A Sample DAS Loop Configuration	41
▼ To Set Up a Typical DAS Loop Configuration	44
Array Configuration Summary	45
4. First-Time Configuration for SCSI Arrays	49
Existing Logical Drive Configuration	50
▼ To View the Logical Drive Configuration	50
▼ To View the Channel Configuration	51
Deleting Logical Drives	51
▼ To Unmap and Delete a Logical Drive	51
Cache Optimization Mode (SCSI)	52
▼ To Verify the Optimization Mode	53
▼ To Change the Optimization Mode	53

Physical Drive Status	54
▼ To Check Physical Drive Availability	54
Channel Settings	56
Configuring Channel Mode	56
▼ To Configure the Channel Mode	57
Host Channel IDs	58
▼ To Add or Delete a Unique Host ID	58
Creating Logical Drives	59
Preparing for Logical Drives Larger Than 253 Gbyte	60
▼ To Change Cylinder and Head Settings	61
▼ To Create a Logical Drive	61
Controller Assignment	70
▼ To Change a Controller Assignment (Optional)	70
Logical Drive Name	71
▼ To Assign a Logical Drive Name (Optional)	71
Partitions	72
▼ To Partition a Logical Drive (Optional)	72
Mapping a Partition to a Host LUN	74
▼ To Map a Logical Drive Partition	75
Labeling a LUN (Solaris Operating System Only)	78
▼ To Label a LUN	78
Solaris Operating System Device Files	79
▼ To Create Device Files for Newly Mapped LUNs	79
Saving Configuration (NVRAM) to a Disk	80
▼ To Save a Configuration to NVRAM	80
5. First-Time Configuration for FC or SATA Arrays	81
Existing Logical Drive Configuration	82
▼ To View the Logical Drive Configuration	82

▼ To View the Channel Configuration	83
Deleting Logical Drives	83
▼ To Unmap and Delete a Logical Drive	83
Cache Optimization Mode (FC and SATA)	84
▼ To Verify the Optimization Mode	85
▼ To Change the Optimization Mode	85
Physical Drive Status	86
▼ To Check Physical Drive Availability	86
Enabling Support for SATA Expansion Units Attached to FC Arrays	87
▼ To Enable or Disable Mixed Drive Support	88
Channel Settings	88
Configuring Channel Mode	88
▼ To Modify a Channel Mode	89
Redundant Communication Channels (RCCOM)	90
Using Four DRV + RCCOM Channels	90
▼ To Configure Channels 4 and 5 as Additional DRV + RCCOM Channels	90
Using Channels 4 and 5 as RCCOM-Only Channels	91
▼ To Configure Channels 4 and 5 as RCCOM-Only Channels	91
Host Channel IDs	92
▼ To Add or Delete a Unique Host ID	93
Channel ID Ranges	94
▼ To Assign an ID From a Different Range	95
Fibre Connection Protocol	97
▼ To Change the Fibre Connection Protocol	97
Creating Logical Drives	97
Preparing for Logical Drives Larger Than 253 Gbyte (Solaris Operating System Only)	98
▼ To Change Cylinder and Head Settings	98

▼ To Create a Logical Drive	99
Controller Assignment	108
▼ To Change a Controller Assignment (Optional)	108
Logical Drive Name	109
▼ To Assign a Logical Drive Name (Optional)	109
Partitions	109
▼ To Partition a Logical Drive (Optional)	110
Mapping a Partition to a Host LUN	112
LUN Mapping	112
LUN Filtering (FC and SATA Only)	114
▼ To Map a Logical Drive Partition	116
▼ To Create Host Filters (FC and SATA Arrays Only)	118
Labeling a LUN (Solaris Operating System Only)	123
▼ To Label a LUN	123
Creating Solaris Operating System Device Files for Newly Mapped LUNs	124
▼ To Create Device Files for Newly Mapped LUNs	124
Saving Configuration (NVRAM) to a Disk	125
▼ To Save a Configuration to NVRAM	125
6. Logical Drives	127
Viewing the Logical Drive Status Table	128
Viewing Physical Drives	128
▼ To View Physical Drives	128
Creating Logical Drives	129
Deleting a Logical Drive	129
Partitioning a Logical Drive	130
Deleting a Logical Drive Partition	130
▼ To Delete a Logical Drive Partition	130
Changing a Logical Drive Name	132

Rebuilding a Logical Drive	132
▼ To Rebuild a Logical Drive	132
Changing a Logical Drive Controller Assignment	133
Expanding the Capacity of a Logical Drive	133
▼ To Expand a Logical Drive	133
Adding Physical Drives	137
▼ To Add a Physical Drive to a Logical Drive	137
Performing a Parity Check	139
Manually Regenerating the Parity of a Logical Drive	139
▼ To Regenerate the Parity of a Logical Drive	140
Overwriting Inconsistent Parity	140
▼ To Enable or Disable Overwrite Inconsistent Parity	141
Generating a Check Parity Error Event	141
▼ To Enable or Disable Reporting Inconsistent Parity Errors as System Events	141
Clearing a Fatal Fail Logical Drive Condition	142
Copying and Replacing Drives With Drives of Larger Capacity	143
▼ To Copy and Replace a Drive	144
Scanning Drives for Bad Blocks	145
▼ To Terminate a Media Scan	146
▼ To Perform a Media Scan	146
Shutting Down a Logical Drive	147
▼ To Shut Down a Logical Drive	147
Changing Write Policy for a Logical Drive	148
▼ To Configure a Logical Drive's Write Policy	148
7. Logical Volumes	151
Understanding Logical Volumes (Multilevel RAID)	152
Logical Volume Limitations	152

Logical Drive and Logical Volume Partitions	153
RAID Expansion	153
Multilevel RAID Arrays	154
Spare Drives	154
Viewing the Logical Volume Status Table	155
Creating a Logical Volume	155
▼ To Create a Logical Volume	155
Deleting a Logical Volume	158
▼ To Delete a Logical Volume	158
Expanding a Logical Volume	158
▼ To Expand a Logical Volume	158
8. Host LUNs	159
Mapping Logical Drive Partitions to Host LUNs	160
Planning for 128 LUNs on a SCSI Array (Optional)	162
▼ To Create 128 LUNs	162
Planning for 1024 LUNs on an FC or SATA Array (Optional, Loop Mode Only)	163
▼ To Create 1024 LUNs	163
Planning for 64 LUNs in Redundant FC or SATA Point-to-Point Configurations	164
Mapping a Partition to a LUN	165
Deleting a Host LUN Mapping	165
▼ To Delete a Host LUN Mapping	165
Creating Host Filter Entries (FC and SATA Only)	166
Determining a Host World Wide Name	166
▼ To Determine a WWN for the Solaris Operating System	166
▼ To Determine a WWN for the Linux, Windows 2000, or Windows 2003 Operating System	167
Manually Adding WWN Entries Using the Host-ID/WWN Name List	168

Viewing and Modifying Host Filter Information	168
▼ To View or Modify Host Filter Information	169
9. Physical Drives	171
Viewing the Status of a Physical Drive	172
▼ To View the Physical Drive Status Table	172
SCSI Drive IDs (SCSI Only)	173
FC Drive IDs (FC and SATA Only)	175
Viewing Physical Drive Information	176
▼ To View Physical Drive Information	176
Assigning a Local Spare Drive	177
▼ To Assign a Local Spare Drive	177
Assigning a Global Spare	178
▼ To Assign a Global Spare	178
Deleting a Spare Drive	178
▼ To Delete a Spare Drive	178
Scanning Drives (SCSI Only)	179
▼ To Scan a New SCSI Drive	179
Adding or Deleting Drive Entries (SCSI Only)	179
▼ To Add a Drive Entry	180
▼ To Remove an Empty Drive Entry	180
Identifying a Failed Drive for Replacement	180
▼ To Identify a Drive	181
Flashing a Selected Physical Drive	181
Flashing All SCSI Drives	182
Flashing All But Selected Drive	183
Fault Protection Measures	183
Cloning a Failing Drive	184
Replacing After Clone	184

▼ To Replace After Clone	184
Perpetual Clone	186
▼ To Enable Perpetual Clone	186
Terminating a Perpetual Clone	188
▼ To Terminate a Perpetual Clone	188
Viewing the Status of a Cloning Operation	189
▼ To View the Status of a Cloning Operation	189
Using SMART Functionality	190
▼ To Enable and Use SMART Detection	191
▼ To Test a Drive for SMART Functionality	192
▼ To Disable SMART Detection	193
Using Media Scan on Individual Drives	193
▼ To Perform a Media Scan	193
▼ To Terminate a Media Scan	194
SCSI Drive Utilities (Reserved)	195
SCSI Drive Low-Level Format Utilities	195
▼ To Low-Level Format a Physical Drive	195
Read/Write Test	196
▼ To Perform a Read/Write Test	196
Changing Disk Reserved Space	197
▼ To Remove Reserved Space from a Drive	198
▼ To Specify Disk Reserved Space	198
10. Host and Drive Channels	199
Host and Drive Channel Status Table	200
▼ To Check and Configure Host and Drive Channels	200
Configuring Channels as Host or Drive	201
Creating Additional Host IDs	201
Deleting a Host Channel SCSI ID	201

- ▼ To Delete a Host Channel SCSI ID 201
- Drive Channel SCSI IDs 202
- Setting SCSI Channel Termination (SCSI Only) (Reserved) 202
 - ▼ To Enable or Disable SCSI Channel Termination (SCSI Only) 203
- Setting Transfer Clock Speed (SCSI Only) 203
 - ▼ To Change the Sync Transfer Clock Speed (SCSI Only) 203
- Setting the SCSI Transfer Width (SCSI Only) 204
 - ▼ To Change the Transfer Width Options (SCSI Only) 204
- Enabling SCSI Bus Byte Parity Checking (SCSI Only) 204
 - ▼ To Enable or Disable Parity Checking 205
- Viewing Chip Information 205
 - ▼ To View Chip Information 205
- Viewing Channel Host-ID WWN Information (FC and SATA Only) 206
 - ▼ To View a Channel's Host-ID/WWN (FC and SATA Only) 206
- Viewing Device Port Names (WWPN) (FC and SATA Only) 206
 - ▼ To View a Channel's Device Port Name List 207
- Setting a Channel's Data Rate (FC and SATA Only) 207
 - ▼ To Set a Channel's Data Rate 208
- Issuing a Loop Initialization Primitive (FC and SATA Only) 209
 - ▼ To Issue a LIP 209

11. Configuration Parameters 211

- Communication Parameters 212
 - RS-232 Port Configuration (Reserved) 212
 - ▼ To Configure the COM Port Baud Rate 212
 - ▼ To Enable or Disable Terminal Emulation Through a Serial Port 213
 - Setting an IP Address 213
 - ▼ To Set an Array's IP Address 214
 - Network Protocol Support 215

- ▼ To Enable and Disable Network Protocols 215
- Setting Telnet Inactivity Timeout Time 216
- ▼ To Set the Telnet Inactivity Timeout 216
- SNMP 217
 - SNMP Traps 217
 - SNMP Trap Objects 219
 - System Group Objects 220
 - Enterprise Objects 224
 - Using Multiple Product MIBs 225
- Configuring the Notification Processing Center (NPC) 226
 - ▼ To Create and Save the NPC Configuration File (`agent.ini`) 227
 - A Simple Sample `agent.ini` File 228
 - A Complete Sample `agent.ini` File 228
 - `agent.ini` File Parameters 229
 - The `SNMP_TRAP` Section 229
 - The `EMAIL` Section 230
 - The `BROADCAST` Section 230
- Caching Parameters 231
 - Enabling and Disabling Write-Back Cache 231
 - ▼ To Change the Write-Back Cache Option 232
 - Setting Optimization 232
 - Setting Periodic Cache Flush Time 232
 - ▼ To Set Periodic Cache Flush Time 232
- Host-Side Parameters Menu Options 233
 - Maximum Queued I/O Count 233
 - ▼ To Set the Maximum Queued I/O Count 233
 - LUNs Per Host SCSI ID 234
 - ▼ To Change the Number of LUNs Per Host SCSI ID 234

Maximum Number of Concurrent Host-LUN Connections	234
▼ To Change the Maximum Number of Concurrent Host-LUN Connections	235
Number of Tags Reserved for Each Host LUN Connection	235
▼ To Modify the Tag Command Queuing on the Host-LUN Connection	236
Peripheral Device Type Parameters (Reserved)	236
Host Cylinder/Head/Sector Mapping Configuration	236
▼ To Configure Sector Ranges, Head Ranges, and Cylinder Ranges	237
Preparing for Logical Drives Larger Than 253 Gbyte on Solaris Systems	237
Configuring In-Band EI Management	238
▼ To Configure In-Band EI Management	238
Fibre Connection Options (FC and SATA Only)	238
▼ To Confirm or Change the Fibre Connection for the Array	239
Drive-Side Parameters Menu	240
Configuring Drive Motor Spin-Up (Reserved)	240
▼ To Spin Up SCSI Hard Drives (Reserved)	240
Configuring Disk Access Delay Time	241
▼ To Set Disk Access Delay Time	241
Configuring Drive I/O Timeout	241
▼ To Choose Drive I/O Timeout	242
Configuring the Maximum Tag Count (Tag Command Queuing)	242
▼ To Change the Maximum Tag Count Setting	243
Configuring Periodic Drive Check Time	243
▼ To Set the Periodic Drive Check Time	243
Configuring Periodic SAF-TE and SES Device Check Time	243
▼ To Set the Periodic SAF-TE and SES Device Check Time	244
Configuring Periodic Auto-Detect Failure Drive Swap Check Time	244
▼ To Set the Auto-Detect Failure Drive Swap Check Time	245
Drive Predictable Failure Mode (SMART)	245

Auto-Assign Global Spare Drive (FC and SATA Only)	245
▼ To Automatically Assign Replacements to Faulty Drives	245
Media Scan at Power-Up	246
Disk Array Parameters Menu	246
Setting Rebuild Priority	246
▼ To Set the Rebuild Priority	246
Verification on Writes	247
▼ To Enable and Disable Verification Methods	247
Redundant Controller Parameters	247
Enabling and Disabling Cache Synchronization	248
▼ To Enable or Disable Cache Synchronization	248
Controller Parameters	248
Configuring a Controller Name	248
▼ To View and Display the Controller Name	249
Password Validation Timeout	249
▼ To Set a Password Validation Timeout	250
Controller Unique Identifier (Reserved)	250
▼ To Specify the Controller Unique Identifier	250
Setting the Controller's Date and Time	251
▼ To Set the Controller's Time Zone	251
▼ To Set the Controller Date and Time	252
12. Peripheral Devices	253
Viewing Peripheral Device Controller Status	253
Viewing SES Status (FC and SATA Only)	254
▼ To Check the Status of SES Components (FC and SATA Only)	255
Identifying Fans (FC and SATA Only)	256
▼ To View the Status of Each Fan	257
SES Temperature Sensor Locations (FC and SATA Only)	259

SES Voltage Sensors (FC and SATA Only)	259
SES Power Supply Sensors (FC and SATA Only)	261
Viewing Peripheral Device SAF-TE Status (SCSI Only)	261
▼ To Check the Status of SAF-TE Components (SCSI Only)	262
Identifying Fans (SCSI Only)	263
SAF-TE Temperature Sensor Locations (SCSI Only)	264
SAF-TE Power Supply Sensors (SCSI Only)	265
Setting Peripheral Device Entry	265
Redundant Controller - Primary	265
Forcing Primary Controller Failure (Reserved)	266
▼ To Force a Primary Controller Failure (Reserved)	266
Forcing Secondary Controller Failure (Reserved)	266
▼ To Force a Secondary Controller Failure (Reserved)	266
▼ To Restore a Force-Failed Primary or Secondary Controller	266
Event Trigger Operations	267
Configuring the Controller Failure Event Trigger	267
▼ To Enable or Disable the Controller Failure Event Trigger	267
Configuring the Battery Backup (BBU) Low Event or BBU Failed Event Trigger	267
▼ To Enable or Disable the BBU Low Event or BBU Failed Event Trigger	268
Configuring the Power Supply Failed Event Trigger	268
▼ To Enable or Disable the Power Supply Failed Event Trigger	268
Configuring the Fan Failure Event Trigger	268
▼ To Enable or Disable the Fan Failure Event Trigger	268
Configuring the Temperature Exceeds Threshold Event Trigger	268
▼ To Configure Over-Temperature Controller Shutdown	269
Operating in a NEBS Environment	269
Viewing Controller Voltage and Temperature Status	270

	▼ To Display Controller Voltage and Temperature Status	271
	▼ To View or Configure Thresholds	271
	Fibre Channel Error Statistics (FC and SATA Only)	273
13.	System Functions and Event Logs	275
	Muting the Beeper	275
	▼ To Change the Beeper Setting	276
	Setting and Changing the Controller Password	276
	▼ To Create a New Password	277
	▼ To Change an Existing Password	277
	▼ To Disable an Existing Password	278
	Resetting the Controller	278
	▼ To Reset the Controller Without Saving Cache Contents	279
	Shutting Down the Controller	279
	▼ To Shut Down a Controller	280
	The Download Firmware Option (Reserved)	280
	The Advanced Maintenance Functions Option (Reserved)	280
	Saving Your Configuration (NVRAM) to Disk	281
	▼ To Save Your Configuration (NVRAM)	283
	Restoring Your Configuration (NVRAM) From Disk	284
	▼ To Restore Saved Configuration Settings	285
	Clearing a Core Dump	286
	Viewing System Information	287
	Viewing Event Logs on the Screen	288
	▼ To View The Array's Event Log	288
14.	Maintaining Your Array	291
	Battery Operation	291
	Battery Status	292

Battery Support for Cache Operations	293
Checking Status Windows	293
Logical Drive Status Table	293
Physical Drive Status Table	296
Channel Status Table	298
Upgrading Firmware	301
Patch Downloads	302
Installing Firmware Upgrades	302
Controller Firmware Upgrade Features	303
Upgrading SES and PLD Firmware	304
Troubleshooting Your Array	305
Controller Failover	305
RAID LUNs Not Visible to Host	306
Rebuilding Logical Drives	306
Automatic Logical Drive Rebuild	306
Manual Rebuild	308
Concurrent Rebuild in RAID 1+0	310
Modifying Drive-Side Parameters	310
Additional Troubleshooting Information	310
A. Basic RAID Concepts	311
RAID Terminology Overview	312
Logical Drives	312
Logical Volumes	313
Channels, Partitions, and LUN Mapping	313
RAID Levels	316
RAID 0	319
RAID 1	319
RAID 1+0	320

RAID 3	321
RAID 5	322
Advanced RAID Levels	323
B. Firmware Specifications	325
C. Record of Settings	331
View and Edit Logical Drives	332
Logical Drive Information	332
Logical Drive Partition Information	333
LUN Mappings	334
View and Edit Drives	335
View and Edit Channels	336
View and Edit Peripheral Devices	337
View System Information	337
Save NVRAM to Disk and Restore From Disk	338
D. Parameter Summary Tables	339
Introducing Default Parameters	339
Basic Default Parameters	340
Default Configuration Parameters	341
Default Peripheral Device Parameters	348
Default System Functions	350
Parameter Defaults That Must Not Change	351
E. Event Messages	353
Controller Events	354
Controller Alerts	354
Controller Warnings	357
Controller Notifications	358

Drive Events	359
Drive Alerts	359
Drive Warnings	362
Drive Notifications	363
Channel Events	363
Channel Alerts	363
Channel Notifications	367
Logical Drive Events	367
Logical Drive Alerts	368
Logical Drive Notifications	371
General Target Events	375
SAF-TE Device Events	375
SAF-TE Device Alerts	375
SAF-TE Device Notifications	376
Controller Self-Diagnostic Events	377
Controller Self-Diagnostic Alerts	377
Controller Self-Diagnostic Notifications	377
I ² C Device Events	378
I ² C Device Notifications	378
SES Device Events	378
SES Device Alerts	378
SES Device Notifications	379
General Peripheral Device Events	380
General Peripheral Device Alerts	380
General Peripheral Device Notifications	382

F. SNMP Management Using HP OpenView 385

Loading MIBs	386
Configuring Events	389

Viewing and Setting System Group Objects 391

Viewing Enterprise Objects 394

Glossary 397

Index 405

Figures

FIGURE 1-1	Local (Dedicated) Spare	6
FIGURE 1-2	Global Spare	7
FIGURE 1-3	Mixing Local and Global Spares	7
FIGURE 2-1	Initial Firmware Screen	12
FIGURE 2-2	Write-through (WT) Cache Status forced by an event trigger	14
FIGURE 2-3	Firmware Main Menu	15
FIGURE 2-4	Progress Indicator and Descriptive Message	18
FIGURE 3-1	A Point-to-Point Configuration with a Dual-Controller Sun StorEdge 3510 FC Array and Two Switches	37
FIGURE 3-2	A Point-to-Point Configuration With a Dual-Controller Sun StorEdge 3511 SATA Array and Two Switches	38
FIGURE 3-3	A DAS Configuration With Four Servers, a Dual-Controller Sun StorEdge 3510 FC Array, and Two Expansion Units	42
FIGURE 3-4	A DAS Configuration With Four Servers, a Dual-Controller Sun StorEdge 3511 SATA Array, and Two Expansion Units	43
FIGURE 4-1	Partitions in Logical Drives	72
FIGURE 4-2	LUNs Resemble Drawers in a File Cabinet Identified by an ID	74
FIGURE 4-3	Mapping Partitions to Host ID/LUNs	75
FIGURE 5-1	Partitions in Logical Drives	110
FIGURE 5-2	LUNs Resemble Drawers in a File Cabinet	113
FIGURE 5-3	Mapping Partitions to Host ID/LUNs	113
FIGURE 5-4	Example of LUN Filtering	115

FIGURE 6-1	Example of Deleted Partitions	131
FIGURE 6-2	Expansion by Copy and Replace	143
FIGURE 7-1	Logical Volume Composed of Multiple Drives	152
FIGURE 8-1	Filing Cabinet Represents the SCSI or FC ID	161
FIGURE 8-2	Mapping Partitions to Host ID/LUNs	161
FIGURE 9-1	Flashing the Drive LED of a Selected Drive	182
FIGURE 9-2	Flashing All Drive LEDs to Detect a Defective Non-Flashing Drive	182
FIGURE 9-3	Flashing All Drive LEDs Except a Selected Drive LED	183
FIGURE 12-1	View Peripheral Device Status	254
FIGURE 12-2	FC and SATA Cooling Fan Locations	258
FIGURE 12-3	Example of the SAF-TE Device Status Window for a Single-Bus Configuration	262
FIGURE 12-4	Example of SAF-TE Device Status Window for a Split-Bus Configuration	263
FIGURE 12-5	Cooling Fan Locations	264
FIGURE 14-1	Automatic Rebuild	308
FIGURE 14-2	Manual Rebuild	309
FIGURE A-1	Logical Drive Including Multiple Physical Drives	312
FIGURE A-2	Allocation of Drives in Logical Drive Configurations	314
FIGURE A-3	Partitions in Logical Drive Configurations	315
FIGURE A-4	Mapping Partitions to Host ID/LUNs	315
FIGURE A-5	Mapping Partitions to LUNs Under an ID	316
FIGURE A-6	RAID 0 Configuration	319
FIGURE A-7	RAID 1 Configuration	320
FIGURE A-8	RAID 1+0 Configuration	321
FIGURE A-9	RAID 3 Configuration	322
FIGURE A-10	RAID 5 Configuration	323

Tables

TABLE 2-1	Firmware Screen Components	12
TABLE 2-2	Navigation Keys	15
TABLE 2-3	Navigation Terminology and Character Conventions	17
TABLE 2-4	Progress Indicator Prefix Meanings	19
TABLE 3-1	Sun StorEdge 3310 SCSI Array and Sun StorEdge 3320 SCSI Array Default Channel Settings	22
TABLE 3-2	Sun StorEdge 3510 FC Array Default Channel Settings	22
TABLE 3-3	Sun StorEdge 3511 SATA Array Default Channel Settings	23
TABLE 3-4	Maximum Number of Supported Physical and Logical Drives, Partitions, and LUN Assignments	24
TABLE 3-5	Actual Capacities per Drive	25
TABLE 3-6	Maximum Usable Storage Capacity Determined by RAID Level	26
TABLE 3-7	Maximum Number of Disks per Logical Drive	26
TABLE 3-8	Maximum Usable Capacity (in Gbyte) per Sun StorEdge 3510 FC Logical Drive	27
TABLE 3-9	Maximum Usable Capacity (in Gbyte) per Sun StorEdge 3310 SCSI and Sun StorEdge 3320 SCSI Logical Drive	27
TABLE 3-10	Maximum Usable Capacity (in Gbyte) per Sun StorEdge 3511 SATA Logical Drive	27
TABLE 3-11	Example Point-to-Point Configuration With Two Logical Drives in a Dual-Controller Array	39
TABLE 3-12	Connection for Four Servers in a DAS Configuration	41
TABLE 3-13	Example Primary and Secondary ID Numbers in a Loop Configuration With Two IDs per Channel	44
TABLE 4-1	Cylinder and Head Mapping for the Solaris Operating System	60

TABLE 5-1	ID Values Assigned to Each ID Range	94
TABLE 5-2	Cylinder and Head Mapping for the Solaris Operating System	98
TABLE 7-1	Parameters Displayed in the Logical Volume Status Window	155
TABLE 8-1	Configuration for 1024 LUNs	164
TABLE 8-2	Example of ID Assignments for 64 LUNs With Multipathing Enabled	165
TABLE 9-1	ID Switch Settings for FC Expansion Units	175
TABLE 11-1	System Group Objects	221
TABLE 11-2	Cylinder and Head Mapping for the Solaris Operating System	238
TABLE 12-1	Fan Status and Fan Speeds	257
TABLE 12-2	Relationship Between Cooling Elements, Fans, and Power Supply Modules	258
TABLE 12-3	Temperature Sensor Locations (FC and SATA)	259
TABLE 12-4	Voltage Sensors for Sun StorEdge 3510 FC Arrays	260
TABLE 12-5	Voltage Sensors for Sun StorEdge 3511 SATA Arrays	260
TABLE 12-6	Power Supply Sensors (FC and SATA)	261
TABLE 12-7	Location of Cooling Fans	263
TABLE 12-8	Temperature Sensor Locations (SCSI)	264
TABLE 12-9	Power Supply Sensors (SCSI)	265
TABLE 14-1	Battery Status Indicators	292
TABLE 14-2	Parameters Displayed in the Logical Drive Status Window	294
TABLE 14-3	Parameters Displayed in the Physical Drive Status Window	297
TABLE 14-4	Parameters Displayed in the Channel Status Table	299
TABLE A-1	RAID Level Overview	317
TABLE A-2	RAID Level Characteristics	318
TABLE A-3	Advanced RAID Levels	323
TABLE B-1	Basic RAID Management	325
TABLE B-2	Advanced Features	326
TABLE B-3	Caching Operation	327
TABLE B-4	RAID Expansion	327
TABLE B-5	Redundant Controller	328
TABLE B-6	Data Safety	328

TABLE B-7	Security	329
TABLE B-8	Environment Management	329
TABLE B-9	User Interface	330
TABLE D-1	Logical Drive Parameters (View and Edit Logical Drives)	340
TABLE D-2	Logical Volume Parameters (View and Edit Logical Volumes)	340
TABLE D-3	Host LUN Parameters (View and Edit Host LUNs)	340
TABLE D-4	Drive Parameters (View and Edit Drives)	340
TABLE D-5	Channel Parameters (View and Edit Channels)	341
TABLE D-6	Communication Parameters > RS-232 Port Configuration	342
TABLE D-7	Communication Parameter > Internet Protocol (TCP/IP)	342
TABLE D-8	Communication Parameters > Network Protocol Support	342
TABLE D-9	Communication Parameter > Telnet Inactivity Timeout Time	343
TABLE D-10	Caching Parameters	343
TABLE D-11	Peripheral Device Type Parameters	344
TABLE D-12	Host-Side and Drive-Side Parameters	345
TABLE D-13	Other Configuration Parameters	347
TABLE D-14	Peripheral Device Type Parameters (View and Edit Peripheral Devices)	348
TABLE D-15	System Function Parameters	350
TABLE D-16	Default Parameters That Must Not Change (unless directed to do so by Technical Support)	351
TABLE E-1	Categories of Event Messages	354

Preface

This guide provides an overview of redundant array of independent disks (RAID) functionality and explains the use of controller firmware commands to configure and monitor Sun StorEdge™ 3000 family arrays.

Note – Different versions of controller firmware apply to different Sun StorEdge 3000 family arrays. Even if the same version number is used, the SunSolve patch number differs for each hardware platform, as does the binary filename. 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 explains what you need to know before you begin accessing and using the firmware application.
- Chapter 2 introduces the initial firmware screen, menu structure, and navigation and screen conventions.
- Chapter 3 describes first-time array configuration.

- Chapter 4 summarizes common procedures for the first-time configuration of Sun StorEdge 3310 SCSI arrays and Sun StorEdge 3320 SCSI arrays.
- Chapter 5 summarizes common procedures for the first-time configuration of Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA arrays.
- Chapter 6 describes the “**view and edit Logical drives**” menu option and related procedures.
- Chapter 7 describes the “**view and edit logical Volumes**” menu option and related procedures.
- Chapter 8 describes the “**view and edit Host luns**” menu option and related procedures.
- Chapter 9 describes the “**view and edit scsi Drives**” menu option and related procedures.
- Chapter 10 describes the “**view and edit channels**” menu option and related procedures.
- Chapter 11 describes the “**view and edit Configuration parameters**” menu option and related procedures.
- Chapter 12 describes the “**view and edit Peripheral devices**” menu option and related procedures.
- Chapter 13 describes the “**system Functions**” menu option, array information, and event logs.
- Chapter 14 describes array maintenance procedures.
- Appendix A introduces RAID terminology and concepts.
- Appendix B provides the array firmware specifications.
- Appendix C summarizes controller parameters for optimization and parameter defaults that should not be changed.
- Appendix D lists the firmware parameter settings for Sun StorEdge 3510 FC arrays, Sun StorEdge 3511 SATA arrays, Sun StorEdge 3310 SCSI arrays, and Sun StorEdge 3320 SCSI arrays.
- Appendix E lists and describes firmware event messages.
- Appendix F describes SNMP management using applications such as HP OpenView.
- Glossary provides RAID terminology and definitions used throughout the product documentation.

Using UNIX Commands

This document does not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ operating system 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 3000 family documentation is available online in both PDF and HTML format at the following location:

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

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

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

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 xxxii.

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>

Section 508 Accessibility Features

The Sun StorEdge 3000 family 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 web sites identified in “Accessing Sun Documentation” on page xxxii. 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 4.2x User's Guide*, part number 817-3711-18.

Before You Begin

Firmware is the intelligence behind a RAID controller. It provides the underlying functionality of the controller, which is presented directly by the firmware menu options and is also used by the command-line interface (CLI), Sun StorEdge Configuration Service, and third-party applications that directly or indirectly use information passed bidirectionally through the firmware's external interface (EI).

Firmware is installed or "flashed" into the array hardware before it is shipped. At any time, you can download and install patches that include later versions of the firmware to take advantage of increased functionality.

Refer to the release notes for your array for an overview of the latest functionality as well as for instructions on how to download and install these patches. Refer to the README file associated with the firmware patch for detailed installation instructions and a list of bugs fixed by that patch.

This manual applies to all Sun StorEdge 3000 family RAID arrays with 4.2x controller firmware:

- Sun StorEdge 3510 FC array
- Sun StorEdge 3511 SATA array
- Sun StorEdge 3310 SCSI array
- Sun StorEdge 3320 SCSI array

However, each platform has its own firmware patch. When you upgrade your firmware, be sure to download and install the proper patch.

Do not attempt to install a patch meant for one platform on a platform of a different type. See "Supported Hardware Platforms" on page 2 for information about the hardware platforms supported by this RAID firmware release.

Several Sun StorEdge 3000 family arrays are also available without firmware; these are connected to a host computer and treated as Just a Bunch of Disks (JBODs). JBODs are managed directly by the host computer's management software and should not be confused with RAID arrays or RAID expansion units, even though their product number designations and appearances may be similar or identical.

Before using the RAID controller firmware, it is important to understand some key concepts underlying the controller's functionality. These concepts are relatively common in storage arrays from many vendors, but may be implemented differently in the Sun StorEdge 3000 family of RAID arrays. This chapter presents an overview of these key concepts. More detailed information about the way these concepts are implemented and used appears later in this guide.

Topics covered in this chapter include:

- "Supported Hardware Platforms" on page 2
- "Key Concepts" on page 4
 - "RAID Planning Considerations" on page 4
 - "Local and Global Spare Drives" on page 5
 - "Using Both Local and Global Spare Drives" on page 7
- "Accessing the Firmware Application" on page 8
 - "Setting an IP Address" on page 8

Supported Hardware Platforms

Four different Sun StorEdge 3000 family arrays feature RAID firmware 4.2x:

- Sun StorEdge 3510 FC array

The Sun StorEdge 3510 FC array is a next-generation Fibre Channel storage system designed to provide direct attached storage (DAS) to entry-level, mid-range, and enterprise servers, or to serve as the disk storage within a storage area network (SAN). This solution features powerful performance and reliability, availability, and serviceability (RAS) features using modern FC technology. As a result, the Sun StorEdge 3510 FC array is ideal for performance-sensitive applications and for environments with many entry-level, mid-range, and enterprise servers, such as:

- Internet
- Messaging
- Database
- Technical
- Imaging

- Sun StorEdge 3511 SATA array

The Sun StorEdge 3511 SATA array shares many features in common with the Sun StorEdge 3510 FC array, but includes internal circuitry that enables it to use low-cost, high-capacity Serial ATA drives. It is best suited for inexpensive secondary storage applications that are not mission-critical where higher capacity drives are needed, and where lower performance and less than 7/24 availability is an option. This includes near-line applications such as:

- Information life cycle management
 - Content addressable storage
 - Backup and restore
 - Secondary SAN storage
 - Near-line DAS storage
 - Static reference data storage
- Sun StorEdge 3310 SCSI array and Sun StorEdge 3320 SCSI array

The Sun StorEdge 3310 SCSI RAID array supports up to two expansion chassis (expansion unit arrays that have a set of drives and no controller) for a total of 36 drives. The RAID array and expansion units connect to the storage devices and consoles by means of standard serial port, Ethernet, and SCSI connections.

The Sun StorEdge 3320 SCSI RAID array supports up to two expansion chassis (expansion unit arrays that have a set of drives and no controller) for a total of 36 drives. The RAID array and expansion units connect to the storage devices and consoles by means of standard serial port, Ethernet, and SCSI connections. This array is similar to the Sun StorEdge 3310 SCSI array except that it uses Ultra-320 SCSI drives.

All of these arrays are rack-mountable, Network Equipment Building System (NEBS) Level 3-compliant, Fibre Channel mass storage subsystems. NEBS Level 3 is the highest level of NEBS criteria used to assure maximum operability of networking equipment in mission-critical environments such as telecommunications central offices.

In addition to the arrays mentioned above, one mixed-platform configuration is supported:

- Sun StorEdge 3511 SATA expansion units connected to a Sun StorEdge 3510 FC RAID array.

This special-purpose configuration, either alone or in combination with Sun StorEdge 3511 SATA expansion units, is described in the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the Sun StorEdge 3510 FC array and the Sun StorEdge 3511 SATA array.

Key Concepts

The following section briefly outlines several key concepts:

- “RAID Planning Considerations” on page 4
- “Local and Global Spare Drives” on page 5
- “Using Both Local and Global Spare Drives” on page 7

Further details are presented later in this guide where the appropriate menu options are described.

See also:

- “Logical Drives” on page 312
- “Logical Volumes” on page 313
- “Channels, Partitions, and LUN Mapping” on page 313

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. You can add expansion units if you need more drives.
- 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 Configurations” on page 21.

- 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 optimization mode is applied to all logical drives. You cannot change the

optimization mode until data is backed up, all logical drives are deleted, and the array is restarted. You can, however, change the stripe size for individual logical drives at the time you create them.

Note – Default stripe sizes result in optimal performance for most applications. Selecting a stripe size that is inappropriate for your optimization mode and RAID level can decrease performance significantly. For example, smaller stripe sizes are ideal for I/Os that are transaction-based and randomly accessed. But when a logical drive configured with a 4-Kbyte stripe size receives files of 128 Kbyte, each physical drive has to write many more times to store it in 4-Kbyte data fragments. Change stripe size only when you are sure it will result in performance improvements for your particular applications.

See “Cache Optimization Mode and Stripe Size Guidelines” on page 30 for more information.

- 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 must 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.

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.

- 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.

A 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 a drive fails in a 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 180.



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.

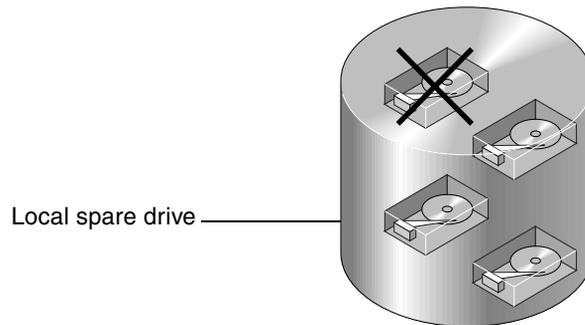


FIGURE 1-1 Local (Dedicated) Spare

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.

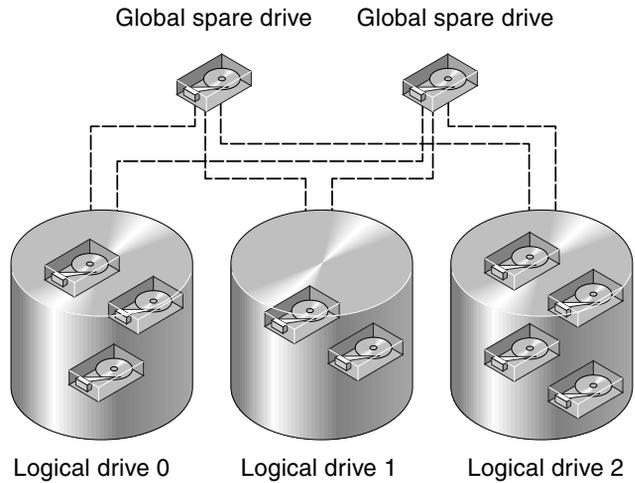


FIGURE 1-2 Global Spare

Using Both Local and Global Spare Drives

In FIGURE 1-3, 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.

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.

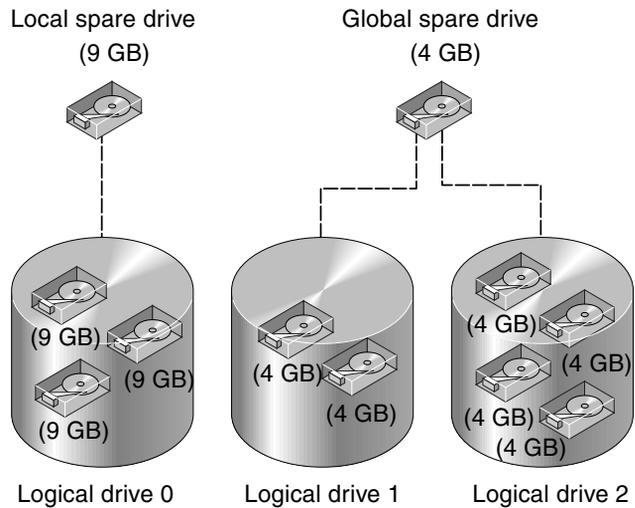


FIGURE 1-3 Mixing Local and Global Spares

In FIGURE 1-3, 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 automatically.

Accessing the Firmware Application

You can access the controller firmware by connecting an RS-232 port on your host to an RS-232 port on your RAID controller with the null-modem cable supplied with your array. The “Connecting Your Array” chapter of the Installation, Operation, and Service Manual for your array contains instructions for setting up communications once this connection is made. Platform-specific instructions are found in the appendix that is appropriate for your hardware and operating system.

You can also access the controller firmware through telnet sessions. The default TCP/IP connection method is to use the IP address, gateway, and netmask assigned by a Dynamic Host Configuration Protocol (DHCP) server. If your network has a DHCP server, you can access the controller’s Ethernet port using that IP address without having to set up the RS-232 port connection described above. The “Connecting Your Array” chapter of the Installation, Operation, and Service Manual for your array contains a full description of the various in-band and out-of-band connections available to you.

Setting an IP Address

To access the array using the Ethernet port, the controller must have an IP address. The default setting uses the Dynamic Host Configuration Protocol (DHCP) to automatically assign an IP address if you have a DHCP server on your network and DHCP support is enabled.

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.

Note – If you assign an IP address to an array to manage it out-of-band, for security reasons consider using an IP address on a private network rather than a publicly routable network. Using the controller firmware to set a password for the controller limits unauthorized access to the array. Changing the firmware’s Network Protocol Support settings can provide further security by disabling the ability to remotely connect to the array using individual protocols such as HTTP, HTTPS, telnet, FTP, and SSH. See “Communication Parameters” on page 212 for more information.

▼ 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.**
Refer to the “Connecting Your Array” chapter of the Installation, Operation, and Service manual for your array for information about the communication parameters to use to ensure communication. Refer to the “Configuring a Sun Server Running the Solaris Operating System” appendix in the same document if you want to configure a tip session to use the COM port.
2. **Choose “view and edit Configuration parameter →Communication Parameters → Internet Protocol (TCP/IP).”**
3. **Select the chip hardware address.**
4. **Choose “Set IP Address →Address.”**
5. **Configure the Ethernet port.**

Note – If your network uses a DHCP or RARP server to automatically provide IP addresses, you can use one of these alternatives to manually configuring your IP address. To configure the port to accept an IP address from a DHCP server, type DHCP and press Return. To configure the port as a RARP client, type RARP and press Return. To disable the LAN port and set all three of the selected LAN port’s fields to Not Set, delete any contents from the Address field and press Return.

6. **If you are manually configuring the LAN port’s IP address:**
 - a. **Type an IP address in the text box and press Return**
 - b. **Choose “Netmask.”**
 - c. **Type the correct netmask for the port in the text box and press Return.**

d. Choose "Gateway."

e. Type the correct gateway IP address for the port and press Return.

7. Press Escape to continue.

A confirmation prompt is displayed.

Change/Set IP Address ?

8. Select Yes to change the address, or No to leave the existing address.

A confirmation prompt informs you that a controller reset is necessary for the new IP address to take effect and asks if you want to reset the controller now.

9. Select Yes to reset the controller.

Basic Firmware Components

This chapter introduces the initial firmware screen, menu structure, and navigation and screen conventions.

The Sun StorEdge 3310 SCSI array, Sun StorEdge 3320 SCSI array, Sun StorEdge 3510 FC array, and Sun StorEdge 3511 SATA array share the same firmware. However, the screens that are displayed and, to a lesser extent, the menu options vary between the different types of array. As a result, some of the examples in the manual might differ from what you see for your array.

Topics covered in this chapter include:

- “Viewing the Initial Firmware Screen” on page 11
- “Navigating Firmware Menus” on page 15
 - “Navigation Terminology and Conventions” on page 17
 - “Menu Options That Toggle Between Conditions” on page 18
- “Progress Indicators” on page 18
- “Device Capacities” on page 19

Viewing the Initial Firmware Screen

The initial firmware screen is displayed, as shown in FIGURE 2-1, when you power on the RAID controller and access the firmware application.

If an event message is displayed, press Escape after you read it to clear it from the screen or Ctrl-C to clear all messages.

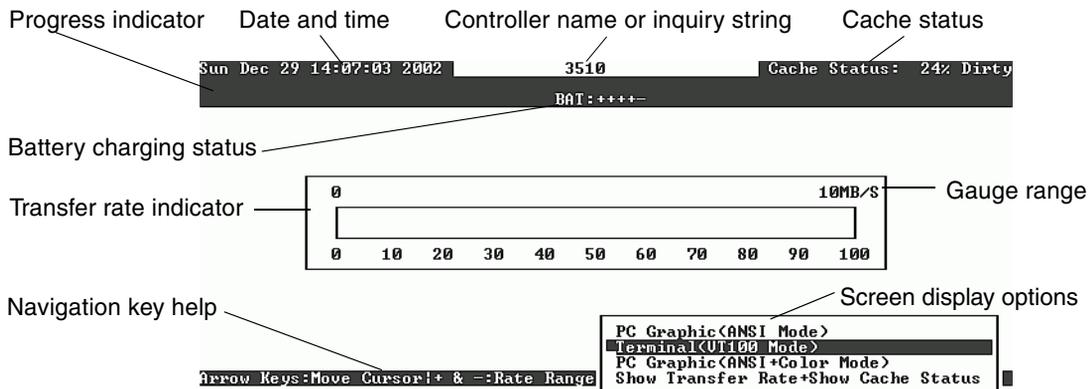


FIGURE 2-1 Initial Firmware Screen

The following table describes the components in the initial firmware screen.

TABLE 2-1 Firmware Screen Components

Component	Description
Date and time	Shows the controller date and time.
Controller name or inquiry string	Identifies the type of controller. This is also where the controller name entered using “view and edit Configuration parameters → Controller Parameters → Controller Name” is displayed.
Cache status	Indicates the percentage of controller cache that differs from what is saved to disk. “Clean” means that all cache has been saved to disk. “WT” indicates that an event trigger has disabled writeback cache and cache status is in write-through mode until the environmental event that caused the event trigger is corrected, at which time cache status will switch back automatically to writeback. See FIGURE 2-2.
Battery charging status	Battery status ranges from BAD to ---- (charging) to +++++ (fully charged). See “Battery Operation” on page 291.
Transfer rate indicator	Indicates the current data transfer rate of communication between the array and the connected hosts. Select “Show Transfer Rate+Show Cache Status” from the list of screen display options and press + or - to increase or decrease the transfer rate Gauge range from its default of 10 Mbyte per second.

TABLE 2-1 Firmware Screen Components (*Continued*)

Component	Description
Progress indicator	Indicates the progress of various tasks as a percentage of completion. The percentage is preceded (?) by an abbreviation of the task being performed: <ul style="list-style-type: none">• i = on-line initialization progress• I = off-line initialization progress• R = rebuild progress• A = add physical drive progress• P = regenerate parity progress• S = media scan progress• E = off-line expansion progress*• e = on-line expansion progress†
Gauge range	Use + or – keys to change the gauge range of the transfer rate indicator when “ Show Transfer Rate → Show Cache Status ” is selected. The default is 10 Mbyte per second
Navigation key help	Explains the keyboard combinations that enable navigation of the interface. The up and down arrow keys, Enter (or Return), Ctrl-L, and the Escape key enable navigation. For details, see “Navigating Firmware Menus” on page 15.
PC graphic (ANSI mode)	Enters the Main Menu and operates in ANSI mode.
Terminal (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	Changes the current data transfer rate of communication between the array and the connected hosts. With I/O running, select this display option, press Enter, and then press the + or - key to increase or decrease the transfer rate gauge range.

* The progress of only 8-9 processes are reported at one time. To see complete progress results, see “Progress Indicators” on page 18

† The progress of only 8-9 processes are reported at one time. To see complete progress results, see “Progress Indicators” on page 18.

If an array is connected to a host using a serial port connection and powered on, the host terminal window displays a series of messages, as shown in the following example.

```

3510          Disk Array is installed with 1024MBytes SDRAM
Total channels: 6
Channel: 0 is a host channel, id: 40
Channel: 1 is a host channel, id: 41
Channel: 2 is a drive channel, id: 14, 15
Channel: 3 is a drive channel, id: 14, 15
Channel: 4 is a host channel, id: 70
Channel: 5 is a host channel, id: 71
Scanning channels. Please wait a few moments!
Preparing to restore saved persistent reservations. Type 'skip' to
skip:

```

If an event trigger is set to disable writeback cache, when it occurs the event message displays as an alert and the Cache Status: area of the initial firmware menu displays WT in the upper right corner. Performance might be decreased while this condition continues. Once the event that triggered the alert is corrected, cache policy reverts to its default setting.

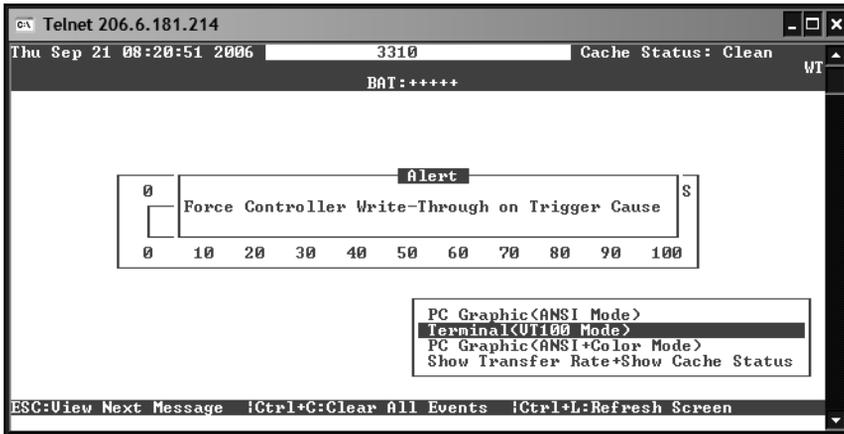


FIGURE 2-2 Write-through (WT) Cache Status forced by an event trigger

Navigating Firmware Menus

To access the firmware menu options, use the up and down arrow keys to choose a screen display mode, and then press Return to enter the Main Menu.

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

Once you have chosen the screen display mode, the Main Menu is displayed.

```
< Main Menu >
view and edit logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit 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 – Because Fibre Channel, SATA, and SCSI arrays share the same controller firmware, most menu options are the same. Parameter values might vary according to the array type, configuration, drive type, and so forth.

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

TABLE 2-2 Navigation Keys

Key	Action
←→↑ ↓	Select menu options or other items that display onscreen. Some screens may not be large enough to display all entries. Use the up or down arrows to “scroll” to these entries.
Return or Enter	Perform the selected menu option or display a submenu

TABLE 2-2 Navigation Keys (*Continued*)

Key	Action
Escape	Return to the previous menu without performing the selected menu option. For some procedures that use the Return or Enter key to mark selected items for inclusion, pressing Escape will perform the next step in the procedure. Pressing Escape also clears event messages.
Ctrl-L (Ctrl key and L key together)	Refresh the screen information
Boldface capital letter in a menu option list	Access a menu command quickly if the menu option has a single capitalized letter that is highlighted.

Note – If you are connected through a serial port connection (including a Solaris tip session) rather than through a telnet session, it is possible that when a controller is inserted, deasserted, or failed over, unwanted “garbage” characters can be displayed. This is due to the controller negotiation that occurs during a power-up or restart cycle. In most cases using the Ctrl-L keyboard shortcut shown above cleans up the extraneous characters. If this fails, the workaround is to close the tip session and start another, or use a telnet session instead.

If an array is connected to a host using a serial port connection and powered on, the host terminal window displays a series of messages, as shown in the following example.

```
3510          Disk Array is installed with 1024MBytes SDRAM
Total channels: 6
Channel: 0 is a host channel, id: 40
Channel: 1 is a host channel, id: 41
Channel: 2 is a drive channel, id: 14, 15
Channel: 3 is a drive channel, id: 14, 15
Channel: 4 is a host channel, id: 70
Channel: 5 is a host channel, id: 71
Scanning channels. Please wait a few moments!
Preparing to restore saved persistent reservations. Type 'skip' to
skip:
```

Do not use the skip option shown at the bottom of the example. This option is reserved for support personnel performing testing.

Note – As you perform the operations described in this guide, you might periodically see event message pop up on the screen. To dismiss an event message after you’ve read it, press Escape. To prevent event messages for displaying so that you can only read them by displaying the event message log, press Ctrl-C. You can press Ctrl-C again at any time to enable pop-up displays of event messages. “Viewing Event Logs on the Screen” on page 287 for more information about event messages.

Navigation Terminology and Conventions

The firmware procedures use terminology and character conventions to indicate a sequence of steps, a specific menu option, or a series of menu options.

TABLE 2-3 Navigation Terminology and Character Conventions

Terminology or Convention	Meaning
Choose	The term <i>choose</i> preceding a menu option indicates that the menu option should be highlighted using the arrow keys and then selected by pressing the Return (or Enter) key. Alternatively, you can choose some menu options by using the following keyboard shortcut:
Single bold-face capital letter in a menu option list	The bold-face capital letter key is a shortcut key. Press the key that corresponds to the capitalized letter to choose that menu option.
“ ” (quotation marks)	Quotation marks indicate a menu option.
“menu option 1→menu option 2 → menu option 3”	This represents a series of nested menus options that are selected with arrow keys. Press Return after each selection to access the next menu item and to complete the series.
Select	The term <i>select</i> preceding a device or other selectable entity indicates that the entity should be highlighted using the arrow keys and then selected by pressing the Return (or Enter) key. For instance, you select a physical drive to add it to a logical drive.

Menu Options That Toggle Between Conditions

Some firmware menu options display their current state. For example, many of the Configuration Parameters resemble this one:

“Auto-Assign Global Spare Drive - Disabled”

When you choose a menu option that displays its current state, you are prompted to change it by choosing Yes or leave it in its current state by choosing No. If you choose Yes, the menu option now shows the new current state. The example now shows:

“Auto-Assign Global Spare Drive - Enabled”

This “toggle-switch” behavior applies to menu options that have only two states, typically Enabled and Disabled.

Progress Indicators

Progress indicators are displayed when necessary to indicate the percentage of completion of a particular task or event. The task might be represented by a descriptive title, such as Drive Copying, or by a prefix abbreviation.

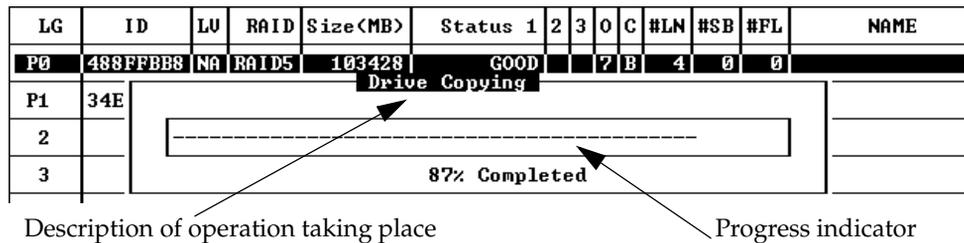


FIGURE 2-4 Progress Indicator and Descriptive Message

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

- Drive Copying
- Flash Erasing
- Flash Programming
- Rebuild Drive
- Add SCSI Drive
- Media Scan

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 TABLE 2-4.

TABLE 2-4 Progress Indicator Prefix Meanings

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

Device Capacities

Firmware screens often present the capacity of devices such as logical drives. All device capacity is displayed in powers of 1024:

- 1 Kbyte = 1024 bytes
- 1 Mbyte = 1024 Kbyte = 1,048,576 bytes
- 1 Gbyte = 1024 Mbyte = 1,073,741,824 bytes
- 1 Tbyte = 1024 Gbyte = 1,099,511,627,776 bytes

Configuration Defaults and Guidelines

This chapter lists default configurations and provides guidelines you need to be aware of when configuring your array.

This chapter covers the following topics:

- “Default Configurations” on page 21
 - “Default Logical Drive Configuration” on page 22
 - “Default Channel Configurations” on page 22
- “Maximum Drive Configurations per Array” on page 24
- “Maximum Number of Disks and Maximum Usable Capacity per Logical Drive” on page 25
- “Controller Operation Guidelines” on page 28
 - “Dual-Controller Guidelines” on page 28
 - “Single-Controller Guidelines” on page 29
- “Cache Optimization Mode and Stripe Size Guidelines” on page 30
- “Cache Write Policy Guidelines” on page 32
- “Fibre Connection Protocol Guidelines” on page 33
- “A Sample SAN Point-to-Point Configuration” on page 35
- “A Sample DAS Loop Configuration” on page 41
- “Array Configuration Summary” on page 45

Default Configurations

This section provides default configuration information for drives and channel settings.

Default Logical Drive Configuration

Sun StorEdge 3000 family arrays are preconfigured with a single RAID 0 logical drive mapped to LUN 0, and no spare drives. This is not a usable configuration. You must delete this logical drive and create new logical drives, as shown in “First-Time Configuration for SCSI Arrays” on page 49 and “First-Time Configuration for FC or SATA Arrays” on page 81.

Default Channel Configurations

Sun StorEdge 3000 family arrays are preconfigured with the channel settings shown in the following tables. The most common reason to change a host channel to a drive channel is to attach expansion units to a RAID array.

Sun StorEdge 3310 SCSI array default channel settings are shown in TABLE 3-1.

TABLE 3-1 Sun StorEdge 3310 SCSI Array and Sun StorEdge 3320 SCSI Array Default Channel Settings

Channel	Default Mode	Primary Controller ID (PID)	Secondary Controller ID (SID)
0	Drive Channel	6	7
1	Host Channel	0	NA
2	Drive Channel	6	7
3	Host Channel	NA	1
6	RCCOM	NA	NA

Sun StorEdge 3510 FC array default channel settings are shown in TABLE 3-2.

TABLE 3-2 Sun StorEdge 3510 FC Array Default Channel Settings

Channel	Default Mode	Primary Controller ID (PID)	Secondary Controller ID (SID)
0	Host Channel	40	NA
1	Host Channel	NA	42
2	Drive Channel + RCCOM	14	15

TABLE 3-2 Sun StorEdge 3510 FC Array Default Channel Settings *(Continued)*

Channel	Default Mode	Primary Controller ID (PID)	Secondary Controller ID (SID)
3	Drive Channel + RCCOM	14	15
4	Host Channel	44	NA
5	Host Channel	NA	46

Sun StorEdge 3511 SATA array default channel settings are shown in TABLE 3-3.

TABLE 3-3 Sun StorEdge 3511 SATA Array Default Channel Settings

Channel	Default Mode	Primary Controller ID (PID)	Secondary Controller ID (SID)
0	Host Channel	40	NA
1	Host Channel	NA	42
2	Drive Channel + RCCOM	14	15
3	Drive Channel + RCCOM	14	15
4	Host Channel	44	NA
5	Host Channel	NA	46

Maximum Drive Configurations per Array

TABLE 3-4 lists the maximum number of physical and logical drives, partitions per logical drive and logical volume, and maximum number of logical unit number (LUN) assignments for each array.

TABLE 3-4 Maximum Number of Supported Physical and Logical Drives, Partitions, and LUN Assignments

Array	Physical Drives	Logical Drives	Partitions per Logical Drive	Partitions per Logical Volume	LUN Assignments
Sun StorEdge 3310 SCSI array and Sun StorEdge 3320 SCSI array	36 (1 array and 2 expansion units)	16	32	32	128
Sun StorEdge 3510 FC array	108 (1 array and 8 expansion units)	32	32	32	128 (point-to-point mode) 64 (point-to-point mode, redundant configuration) 1024 (loop mode) 512 (loop mode, redundant configuration)
Sun StorEdge 3511 SATA array	72 (1 array and 5 expansion units)	32	32	32	128 (point-to-point mode) 64 (point-to-point mode, redundant configuration) 1024 (loop mode) 512 (loop mode, redundant configuration)
Sun StorEdge 3510 FC array with Sun StorEdge 3511 SATA expansion units ¹	72 (1 array and 5 expansion units)	32	32	32	128 (point-to-point mode) 64 (point-to-point mode, redundant configuration) 1024 (loop mode) 512 (loop mode, redundant configuration)

¹ Sun StorEdge 3511 SATA expansion units can be connected to a Sun StorEdge 3510 FC array, either alone or in combination with Sun StorEdge 3510 FC expansion units

Maximum Number of Disks and Maximum Usable Capacity per Logical Drive

The following tables show the maximum number of disks per logical drive, and the maximum usable capacity of a logical drive, depending on RAID level and optimization mode.

The maximum capacity per logical drive supported by the RAID firmware is:

- 16 Tbyte with random optimization
- 64 Tbyte with sequential optimization

Actual logical drive maximum capacities are usually determined by practical considerations or the amount of disk space available.



Caution – In FC and SATA configurations with large drive capacities, the size of the logical drive might exceed the device capacity limitation of your operating system. Be sure to check the device capacity limitation of your operating system before creating the logical drive. If the logical drive size exceeds the capacity limitation, you must partition the logical drive.

TABLE 3-5 shows the usable capacity of the drives available in Sun StorEdge 3000 family arrays.

Note – The 250 Mbyte of reserved space on each drive used for storing controller metadata is not included in this table, since it is not available for storing data.

TABLE 3-5 Actual Capacities per Drive

Drive Size	Usable Capacity (Mbyte)
36 Gbyte	34,482
73 Gbyte	69,757
146 Gbyte	139,759
250 Gbyte	238,216
300 Gbyte	285,852
400 Gbyte	381,291

TABLE 3-6 shows the maximum usable storage capacity for Sun StorEdge 3310 SCSI arrays, Sun StorEdge 3320 SCSI arrays, Sun StorEdge 3510 FC arrays, and Sun StorEdge 3511 SATA arrays, using the maximum number of expansion units, fully populated with the largest currently available drives.

TABLE 3-6 Maximum Usable Storage Capacity Determined by RAID Level

Array	Number of Disks	Drive Size	RAID 0 (Tbyte)	RAID 1 (Tbyte)	RAID 3 or RAID 5 (Tbyte)
Sun StorEdge 3310 SCSI array and Sun StorEdge 3320 SCSI array	36	300 Gbyte	9.81	4.90	9.54
Sun StorEdge 3510 FC array	108	146 Gbyte	14.39	7.20	14.26
Sun StorEdge 3511 SATA array	72	400 Gbyte	26.18	13.09	25.82

TABLE 3-7 shows the maximum number of disks that can be used in a single logical drive, based upon the drive size, and the optimization method chosen.

TABLE 3-7 Maximum Number of Disks per Logical Drive

Drive Size	SCSI (Random and Sequential Optimization)	FC (Random or Sequential Optimization)	SATA (Random Optimization)	SATA (Sequential Optimization)
36 Gbyte	36	108	N/A	N/A
73 Gbyte	36	108	N/A	N/A
146 Gbyte	36	108	N/A	N/A
250 Gbyte	n/a	n/a	66	72
300 Gbyte	36	55 random 108 sequential	N/A	N/A
400 Gbyte	n/a	n/a	41	72

Note – Except for SATA arrays using random optimization, it is possible (though impractical) to employ all available disks in a single logical drive

TABLE 3-8 shows the maximum usable capacity of a single logical drive in a Sun StorEdge 3510 FC array, depending on drive size.

TABLE 3-8 Maximum Usable Capacity (in Gbyte) per Sun StorEdge 3510 FC Logical Drive

Drive Size	RAID 0	RAID 1	RAID 3/5
36 Gbyte	3636	1818	3603
73 Gbyte	7357	3678	7289
146 Gbyte	14740	7370	14603
300 Gbyte	30148	15074	29869

TABLE 3-9 shows the maximum usable capacity of a single logical drive in a Sun StorEdge 3310 SCSI array, depending on drive size.

TABLE 3-9 Maximum Usable Capacity (in Gbyte) per Sun StorEdge 3310 SCSI and Sun StorEdge 3320 SCSI Logical Drive

Drive Size	RAID 0	RAID 1	RAID 3 or RAID 5
36 Gbyte	1212	606	1178
73 Gbyte	2452	1226	2384
146 Gbyte	4913	2456	4776
300 Gbyte	10049	5024	9770

TABLE 3-10 shows the maximum usable capacity of a single logical drive in a Sun StorEdge 3511 SATA array, depending on drive size.

TABLE 3-10 Maximum Usable Capacity (in Gbyte) per Sun StorEdge 3511 SATA Logical Drive

Drive Size	RAID 0 (Random)	RAID 0 (Sequential)	RAID 1 (Random)	RAID 1 (Sequential)	RAID 3 or RAID 5 (Random)	RAID 3 or RAID 5 (Sequential)
250 Gbyte	15353	16749	7676	8374	15121	16516
400 Gbyte	15266	26809	7633	13404	14894	26437

Controller Operation Guidelines

This section provides guidelines for dual-controller and single-controller operation.

Dual-Controller Guidelines

Keep the following operation details in mind when configuring a dual-controller array.

- The controller firmware assumes that two controllers are available, or might be made available during operation at any time. In a one-rack-unit (1U) single-controller configuration, a two-rack-unit (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 shows 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, the controllers autonegotiate and designate one controller as primary and the other controller as secondary.
- The two controllers behave as one primary controller. Once redundancy takes effect, configuration can be applied only to 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.



Caution – Major upgrades of controller firmware, or replacing a controller with one that has a significantly different version of firmware, might involve differences in non-volatile RAM (NVRAM) that require following special upgrade procedures. For more information, refer to the *Sun StorEdge 3000 Family FRU Installation Guide* and to the release notes for your array.

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

- 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 and 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 drives 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 selected. In this configuration, assigning all the logical drives to one controller means that the other controller remains idle, becoming active only if the primary controller fails.

Single-Controller Guidelines

Keep the following operation details in mind when configuring a single-controller array.

- The controller must be the primary controller or the controller cannot operate. The primary controller controls all logical drive and firmware operations. Keep the controller as the primary controller at all times and assign all logical drives to the primary controller.

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

- In a single-controller configuration, disable the Write-Back Cache feature to avoid the possibility of data corruption in the event of a controller failure. This has a negative effect on performance. To avoid either issue, use dual controllers.

Using two single controllers in a clustering environment with host-based mirroring provides some of the advantages of using a dual controller. However you still need to disable the Write-Back Cache in case one of the single controllers fails, to avoid the risk of data corruption. For this reason, a dual controller configuration is preferable.

- For a single-controller configuration, the Peripheral Device Status shows a status of 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.
- If you are using a single controller, save your NVRAM after any configuration changes so that you can restore it in the event of a controller failure and replacement. See “Saving Your Configuration (NVRAM) to Disk” on page 281 and “Restoring Your Configuration (NVRAM) From Disk” on page 284 for more information.



Caution – Major upgrades of controller firmware, or replacing a controller with one that has a significantly different version of firmware, might involve differences in non-volatile RAM (NVRAM) that require following special upgrade procedures. For more information, refer to the *Sun StorEdge 3000 Family FRU Installation Guide* and to the release notes for your array.

- Also keep a written record of your configuration, including the firmware version number, using tables similar to those found in Appendix C. Major firmware upgrades or downgrades can require manually recreating your configuration from these records, since no synchronization from a second controller is possible. Refer to the *Sun StorEdge 3000 Family FRU Installation Guide* for information about replacing controllers.

Cache Optimization Mode and Stripe Size Guidelines

Before creating or modifying logical drives, determine the appropriate optimization mode for the RAID array. The controller supports two optimization modes, sequential I/O and random I/O. Sequential I/O is the default mode.

Note – Due to firmware improvements beginning with version 4.11, sequential optimization yields better performance than random optimization for most applications and configurations. Use sequential optimization unless real-world tests in your production environment show better results for random optimization.

When you specify sequential or random cache optimization, the controller determines a default stripe size for newly-created logical drives. But you can specify whatever stripe size you choose for each logical drive when you create it, enabling you to maximize performance by matching stripe size with your application requirements. Since different applications may use different logical drives, this functionality provides you with greatly increased flexibility.

- For sequential optimization, available stripe size choices include 16 Kbyte, 32 Kbyte, 64 Kbyte, 128 Kbyte, and 256 Kbyte.
- The default stripe size for sequential optimization is 128 Kbyte for all logical drives except RAID 3, which is 16 Kbyte.
- For sequential optimization, the cache block size is 128 Kbyte.
- For random optimization, available stripe size choices include 4 Kbyte, 8 Kbyte, 16 Kbyte, 32 Kbyte, 64 Kbyte, 128 Kbyte, and 256 Kbyte.
- The default stripe size for random optimization is 32 Kbyte for all logical drives except RAID 3, which is 4 Kbyte.
- For random optimization, the cache block size is 32 Kbyte.

See “Cache Optimization Mode (SCSI)” on page 52 for information about how to set the cache optimization mode on a Sun StorEdge 3310 SCSI array or Sun StorEdge 3320 SCSI array. See “Cache Optimization Mode (FC and SATA)” on page 84 for information about how to set the cache optimization mode for a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array.

The RAID array’s cache optimization mode determines the cache block size used by the controller for all logical drives. An appropriate cache block size improves performance when a particular application uses either large or small stripe sizes:

- Video playback, multimedia post-production audio and video editing, and similar applications read and write large files in sequential order.
- Transaction-based and database update applications read and write small files in random order.

Once logical drives are created, you cannot use the RAID firmware’s Optimization for Random I/O or Optimization for Sequential I/O menu option to change the optimization mode without deleting all logical drives. You can, however, use the Sun StorEdge CLI `set cache-parameters` command to change the optimization mode while logical drives exist. Refer to the *Sun StorEdge 3000 Family CLI 2.0 User’s Guide* for more information.

Note – Using the Sun StorEdge CLI `set cache-parameters` command to change optimization mode can result in a pre-existing logical drive having a stripe size that, because it is inappropriate for that optimization mode, could not have been selected at the time the logical drive was created. This combination will not yield the best performance possible, but there is no risk of data loss or other data-related problems. You can avoid this inefficiency by choosing stripe sizes and an optimization mode that are appropriate for your applications.

Since the cache block size works in conjunction with stripe size, the optimization mode you choose determines default logical drive stripe sizes that are consistent with the cache block size setting. But you can now fine-tune performance by specifying each logical drive’s stripe size so that it matches your application needs, using a firmware menu option that is available at the time you create the logical drive. See “Cache Optimization Mode and Stripe Size Guidelines” on page 30 for more information.

Note – Once the stripe size is selected and data is written to logical drives, the only way to change the stripe size of an individual logical drive is to back up all its data to another location, delete the logical drive, and create a logical drive with the stripe size that you want.

See “(Optional) Configure the logical drive stripe size.” on page 67 for information about how to set the stripe size for a logical drive you are creating on a Sun StorEdge 3310 SCSI array or Sun StorEdge 3320 SCSI array. See “(Optional) Configure the logical drive stripe size.” on page 105 for information about how to set the stripe size for a logical drive you are creating on a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array.

Cache Write Policy Guidelines

The cache write policy determines when cached data is written to the disk drives. The ability to hold data in cache while it is being written to disk can increase storage device speed during sequential reads. Write policy options include write-through and write-back.

When write-through cache is specified, the controller writes the data to the disk drive before signaling the host operating system that the process is complete. Write-through cache has slower write operation and throughput performance than write-back cache, but it is safer, with minimum risk of data loss on power failure. Because a battery module is installed, power is supplied to the data cached in memory and the data can be written to disk after power is restored.

When write-back cache is specified, the controller receives the data to write to disk, stores it in the memory buffer, and immediately sends the host operating system a signal that the write operation is complete, before the data is actually written to the disk drive. Write-back caching improves the performance of write operations and the throughput of the controller card.

Write-back cache is enabled by default. When you disable write-back cache, write-through cache is automatically enabled. The setting you specify becomes the default global cache setting for all logical drives. With RAID firmware version 4.11 and later, the cache setting can now be individually tailored for each logical drive. When you configure a logical drive, you can set its individual cache write policy to default, write-back, or write-through.

If you specify default for an individual logical drive, the global write policy is assigned to it. Then, if the global cache write policy that applies to the entire RAID array is changed, any logical drive that has been assigned the default setting write policy is also changed.

If you specify write-back or write-through for an individual logical drive, the cache write policy for that drive remains the same regardless of any changes to the global cache write policy.

If you have specified a global write-back policy, you can also configure the RAID array to automatically change from a write-back cache policy to a write-through cache policy when one or more of the following trigger events occur:

- Controller failure
- Battery-backup unit failure or battery not fully charged
- Power supply failure
- Fan failure

Once the condition that led to the trigger event is rectified, the cache policy automatically returns to its previous setting. For more information on configuring the write policy to automatically switch from write-back cache to write-through cache, see “Event Trigger Operations” on page 267.

Fibre Connection Protocol Guidelines

Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA arrays support the following connection protocols:

- Point-to-Point

This protocol can be used only with a switched fabric network configuration, also called a storage area network (SAN). The point-to-point protocol supports full duplex communication, but only allows one ID per channel.

- Fibre Channel–Arbitrated Loop (FC-AL), also known as loop mode

The loop mode can be used with direct-attached storage (DAS) or SAN configurations. Loop mode supports only half-duplex communication, but allows up to eight IDs per channel.

The following guidelines apply when implementing point-to-point configurations and connecting to fabric switches.

- The default mode is Loop only. If you prefer to use a point-to-point configuration, change the Fibre Channel Connection mode to Point-to-point only. See “Fibre Connection Protocol” on page 97 for information about how to change this setting.

Note – If you connect to a fabric switch without changing the default loop mode, the array automatically shifts to public loop mode. As a result, communication between the array and the switched fabric runs in half-duplex (send or receive) mode instead of providing the full-duplex (send and receive) performance of point-to-point mode.

- Check the host IDs on all the channels to ensure that there is only one ID per channel (on the primary controller or on the secondary controller) for point-to-point mode. When viewing the host IDs, there should be one primary controller ID (PID) or one secondary controller ID (SID); the alternate port ID should display N/A. Proper point-to-point mode allows only one ID per channel.
- If you change the mode to Point-to-point only and attempt to add a second ID, the controller does not allow you to add an ID to the same channel. For example, if you have ID 40 on CH 0 PID, and N/A on CH 0 SID, the controller does not allow you to add another PID to CH 0.

The controller displays a warning if the user is in point-to-point mode and tries to add an ID to the same channel but on the other controller. The warning is displayed because you have the ability to disable the internal connection between the channels on the primary and secondary controller using the `set inter-controller link` CLI command and, by doing this, you can have one ID on the primary and another ID on the secondary as a legal operation.

However, if you ignore this warning and add an ID to the other controller, the RAID controller does not allow a login as an FC-AL port because this would be illegal in a point-to-point configuration.

- You can add up to eight IDs per channel (four IDs on each controller), which forces the fabric switch port type to become FC-AL (Loop). To ensure F-port behavior (full fabric/full duplex) when attaching to a switch, only one ID can be present on each channel and the array port must be set to point-to-point mode.
- With four host channels and four host IDs, you should load-balance the host ID setup so that half the IDs are on the primary controller and half the IDs are on the secondary controller. When setting up LUNs, map each LUN to either two PIDs or two SIDs. For example, to provide redundancy, map half of the LUNs across Channel 0 (PID 40) and Channel 4 (PID 42), and then map the other half of your LUNs across Channel 1 (SID 41) and Channel 5 (SID 43). The hosts are in turn dual-pathed to the same two switched fabrics.
- Point-to-point mode enables a maximum of 128 LUNs per array. In a redundant configuration, 32 LUNs are dual-mapped across two channels on the primary controller, and another 32 LUNs are dual-mapped across the secondary controller, for a total of 64 distinct LUNs.

To use more than 64 LUNs, you must change to Loop only mode, add host IDs to one or more channels, and add 32 LUNs for each host ID.

Note – When in loop mode and connected to a fabric switch, each host ID is displayed as a loop device on the switch so that, if all 16 IDs are active on a given channel, the array looks like a loop with 16 nodes attached to a single switch FL port.

Note – In public loop mode, the array can have a maximum of 1024 LUNs, where 512 LUNs are dual-mapped across two channels, primary and secondary controller, respectively.

A Sample SAN Point-to-Point Configuration

A point-to-point configuration has the following characteristics:

- In SAN configurations, the switches communicate with the host ports on a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array using a fabric point-to-point (F_port) mode.
- When you use fabric point-to-point (F_port) connections between a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array and fabric switches, the maximum number of LUNs is limited to 128 LUNs for a nonredundant configuration and 64 LUNs for a redundant configuration.
- Fibre Channel standards allow only one ID per port when operating point-to-point protocols, resulting in a maximum of four IDs with a maximum of 32 LUNs for each ID, and a combined maximum of 128 LUNs.
- The working maximum number of LUNs is actually 64 LUNs in a configuration where you configure each LUN on two different channels for redundancy and to avoid a single point of failure.

In a dual-controller array, one controller automatically takes over all operation of a second failed controller in all circumstances. However, when an I/O controller module needs to be replaced and a cable to an I/O port is removed, that I/O path is broken unless multipathing software has established a separate path from the host to the operational controller. Supporting hot-swap servicing of a failed controller requires the use of multipathing software, such as Sun StorEdge Traffic Manager software, on the connected servers.

Note – Multipathing for Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA arrays is provided by Sun StorEdge Traffic Manager software. Refer to the release notes for your array for information about which versions of Sun StorEdge Traffic Manager software are supported for your host.

Remember these important considerations:

- A single logical drive can be mapped to only one controller, either the primary controller or the secondary controller.

- In a point-to-point configuration, only one host ID per channel is allowed. The host ID can be assigned to the primary controller and be a PID, or it can be assigned to the secondary controller and be a SID.
- If you have two switches and set up multipathing (to keep all logical drive connections operational for any switch failure or the removal of any I/O controller module), ensure that each logical drive is mapped to two ports, one on each I/O controller module, and on two channels. The cables from the two ports mapped to each logical drive must be cabled to two separate switches. See FIGURE 3-1 and FIGURE 3-2 for examples of this configuration.

The following figures show the channel numbers (0, 1, 4, and 5) of each host port and the host ID for each channel. N/A means that the port does not have a second ID assignment. The primary controller is the top I/O controller module, and the secondary controller is the bottom I/O controller module.

The dashed lines between two ports indicate a port bypass circuit that functions as a mini-hub and has the following advantages:

- The port bypass circuit on each channel connects the upper and lower ports on the same channel and provides access to both controllers at the same time.
- Since there are host connections to two channels, if one host connection is removed, the other host connection remains operational.
- Therefore, if you have a redundant multipathing configuration in which you have two host connections to each logical drive and one connection fails, the remaining path maintains a connection to the logical drive.

In FIGURE 3-1 and FIGURE 3-2, with multipathing software to reroute the data paths, each logical drive remains fully operational when the following conditions occur:

- One switch fails or is disconnected, and the logical drive is routed to the second switch. For example, if switch 0 fails, switch 1 automatically accesses logical drive 0 through the cabling to the lower port on PID 42.
- One I/O controller module fails, and all the host IDs for that controller are reassigned (moved) to the second I/O controller module. For example, if the upper I/O controller module is removed, host IDs 40 and 44 are automatically moved to the lower module and are managed by the second controller.
- An I/O controller module fails or one cable is removed from an I/O controller module, and all I/O traffic to the disconnected channel is rerouted through the second port/host LUN assigned to the logical drive. For example, if you remove the cable to channel 4, the data path for logical drive 1 switches to the port on channel 5.

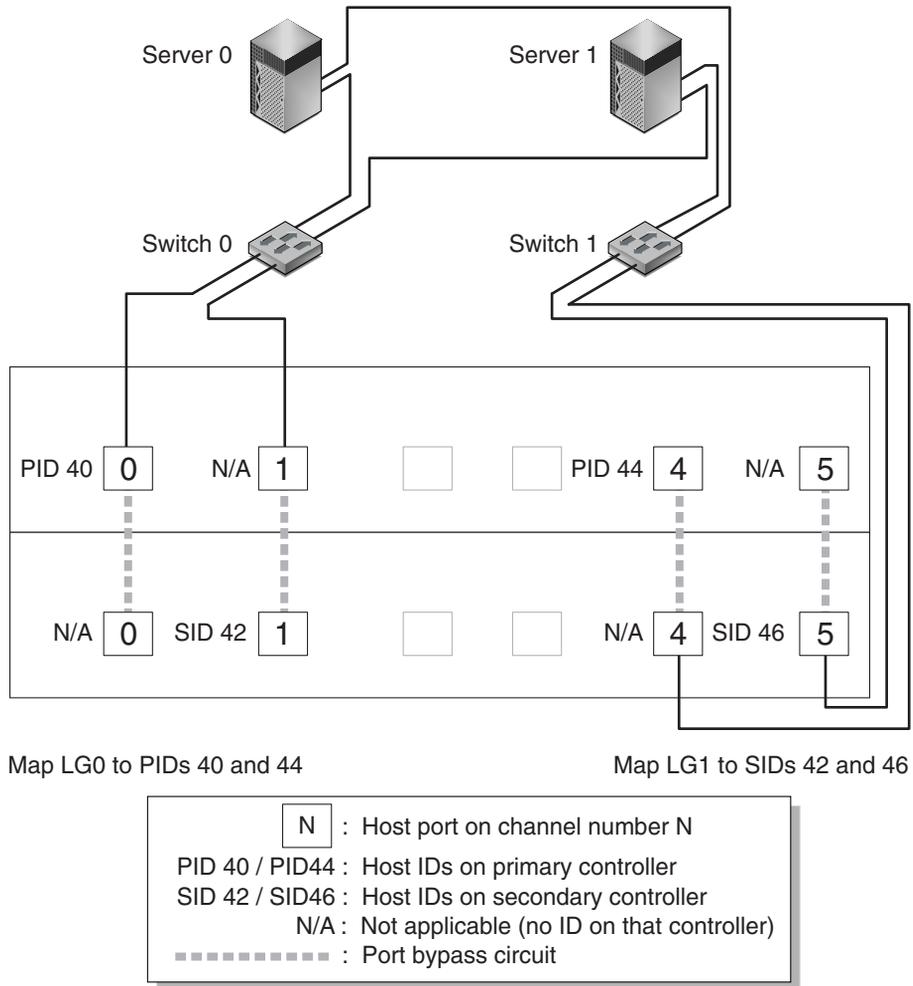


FIGURE 3-1 A Point-to-Point Configuration with a Dual-Controller Sun StorEdge 3510 FC Array and Two Switches

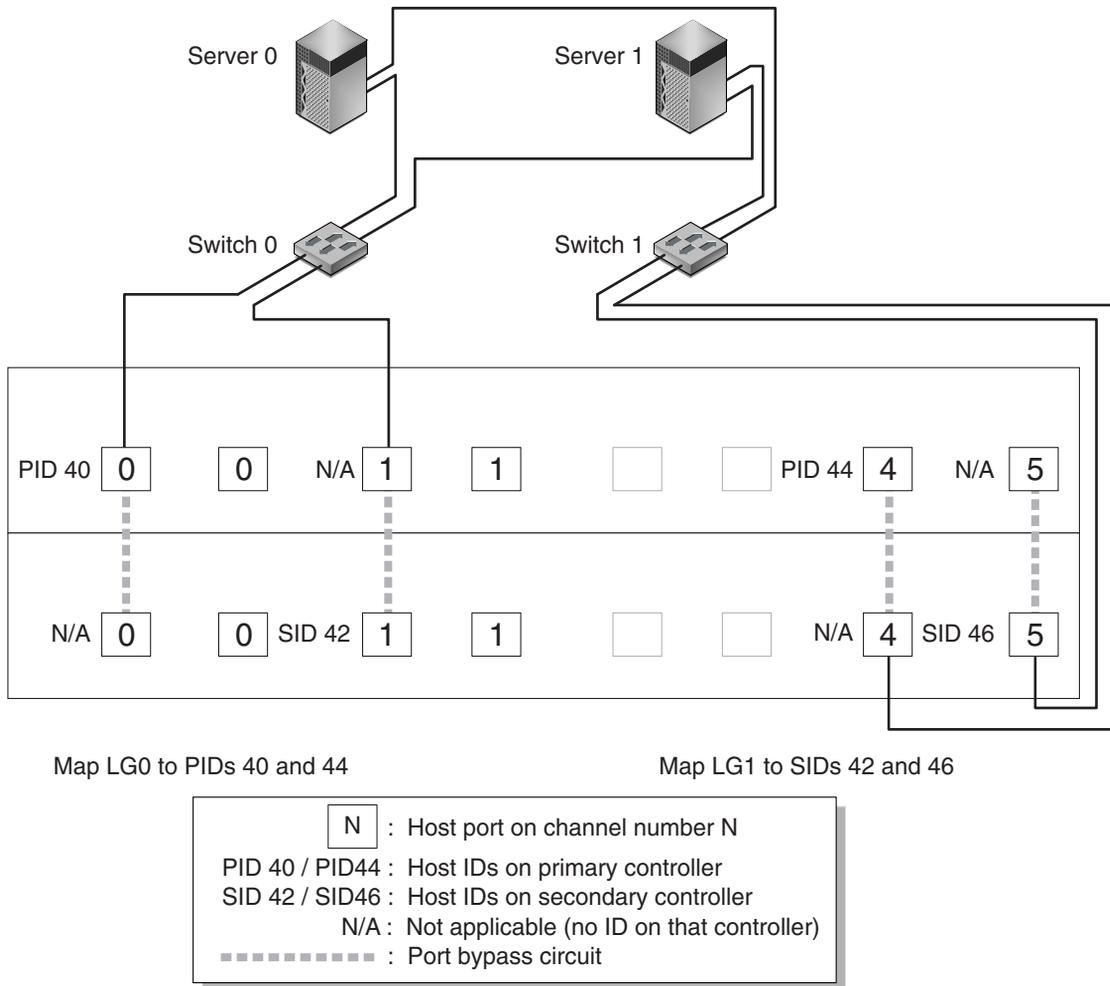


FIGURE 3-2 A Point-to-Point Configuration With a Dual-Controller Sun StorEdge 3511 SATA Array and Two Switches

Note – These illustrations show the default controller locations; however, the primary controller and secondary controller locations can occur in either slot and depend on controller resets and controller replacement operations.

TABLE 3-11 summarizes the primary and secondary host IDs assigned to logical drives 0 and 1, based on FIGURE 3-1 and FIGURE 3-2.

TABLE 3-11 Example Point-to-Point Configuration With Two Logical Drives in a Dual-Controller Array

Task	Logical Drive	LUN IDs	Channel Number	Primary ID Number	Secondary ID Number
Map 32 partitions of LG 0 to CH 0	LG 0	0-31	0	40	N/A
Duplicate-map 32 partitions of LG 0 to CH 1	LG 0	0-31	1	41	N/A
Map 32 partitions of LG 1 to CH 4	LG 1	0-31	4	N/A	50
Duplicate-map 32 partitions of LG 1 to CH 5	LG 1	0-31	5	N/A	51

▼ To Set Up a Typical Point-to-Point SAN Configuration

Perform the following steps, which are described in more detail later in this guide, to set up a typical point-to-point SAN configuration based on FIGURE 5-1 and FIGURE 5-2.

- 1. Check the position of installed small form-factor pluggable transceivers (SFPs). Move them as necessary to support the connections needed.**

You need to add SFP connectors to support more than four connections between servers and a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array. For example, add two SFP connectors to support six connections and add four SFP connectors to support eight connections.

- 2. Connect expansion units, if needed.**
- 3. Create at least two logical drives (logical drive 0 and logical drive 1) and configure spare drives.**
Leave half of the logical drives assigned to the primary controller (default assignment). Assign the other half of the logical drives to the secondary controller to load-balance the I/O.
- 4. Create up to 32 partitions (LUNs) in each logical drive.**
- 5. Change the Fibre Connection Option to "Point to point only" ("view and edit Configuration parameters →Host-side SCSI Parameters →Fibre Connections Option").**

6. **For ease of use in configuring LUNs, change the host IDs on the four channels to the following assignments:**

Channel 0: PID 40 (assigned to the primary controller)

Channel 1: PID 41 (assigned to the primary controller)

Channel 4: SID 50 (assigned to the secondary controller)

Channel 5: SID 51 (assigned to the secondary controller)

Note – Do not use the “**Loop preferred, otherwise point to point**” menu option. This command is reserved for special use and should be used only if directed by technical support.

7. **Map logical drive 0 to channels 0 and 1 of the primary controller.**

Map LUN numbers 0 through 31 to the single ID on each host channel.

8. **Map logical drive 1 to channels 4 and 5 of the secondary controller.**

Map LUN numbers 0 through 31 to the single ID on each host channel. Since each set of LUNs is assigned to two channels for redundancy, the total working maximum number of LUNs is 64 LUNs.

Note – The LUN ID numbers and the number of LUNs available per logical drive can vary according to the number of logical drives and the ID assignments you want on each channel.

9. **Connect the first switch to ports 0 and 4 of the upper controller.**
10. **Connect the second switch to ports 1 and 5 of the lower controller.**
11. **Connect each server to each switch.**
12. **Install and enable multipathing software on each connected server.**

The multipathing software prevents path failure but does not alter the controller redundancy through which one controller automatically takes over all functions of a second failed controller.

A Sample DAS Loop Configuration

The typical direct attached storage (DAS) configuration shown in FIGURE 3-3 and FIGURE 3-4 includes four servers, a dual-controller array, and two expansion units. Expansion units are optional.

Servers, as shown in FIGURE 3-3 and FIGURE 3-4, are connected to the channels shown in TABLE 3-12.

TABLE 3-12 Connection for Four Servers in a DAS Configuration

Server Number	Upper I/O Controller Module	Lower I/O Controller Module
1	0	5
2	4	1
3	5	0
4	1	4

Establishing complete redundancy and maintaining high availability requires the use of multipathing software such as Sun StorEdge Traffic Manager software. To configure for multipathing:

1. Establish two connections between each server and the array.
2. Install and enable multipathing software on the server.
3. Map the logical drive each server is using to the controller channels that the server is connected to.

DAS configurations are typically implemented using a fabric loop (FL_port) mode. A loop configuration example is described under “A Sample DAS Loop Configuration” on page 41.

Fabric loop (FL_port) connections between a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array and multiple servers allow up to 1024 LUNs to be presented to servers. For guidelines on how to create 1024 LUNs, see “Planning for 1024 LUNs on an FC or SATA Array (Optional, Loop Mode Only)” on page 163.

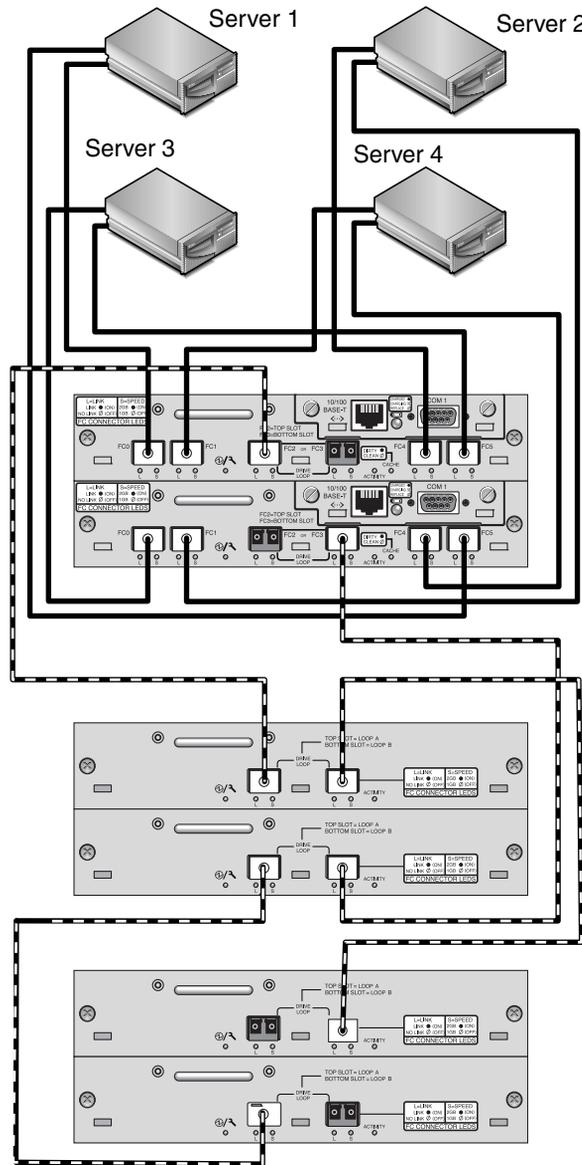


FIGURE 3-3 A DAS Configuration With Four Servers, a Dual-Controller Sun StorEdge 3510 FC Array, and Two Expansion Units

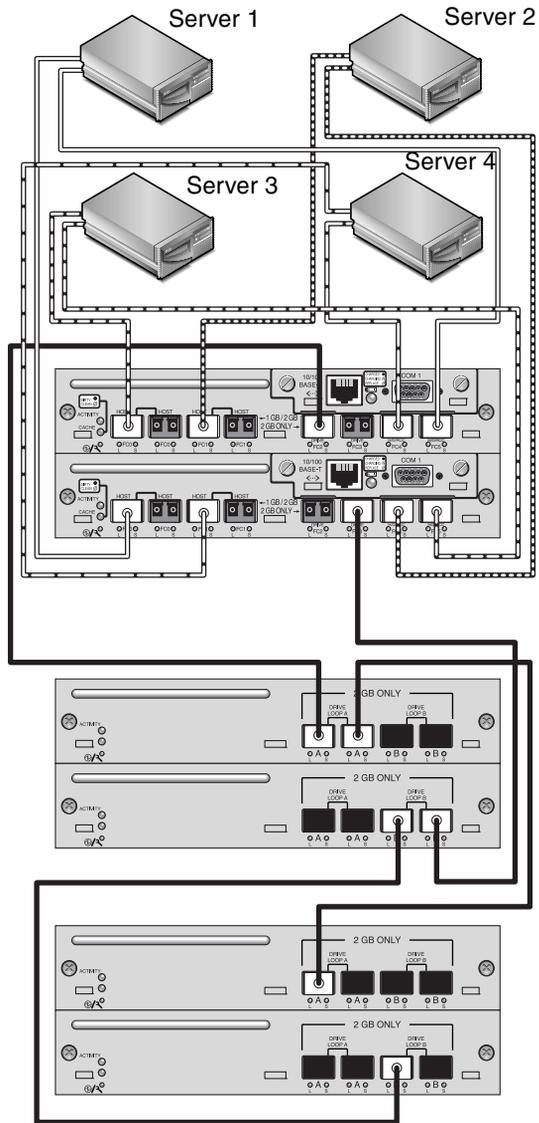


FIGURE 3-4 A DAS Configuration With Four Servers, a Dual-Controller Sun StorEdge 3511 SATA Array, and Two Expansion Units

▼ To Set Up a Typical DAS Loop Configuration

Perform the following steps, which are described in more detail later in this manual, to set up a DAS loop configuration based on FIGURE 3-3 and FIGURE 3-4.

1. **Check the location of installed SFPs. Move them as necessary to support the connections needed.**

You need to add SFP connectors to support more than four connections between servers and a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array. For example, add two SFP connectors to support six connections and add four SFP connectors to support eight connections.

2. **Connect expansion units, if needed.**
3. **Create at least one logical drive per server, and configure spare drives as needed.**
4. **Create one or more logical drive partitions for each server.**
5. **Confirm that the Fibre Connection Option is set to Loop only.**

Note – Do not use the “**Loop preferred, otherwise point to point**” menu option. This command is reserved for special use and should be used only if directed by technical support.

6. **Set up to eight IDs on each channel, if needed (see TABLE 3-13).**

TABLE 3-13 Example Primary and Secondary ID Numbers in a Loop Configuration With Two IDs per Channel

Channel Number	Primary ID Number	Secondary ID Number
0	40	41
1	43	42
4	44	45
5	47	46

7. **Map logical drive 0 to channels 0 and 5 of the primary controller.**
8. **Map logical drive 1 to channels 1 and 4 of the secondary controller.**
9. **Map logical drive 2 to channels 0 and 5 of the primary controller.**
10. **Map logical drive 3 to channels 1 and 4 of the secondary controller.**

11. Connect the first server to port FC 0 of the upper controller and port FC5 of the lower controller.
 12. Connect the second server to port FC 4 of the upper controller and port FC1 of the lower controller.
 13. Connect the third server to port FC 5 of the upper controller and port FC0 of the lower controller.
 14. Connect the fourth server to port FC 1 of the upper controller and port FC4 of the lower controller.
 15. Install and enable multipathing software on each connected server.
-

Array Configuration Summary

This section lists the typical sequence of steps for completing a first-time configuration of the array. For detailed steps and more information, see the referenced sections.

Typical steps for completing a first-time configuration of the array are as follows:

1. Set up the serial port connection.
2. Set an IP address for the controller.
See “Setting an IP Address” on page 213.
3. Determine whether sequential or random optimization is more appropriate for your applications and configure your array accordingly.
See “Cache Optimization Mode and Stripe Size Guidelines” on page 30 for more information. Also see “Cache Optimization Mode (SCSI)” on page 52 for information about how to configure a SCSI array’s optimization mode, or “Cache Optimization Mode (FC and SATA)” on page 84 for information about how to configure an FC or SATA array’s optimization mode.
4. Check physical drive availability.
See “To Check Physical Drive Availability” on page 54 for a SCSI array. See “Physical Drive Status” on page 86 for FC or SATA arrays.
5. (Optional) Configure host channels as drive channels.
See “Channel Settings” on page 56 for a SCSI array. See “Channel Settings” on page 88 for FC or SATA arrays.

6. For a Fibre Channel or SATA array, confirm or change the Fibre Connection Option (point-to-point or loop).
See “Fibre Connection Protocol Guidelines” on page 33 and “Fibre Connection Protocol” on page 97 for the procedure to configure the Fibre Connection protocol.
7. Revise or add host IDs on host channels.
See “To Add or Delete a Unique Host ID” on page 58 for SCSI arrays. See “To Add or Delete a Unique Host ID” on page 93 for FC or SATA arrays.
The IDs assigned to controllers take effect only after the controller is reset.
8. Delete default logical drives and create new logical drives as required.
See “Deleting Logical Drives” on page 51 and “Creating Logical Drives” on page 59 for SCSI arrays. See “Deleting Logical Drives” on page 83 and “Creating Logical Drives” on page 97 for FC or SATA arrays.
9. (Optional) In dual-controller configurations only, assign logical drives to the secondary controller to load-balance the two controllers.
See “Controller Assignment” on page 70 for a SCSI array. See “Controller Assignment” on page 108 for FC or SATA arrays.
10. (Optional) Partition the logical drives.
See “Partitions” on page 72 for SCSI arrays. See “Partitions” on page 109 for Fibre Channel and SATA arrays.
11. Map each logical drive partition to an ID on a host channel.
For more information, see “Mapping a Partition to a Host LUN” on page 74 for SCSI arrays.

Note – Each operating system has a method for recognizing storage devices and LUNs and might require the use of specific commands or the modification of specific files. Be sure to check the information for your operating system to ensure that you have performed the necessary procedures.

For information about different operating system procedures, refer to the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array.

12. (Optional) Create and apply host LUN filters to FC or SATA logical drives.
See “Mapping a Partition to a Host LUN” on page 112 for Fibre Channel and SATA arrays.
13. Reset the controller.
The configuration is complete.

14. Save the configuration to a disk.

See “Saving Configuration (NVRAM) to a Disk” on page 80.

15. Ensure that the cabling from the RAID array to the hosts is complete.

Refer to the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array.

First-Time Configuration for SCSI Arrays

The Sun StorEdge 3310 SCSI array and Sun StorEdge 3320 SCSI array are each preconfigured with a single RAID 0 logical drive mapped to LUN 0, and no spare drives. This is not a working configuration. Unmap and delete this logical drive, using the procedure in “To Unmap and Delete a Logical Drive” on page 51, and replace it with logical drives that suit your requirements.

This chapter shows you how to configure your array for the first time, or reconfigure it. It describes the normal sequence of events you follow to configure an array:

- “To View the Logical Drive Configuration” on page 50
- “To View the Channel Configuration” on page 51
- “To Unmap and Delete a Logical Drive” on page 51
- “To Verify the Optimization Mode” on page 53
- “To Change the Optimization Mode” on page 53
- “To Check Physical Drive Availability” on page 54
- “To Configure the Channel Mode” on page 57
- “To Add or Delete a Unique Host ID” on page 58
- “To Change Cylinder and Head Settings” on page 61
- “To Create a Logical Drive” on page 61
- “To Change a Controller Assignment (Optional)” on page 70
- “To Assign a Logical Drive Name (Optional)” on page 71
- “To Partition a Logical Drive (Optional)” on page 72
- “To Map a Logical Drive Partition” on page 75
- “To Label a LUN” on page 78
- “To Create Device Files for Newly Mapped LUNs” on page 79
- “To Save a Configuration to NVRAM” on page 80

Before configuring your array, carefully read chapters 1, 2, and 3.

Note – As you perform the operations described in this and other chapters, you might periodically see event message pop up on the screen. To dismiss an event message after you’ve read it, press Escape. To prevent event messages for displaying so that you can only read them by displaying the event message log, press Ctrl-C. You can press Ctrl-C again at any time to enable pop-up displays of event messages. “Viewing Event Logs on the Screen” on page 287 for more information about event messages.

Existing Logical Drive Configuration

If you are configuring your array for the first time, there is no need to review the existing configuration before you delete it.

If you are reconfiguring logical drives, it is a good idea to view the existing logical drive configuration to determine its status and any changes you want to make to the RAID level, size, number of physical drives that make up a selected logical drive, and spare drives. You also need to view the channel configuration to determine whether you want to make any changes to the channel mode and channel host IDs.

▼ To View the Logical Drive Configuration

1. From the Main Menu, choose “view and edit Logical drives” to display the Logical Drive Status Table.

For a description of this table’s categories, see “Logical Drive Status Table” on page 293.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	29258AC6	NA	RAID5	103428	GOOD				7	B	4	1	0	rtv
1			NONE											
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

2. Note the changes you want to make to the existing configuration.

▼ To View the Channel Configuration

1. From the Main Menu, choose “view and edit channels” to display the Channel Status Table.

For a description of this table’s categories, see “Channel Status Table” on page 298

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Drive	6	7	80.0MHz	Wide	L	Off	80.0MHz	Wide
1	Host	0	NA	80.0MHz	Wide	L	Off	40.0MHz	Wide
2	Drive	6	7	80.0MHz	Wide	L	Off	80.0MHz	Wide
3	Host	NA	1	80.0MHz	Wide	L	Off	Async	Narrow
6<C>	RCCOM	NA	NA	AUTO	Serial	F	NA	1 GHz	Serial

2. Note changes you want to make to the existing configuration.

Deleting Logical Drives

To assign a different RAID level or a different set of drives to a logical drive, or to change local spare drives, you must first unmap and delete the logical drive, and then create a new logical drive.



Caution – This operation erases all data on the logical drive. Therefore, if any data exists on the logical drive, copy it to another location or back it up before it is deleted.

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

▼ To Unmap and Delete a Logical Drive

1. From the Main Menu, choose “view and edit Host luns” to display a list of channel and host IDs.
2. Choose a channel and host ID combination from the list.

A list of channel and host IDs is displayed. You might need to scroll through the list to display some of the channels and host IDs.

3. Select a host LUN and choose Yes to unmap the host LUN from the channel/host ID.

<pre> < Main Menu > view and edit Logical drives view and edit logical Volumes view and edit Host luns u v v v s view system Information view and edit Event logs </pre>		LUN	LU/LD	DRU	Partition	Size<MB>	RAID
	0	LD	0	0	68953	RAID0	
	<pre> Unmap Host Lun ? Yes No </pre>						
	3						
	4						
	5						
	6						
	7						

4. Repeat Step 3 to unmap all remaining host LUNs that are mapped to the logical drive you want to delete.
5. Press Escape to return to the Main Menu.
6. From the Main Menu, choose “view and edit Logical drives.”
7. Select the logical drive that you unmapped and want to delete.
8. Choose “Delete logical drive” and, if it is safe to delete the logical drive, choose Yes to confirm the deletion.

Cache Optimization Mode (SCSI)

Before creating any logical drives, determine the appropriate optimization mode for the array. The type of application accessing the array determines whether to use sequential or random optimization. See “Cache Optimization Mode and Stripe Size Guidelines” on page 30 for a detailed description of sequential and random optimization.

Note – Due to firmware improvements beginning with version 4.11, sequential optimization yields better performance than random optimization for most applications and configurations. Use sequential optimization unless real-world tests in your production environment show better results for random optimization.

If you are modifying an existing configuration and do not want to delete your existing logical drives, verify your optimization mode but do not change it.

▼ To Verify the Optimization Mode

1. From the Main Menu, choose “view and edit Configuration parameters →Caching Parameters.”

Sequential I/O is the default optimization mode.

```
< Main Menu >
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit channels
view and edit Configuration parameters
v
s
v
v
H
D
D
R
C
Communication Parameters
Caching Parameters
Write-Back Cache Enabled
Optimization for Random I/O
Periodic Cache Flush Time - Disabled <Default>
```

2. To accept the optimization mode that is displayed, press Escape.

▼ To Change the Optimization Mode

Once logical drives are created, you cannot use the RAID firmware to change the optimization mode without deleting all logical drives. You can, however, use the Sun StorEdge CLI `set cache-parameters` command to change the optimization mode while logical drives exist. Refer to the *Sun StorEdge 3000 Family CLI 2.0 User's Guide* for more information.

If you have not deleted all logical drives, a notice will inform you of this requirement and you will not be able to change the optimization mode. See “Deleting Logical Drives” on page 51 for the procedure to delete logical drives.

1. From the Main Menu, choose “view and edit Configuration parameters →Caching Parameters” to display the current optimization mode.
2. Select “Optimization for Sequential I/O” or “Optimization for Random I/O” as applicable.

If you have not deleted all logical drives, a notice will inform you of this requirement and you will not be able to change the optimization mode.

3. Choose Yes to change the Optimization mode from Sequential I/O to Random I/O, or from Random I/O to Sequential I/O.

You are prompted to reset the controller:

NOTICE: Change made to this setting will NOT take effect until all Logical Drives are deleted and then the controller is RESET. Prior to resetting the controller, operation may not proceed normally.

Do you want to reset the controller now ?

4. Choose Yes to reset the controller.

If you do not reset the controller now, the optimization mode remains unchanged.

Physical Drive Status

Before configuring physical drives into a logical drive, you must determine the availability of the physical drives in your enclosure. Only drives with a status of FRMT DRV are available.

Note – A drive that does not show a status of FRMT DRV needs to have reserved space added. See “Changing Disk Reserved Space” on page 197 for more information.

▼ To Check Physical Drive Availability

1. From the Main Menu, choose “view and edit Drives” to display a list of all installed physical drives.

Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
0	0	34732	160MB	NONE	FRMT DRV	SEAGATE ST336607LSUN36G
0	1	34732	160MB	NONE	FRMT DRV	SEAGATE ST336607LSUN36G
0	2	34732	160MB	NONE	FRMT DRV	SEAGATE ST336607LSUN36G
0	3	34732	160MB	NONE	FRMT DRV	SEAGATE ST336607LSUN36G
0	4	34732	160MB	NONE	FRMT DRV	SEAGATE ST336607LSUN36G
0	5	34732	160MB	NONE	FRMT DRV	SEAGATE ST336607LSUN36G
0	14				SAF-TE	SUN StorEdge 3310 A
2	0	34732	160MB	NONE	FRMT DRV	SEAGATE ST336607LSUN36G

2. Use the arrow keys to scroll through the table and check that all installed drives are listed.

When the power is initially turned on, the controller scans all installed physical drives that are connected through the drive channels.

Note – If a drive is installed but is not listed, it might be defective or installed incorrectly. If a physical was connected after the controller completed initialization, use the “Scan scsi drive” menu option to enable the controller to recognize the newly added physical drive and to configure it. See “To Scan a New SCSI Drive” on page 179 for information about scanning a new SCSI drive.

3. To view more information about a drive:

- a. Select the drive.
- b. Choose “View drive information.”

Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
0	0	34732	160MB	0	ON-LINE	SEAGATE ST336607LSUN36G
View drive information					N-LINE	SEAGATE ST336607LSUN36G
Scan scsi drive					N-LINE	SEAGATE ST336607LSUN36G
add drive Entry					N-LINE	SEAGATE ST336607LSUN36G
Identify scsi drive					N-LINE	SEAGATE ST336607LSUN36G
clone Failing drive					AND-BY	SEAGATE ST336607LSUN36G
media scan					AND-BY	SEAGATE ST336607LSUN36G
disk Reserved space - 256 mb					AND-BY	SEAGATE ST336607LSUN36G
0	5	34732	160MB	1	ON-LINE	SEAGATE ST336607LSUN36G
0	14				SAF-TE	SUN StorEdge 3310 A
2	0	34732	160MB	2	ON-LINE	SEAGATE ST336607LSUN36G

Additional information is displayed about the drive you selected.

Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
0	0	34732	160MB	0	ON-LINE	SEAGATE ST336607LSUN36G
View drive information					N-LINE	SEAGATE ST336607LSUN36G
S						
a						
I						
c						
n						
Revision Number					0307	ST336607LSUN36G
Serial Number					3JA1BJ2P00007338	ST336607LSUN36G
Disk Capacity (blocks)					71132958	ST336607LSUN36G
disk Reserved space - 256 mb					AND-BY	SEAGATE ST336607LSUN36G
0	5	34732	160MB	1	ON-LINE	SEAGATE ST336607LSUN36G
0	14				SAF-TE	SUN StorEdge 3310 A
2	0	34732	160MB	2	ON-LINE	SEAGATE ST336607LSUN36G

Channel Settings

The Sun StorEdge 3310 SCSI array and Sun StorEdge 3320 SCSI array are preconfigured with the channel settings shown in “Default Channel Configurations” on page 22. Follow the procedures for configuring a channel mode if you plan on adding a host connection or expansion unit. To make changes to channel host IDs, follow the procedures for adding or deleting a host ID.

Configuring Channel Mode

When configuring the channel mode, the following rules apply:

- You must have at least one host channel.
- Channel 0 must remain a dedicated drive channel.
- Channels 1, 2, and 3 can be configured as host or drive channels. In a dual-bus configuration, channel 2 must be a drive channel.
- In a dual-bus configuration, you can only change the mode on channel 3.
- Channel 6 must remain a dedicated RCCOM (redundant controller communications) channel.

Note – RCCOM provides the communication channels by which two controllers in a redundant RAID array communicate with one another. This communication enables the controllers to monitor each other, and includes configuration updates, and control of cache. By default, channel 6 is configured as RCCOM.

▼ To Configure the Channel Mode

1. From the Main Menu, choose “view and edit channels” to display the Channel Status Table.

Main Menu										
view and edit Logical drives										
view and edit logical Volumes										
view and edit Host luns										
view and edit Drives										
view and edit channels										
Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid	
1	Drive	6	7	80.0MHz	Wide	L	Off	80.0MHz	Wide	
2	Host	0	NA	80.0MHz	Wide	L	Off	40.0MHz	Wide	
3	Drive	6	7	80.0MHz	Wide	L	Off	80.0MHz	Wide	
6<	Host	12	1	80.0MHz	Wide	L	Off	Async	Narrow	
6<<	RCCOM	NA	NA	AUTO	Serial	F	NA	1 GHz	Serial	

2. Select the channel that you want to modify, which displays a menu of channel options.

Main Menu										
view and edit Logical drives										
view and edit logical Volumes										
view and edit Host luns										
view and edit Drives										
view and edit channels										
Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid	
1	Drive	6	7	80.0MHz	Wide	L	Off	80.0MHz	Wide	
1	channel Mode				Wide	L	Off	40.0MHz	Wide	
2	Primary controller scsi id				Wide	L	Off	80.0MHz	Wide	
3	Secondary controller scsi id				Wide	L	Off	Async	Narrow	
6<	scsi Terminator				erial	F	NA	1 GHz	Serial	
6<	sync transfer Clock									
6<	Wide transfer									
6<	parity check - Enabled									
6<	view chip inFormation									

3. Choose “channel Mode” to change the channel from host to drive, or drive to host, and then choose Yes to confirm the mode change.

This change does not take effect until the controller is reset.

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?

4. Choose Yes to reset the controller.

Host Channel IDs

Host channel IDs identify the controller to the host. Some applications require that specific IDs be assigned to host channels in order to recognize the array. Sun StorEdge 3310 SCSI array and Sun StorEdge 3320 SCSI array default host channel IDs are shown in TABLE 3-1 under “Default Channel Configurations” on page 22.

Each host ID can have up to 32 partitions, which are then mapped to LUNs to create a total not to exceed 128. The default host channel ID settings enable you to map up to a total of 64 LUNs. To map up to 128 LUNs, you must add host IDs. At least four host IDs are required; no more than six host IDs are supported.

For details on mapping 128 LUNs, refer to “Planning for 128 LUNs on a SCSI Array (Optional)” on page 162.

Each host channel has a unique primary and secondary ID available. You can:

- Edit each host ID to change the ID of each controller host channel that is seen by the host. To change an ID, you must delete it first and then add the new ID.
- Add host IDs (by adding a second host ID to channels 1 and 3, and additional host IDs if you make Channel 2 into a host channel).

Note – Channel ID values range from 0 to 15.

▼ To Add or Delete a Unique Host ID

Note – To change an ID, you must delete the old ID first and then add the new ID

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

< Main Menu >										
view and edit Logical drives										
view and edit logical Volumes										
view and edit Host luns										
view and edit Drives										
view and edit channels										
U V S V U	Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
	2	Drive	6	7	80.0MHz	Wide	L	Off	80.0MHz	Wide
	1	Host	0	NA	80.0MHz	Wide	L	Off	40.0MHz	Wide
	2	Drive	6	7	80.0MHz	Wide	L	Off	80.0MHz	Wide
	3	Host	12	1	80.0MHz	Wide	L	Off	Async	Narrow
	6<C>	RCCOM	NA	NA	AUTO	Serial	F	NA	1 GHz	Serial

2. Select the host channel on which you want to add an ID.

3. Choose “view and edit scsi Id.”

If host IDs have already been configured on the host channel, they will be displayed. If no host IDs have been configured, the following message is displayed.

No SCSI ID Assignment - Add Channel SCSI ID?

- 4. If a host ID has already been assigned to that channel, select an ID and press Return to view a menu for adding or deleting SCSI IDs.**
- 5. To add an ID, select “Add Channel SCSI ID.” To delete an ID, select “Delete Channel SCSI ID.”**
- 6. If adding an ID, select a controller from the list to display a list of SCSI IDs. If deleting an ID, select Yes to delete the ID.**
- 7. If adding an ID, select an ID from the list, and then choose Yes to confirm the addition.**
- 8. If you are only changing one Channel ID, choose Yes to the following confirmation message to reset the controller.**

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?

9. If you are changing more than one Channel ID, do not reset the controller until all IDs are changed.

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

Creating Logical Drives

The RAID array is preconfigured with one RAID 0 logical drive as described in “Default Logical Drive Configuration” on page 22. Each logical drive consists of a single partition by default.

This section describes how to modify the RAID level or add logical drives. In these procedures, you configure a logical drive to contain one or more physical drives based on the desired RAID level, and divide the logical drive into additional partitions.

Note – Depending on the size and RAID level, it can take several hours to build a logical drive. Online initialization, however, enables you to begin configuring and using the logical drive before initialization is complete.

If you do not use on-line initialization, be sure to allow enough time when you create logical drives. Creating a 2-Tbyte RAID 5 logical drive can take up to:

- 2.25 hours for Sun StorEdge 3310 SCSI arrays and Sun StorEdge 3510 FC arrays
- 10.3 hours for a Sun StorEdge 3511 SATA array

Preparing for Logical Drives Larger Than 253 Gbyte

The Solaris operating system requires drive geometry for various operations, including `newfs`. For the appropriate drive geometry to be presented to the Solaris operating system for logical drives larger than 253 Gbyte, use the default settings shown below to cover all logical drives over 253 Gbyte. These settings work for smaller configurations as well. The controller automatically adjusts the sector count so the operating system can read the correct drive capacity.

For Solaris operating system configurations, use the values in the following table.

TABLE 4-1 Cylinder and Head Mapping for the Solaris Operating System

Logical Drive Capacity	Cylinder	Head	Sector
< 253 GB	< 65536 (default)	variable	variable (default)
253 GB–1 TB	< 65536 (default)	64 (default)	variable (default)

See “Host Cylinder/Head/Sector Mapping Configuration” on page 236 for more information. See “To Change Cylinder and Head Settings” on page 98 for instructions on how to apply these settings to FC and SATA arrays.

After settings are changed, they apply to all logical drives in the chassis.

Note – Refer to your operating system documentation for limitations on device sizes.

▼ To Change Cylinder and Head Settings

1. Choose “view and edit Configuration parameters →Host-side Parameters →Host Cylinder/Head/Sector Mapping Configuration →Sector Ranges - →Variable,” and then choose Yes to confirm your choice.
2. Choose “Head Ranges - →64 Heads,” and then choose Yes to confirm your choice.
3. Choose “Cylinder Ranges - →< 65536,” and then choose Yes to confirm your choice.

▼ To Create a Logical Drive

Note – To reassign drives and add local or global spare drives to the preconfigured array, you must first unmap and then delete the existing logical drives. For more information about deleting a logical drive, see “Deleting Logical Drives” on page 51.

1. From the Main Menu, choose “view and edit Logical drives.”
Unassigned logical drives show a RAID level of NONE.
2. Select the first available unassigned logical drive (LG).

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	691A4994	NA	RAID5	68952	GOOD	I		7	T		3	0	0	
1			NONE											
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

You can create as many as 16 logical drives using physical drives on any loop.

3. When prompted to Create Logical Drive? choose Yes to confirm your choice and display a pull-down list of supported RAID levels.
4. Select a RAID level from the list to assign to the logical drive.

Note – RAID 5 is used as an example in the following steps.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	RAID 5	#FL	NAME
P0	691A4994	NA	RAID5	68952	GOOD	I			7	RAID 3	0	
1			NONE							RAID 1		
2			NONE							RAID 0		
3			NONE							NRAID		
4			NONE									
5			NONE									
6			NONE									
7			NONE									

Note – NRAID does not provide data redundancy. The NRAID option that appears in some firmware menus does not provide the protection of other RAID levels and is rarely used.

For more information about RAID levels, see “RAID Levels” on page 316.

5. Select the drives you want to include in the logical drive from the list of available physical drives, using the steps below.

You must select at least the minimum number of drives required for the selected RAID level.

- RAID 3 and RAID 5 require a minimum of three physical drives.
- RAID 0 requires a minimum of two physical drives.
- RAID 1 requires a minimum of two physical drives. Additional drives must be added in increments of two.

For redundancy, you can create a logical drive containing drives distributed over separate channels. You can then create several partitions on each logical drive. In a RAID 1 or RAID 0+1 configuration, the order in which you select the physical drives for a logical drive determines the channels to which the physical drives are assigned. If you want drives to be mirrored over two channels, select them in the appropriate order. For example:

- The first drive you select is assigned to channel 2, ID0.
- The second drive you select is assigned to channel 0, ID0.
- The third drive you select is assigned to channel 2, ID1.
- The fourth drive you select is assigned to channel 0, ID1.

- a. Use the up and down arrow keys and press Return to select the drives you want to include in the logical drive.

An asterisk mark (*) is displayed in the Chl (Channel) column of each selected physical drive.

Chl	ID	Size<MB>	Speed	LG_DRU	Status	Vendor and Product ID
* 0	3	34732	160MB	NONE	FRMT DRU	SEAGATE ST336607LSUN36G
* 0	4	34732	160MB	NONE	FRMT DRU	SEAGATE ST336607LSUN36G
* 0	5	34732	160MB	NONE	FRMT DRU	SEAGATE ST336607LSUN36G
2	0	34732	160MB	NONE	FRMT DRU	SEAGATE ST336607LSUN36G
2	3	34732	160MB	NONE	FRMT DRU	SEAGATE ST336607LSUN36G
2	4	34732	160MB	NONE	FRMT DRU	SEAGATE ST336607LSUN36G
2	5	34732	160MB	NONE	FRMT DRU	SEAGATE ST336607LSUN36G

- b. To deselect a drive, press Return again on the selected drive.

The asterisk marking that drive disappears.

- c. After all physical drives have been selected for the logical drive, press Escape to display a menu of logical drive options.

Several optional menu options are displayed. You can choose these menu options to define aspects of the logical drive you are creating:

- The “Maximum Drive Capacity” menu option enables you to specify the size of the logical drive.
- The “Assign Spare Drives” menu option enables you to specify a local spare drive to be used if an existing physical drive in the logical drive you are creating becomes defective.
- The “Disk Reserved Space” menu option displays the size of the reserved space used to store logical drive metadata. While it is possible to delete or change reserved space size, do not do so. Refer to “To Specify Disk Reserved Space” on page 198 for more information.
- The “Logical Drive Assignments” menu option enables you to assign the logical drive you are creating to either the primary or secondary controller.
- The “Write Policy:” menu option enables you to set the cache write policy for the logical drive you are creating.
- The “Initialize Mode:” menu option enables you specify whether the logical drive you are creating is initialized on-line or off-line.
- The “Stripe Size:” menu option enables you to specify the stripe size for the logical drive you are creating.

These menu options are described in the remainder of this section.

6. (Optional) Set the maximum logical drive capacity, using the following procedure:

a. Choose "Maximum Drive Capacity."

Note – Changing the maximum drive capacity reduces the size of the logical drive and leaves some disk space unused.

b. Type in the maximum capacity of each physical drive that makes up the logical drive you are creating.

LG	ID	LU	RAID	Size<MB>	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	161637C1	NA	RAID5	103428	GOOD				7	B	4	0	0	
1			NONE											
	<div style="border: 1px solid black; padding: 5px;"> <p>Maximum Drive Capacity : 34476MB</p> <p>Maximum Available Drive Capacity<MB>: 34476</p> <p>Maximum Drive Capacity<MB> : 1024_</p> </div>													
6			NONE											
7			NONE											

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.

7. (Optional) Add a local spare drive from the list of unused physical drives by following these steps:

a. Choose "Assign Spare Drives" to display a list of available physical drives you can use as a local spare.

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

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

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

- b. Select a physical drive from the list to use as a local spare.

Slot	Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
*	2<3>	8	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	2<3>	10	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G

- c. Press Escape to return to the menu of logical drive options.

Note – The Disk Reserved Space option is not supported while you are creating logical drives.

If you use two controllers for a redundant configuration, you can assign a logical drive to either of the controllers to balance the workload. By default, all logical drives are assigned to the primary controller.

Logical drive assignments can be changed later, but that operation requires a controller reset to take effect.

8. (Optional) For dual-controller configurations, you can assign this logical drive to the secondary controller by following these steps:



Caution – In single-controller configurations, assign logical drives only to the primary controller.

- a. Choose “Logical Drive Assignments.”

A confirmation message is displayed.

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

- b. Choose Yes to assign the logical drive to the redundant controller.

9. (Optional) Configure the logical drive’s write policy.

Write-back cache is the preconfigured global logical drive write policy, which is specified on the Caching Parameters submenu. (See “Enabling and Disabling Write-Back Cache” on page 231 for the procedure on setting the global caching parameter.) This option enables you to assign a write policy per logical drive that is either the same as or different than the global setting. Write policy is discussed in more detail in “Cache Write Policy Guidelines” on page 32.

- a. Choose “Write Policy -.”

Note – The Default write policy displayed is the global write policy assigned to all logical drives.

The following write policy options are displayed:

- **Default (default).** Assigns the global write policy. If the global setting for write policy is changed, automatically changes the write policy for this logical drive.
As described in “Cache Write Policy Guidelines” on page 32, the array can be configured to dynamically switch write policy from write-back cache to write-through cache if specified events occur. Write policy is only automatically switched for logical drives with write policy configured to Default. See “Event Trigger Operations” on page 267 for more information.
- **Write-Back.** Assigns write-back cache regardless of any changes to the global write policy.
- **Write-Through.** Assigns write-through cache regardless of any changes to the global write policy.

b. Choose a write policy option.



Note – You can change the logical drive logical drive’s write policy at any time, as explained in “Changing Write Policy for a Logical Drive” on page 148.

10. (Optional) Set the logical drive initialization mode by choosing “Initialize Mode” from the list of logical drive options, and then choosing Yes to change the initialization mode.

The assigned initialization mode is displayed in the list of logical drive options.

You can choose between these two logical drive initialization options:

- On-Line (default)

This option enables you to configure and use the logical drive before initialization is complete. Because the controller is building the logical drive while performing I/O operations, initializing a logical drive on-line requires more time than off-line initialization.

- **Off-Line**

This menu option enables you to configure and use the drive only after initialization is complete. Because the controller is building the logical drive without having to also perform I/O operations, off-line initialization requires less time than on-line initialization.

Because logical drive initialization can take a considerable amount of time, depending on the size of your physical disks and logical drives, you can choose on-line initialization so that you can use the logical drive before initialization is complete.

11. (Optional) Configure the logical drive stripe size.

Depending on the optimization mode selected, the array is configured with the default stripe sizes shown in “Cache Optimization Mode and Stripe Size Guidelines” on page 30. When you create a logical drive, however, you can assign a different stripe size to that logical drive.

Note – Default stripe sizes result in optimal performance for most applications. Selecting a stripe size that is inappropriate for your optimization mode and RAID level can decrease performance significantly. For example, smaller stripe sizes are ideal for I/Os that are transaction-based and randomly accessed. But when a logical drive configured with a 4-Kbyte stripe size receives files of 128 Kbyte, each physical drive has to write many more times to store it in 4-Kbyte data fragments. Change stripe size only when you are sure it will result in performance improvements for your particular applications.

See “Cache Optimization Mode and Stripe Size Guidelines” on page 30 for more information.

Note – Once a logical drive is created, its stripe size cannot be changed. To change the stripe size, you must delete the logical drive, and then recreate it using the new stripe size.

a. Choose Stripe Size.

A menu of stripe size options is displayed.

- If you specified random optimization, in addition to Default, you can choose among disk stripe sizes of 4 Kbyte, 8 Kbyte, 16 Kbyte, 32 Kbyte, 64 Kbyte, 128 Kbyte, and 256 Kbyte.
- If you specified sequential optimization, in addition to Default, you can choose among disk stripe sizes of 16 Kbyte, 32 Kbyte, 64 Kbyte, 128 Kbyte, and 256 Kbyte.

- b. Choose Default to assign the stripe size per Optimization mode, or choose a different stripe size from the menu.

Default stripe size per optimization mode is shown in “Cache Optimization Mode and Stripe Size Guidelines” on page 30.

The selected stripe size is displayed in the list of logical drive options.

12. Once all logical drive options have been assigned, press Escape to display the settings you have chosen.

```
Raid Level           : RAID 5
Online SCSI Drives  : 3
Maximum Drive Capacity : 20000 MB
Disk Reserved Space  : 256 MB
Spare SCSI Drives    : 0
Logical Drive Assignment: Secondary Controller
Write Policy         : Default<Write-Back>
Initialize Mode      : On-Line
Stripe Size          : 128K Bytes

Create Logical Drive ?
  Yes                No
```

13. Verify that all the information is correct, and then choose Yes to create the logical drive.

Note – If the logical drive has not been configured correctly, select No to return to the logical drive status table so you can configure the drive correctly.

Messages indicate that the logical drive initialization has begun, and then that it has completed.

14. Press Escape to close the drive initialization message.

A progress bar displays the progress of initialization as it occurs.

You can press Escape to remove the initialization progress bar and continue working with menu options to create additional logical drives. The percentage of completion for each initialization in progress is displayed in the upper left corner of the window as shown in the following example screen.

12:0%														
BAT:++++														
LG	ID	LU	RAID	Size<MB>	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	13843684	NA	RAID5	103428	GOOD				7	B	4	1	0	
1					On-Line Initializing Notification									
2				LG:2 Logical Drive NOTICE: Starting Creation										
3														
4			NONE											
5			NONE											
6			NONE											
7			NONE											

The following message is displayed when the initialization is completed:

LG	ID	LU	RAID	Size<MB>	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	691A4994	NA	RAID5	68952	GOOD				7	T	3	0	0	
S1	518F95F				Notification									
2				On-Line Initialization of Logical Drive 1 Completed										
3														
4			NONE											
5			NONE											
6			NONE											
7			NONE											

15. Press Escape to dismiss the notification.

The newly created logical drive is displayed in the status window.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	691A4994	NA	RAID5	68952	GOOD				7	T	3	0	0	
S1	518F95F1	NA	RAID5	68952	GOOD				7	B	3	0	0	
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

Controller Assignment

By default, logical drives are automatically assigned to the primary controller. If you assign half of the logical drives to the secondary controller in a dual controller array, 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).



Caution – In single-controller configurations, 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. In a single-controller configuration, if you disable the Redundant Controller function and reconfigure the controller with the Autoconfigure option or as a secondary controller, the controller module becomes inoperable and will need to be replaced.

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 a Partition to a Host LUN” on page 74).

▼ To Change a Controller Assignment (Optional)



Caution – Only assign logical drives to primary controllers in single-controller configurations.

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the drive you want to reassign.
3. Choose “logical drive Assignments,” and then choose Yes to confirm the reassignment.

The reassignment is evident from the “view and edit Logical drives” screen. A “P” in front of the LG number, such as “P0,” 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 the secondary controller.

Logical Drive Name

You can assign a name to each logical drive. These logical drive names are used only in RAID firmware administration and monitoring and do not appear anywhere on the host. After you assign a drive name, you can change it at any time.

▼ To Assign a Logical Drive Name (Optional)

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select a logical drive.
3. Choose “logical drive Name.”
4. Type the name you want to give the logical drive in the New Logical Drive Name field and press Return to save the name.

LG	ID	LU	RAID	Size<MB>	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	691A4994	NA	RAID5	68952	GOOD				7	T	3	0	0	
S1	518F95F1	NA	RAID5	68952	GOOD	I			7	B	3	0	0	
<div style="border: 1px solid black; padding: 5px;"> <p>View scsi drives Delete logical drive Partition logical drive logical drive Name</p> </div>														
<div style="border: 1px solid black; padding: 5px;"> <p>I r s w</p> <p>Current Logical Drive Name: New Logical Drive Name: New Name_</p> </div>														
7			NONE											

Partitions

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 and up to 128 LUN assignments. For guidelines on setting up 128 LUNs, see “Planning for 128 LUNs on a SCSI Array (Optional)” on page 162.



Caution – If you modify the size of a partition or logical drive, all data on the drive is lost.

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

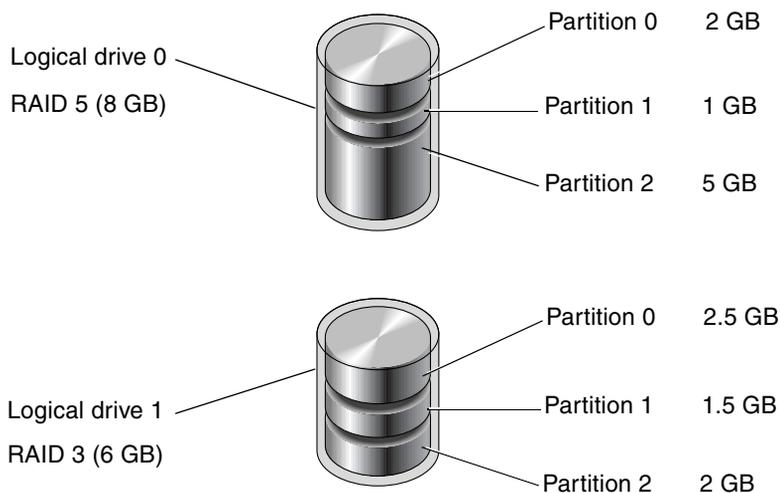


FIGURE 4-1 Partitions in Logical Drives

▼ To Partition a Logical Drive (Optional)



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

1. From the Main Menu, choose "view and edit Logical drives."
2. Select the logical drive you want to partition.
3. Choose "Partition logical drive."

If the logical drive has not already been partitioned, the following warning is displayed:

```
This operation may result in the LOSS OF ALL DATA on the Logical
Disk.

Partition Logical Drive?
```

4. Choose Yes to confirm.

A list of the partitions on 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 a partition.

A Partition Size dialog is displayed.

6. Type the desired size of the selected partition.

The following warning is displayed:

```
This operation will result in the LOSS OF ALL DATA on the partition.
Partition Logical Drive?
```

7. Choose Yes to confirm.

The remaining capacity of the logical drive is automatically allocated to the next partition. In the following example, a partition size of 20000 Mbyte was entered; the remaining storage of 20000 Mbyte is allocated to the partition below the newly created partition.

Partition	Offset(MB)	Size(MB)
0	0	20000
1	20000	20000
2		
3		
4		
5		
6		
7		

8. Repeat Step 5 through Step 7 to partition the remaining capacity of your logical drive.

For information on deleting a partition, see “Deleting a Logical Drive Partition” on page 130.

Mapping a Partition to a Host LUN

A partition is a division of the logical drive that appears as a physical drive to any host that has access to that partition. For Sun StorEdge 3310 SCSI arrays and Sun StorEdge 3320 SCSI arrays, you can create a maximum of 32 partitions per logical drive. So that host bus adapters (HBAs) recognize the partitions when the host bus is reinitialized, each partition must be mapped to a host LUN (logical unit number).

Channel IDs represent the physical connection between the HBA and the array. The host ID is an identifier assigned to the channel so that the host can identify LUNs. The following figure shows the relationship between a host ID and a LUN.

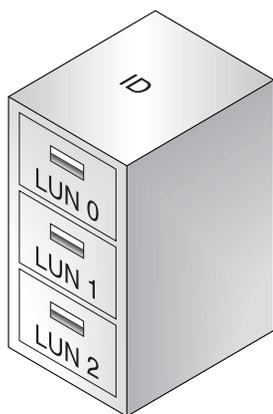


FIGURE 4-2 LUNs Resemble Drawers in a File Cabinet Identified by an ID

The ID is like a cabinet, and the drawers are the LUNs.

- Each cabinet (ID) can have up to 32 drawers (LUNs).
- The maximum number of LUNs that can be created for a Sun StorEdge 3310 SCSI array or Sun StorEdge 3320 SCSI array is 128. To create 128 LUNs for a SCSI array, see “Planning for 128 LUNs on a SCSI Array (Optional)” on page 162.

The following figure illustrates mapping partitions to host ID/LUNs.

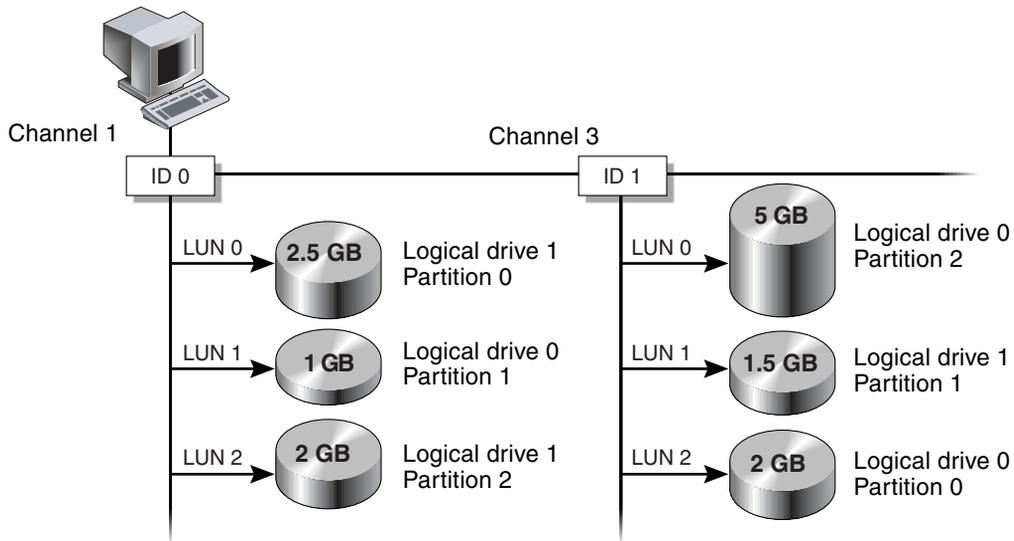


FIGURE 4-3 Mapping Partitions to Host ID/LUNs

All hosts on the mapped host channel have full access to all partitions mapped to LUNs on that channel. To provide redundant connections between a host and a partition, map the partition to a LUN on both of the host channels that connect to that host. Only one partition can be mapped to each LUN.

Note – When you modify a partition, you must first unmap the LUN.

Note – If you plan to map 128 LUNs, the process is easier if you use Sun StorEdge Configuration Service. Refer to the *Sun StorEdge 3000 Family Configuration Service User's Guide* for more information.

▼ To Map a Logical Drive Partition

1. From the Main Menu, choose “view and edit Host luns.”

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

2. Select a channel and host ID on the primary controller.

- If the Logical Drive and Logical Volume menu options are displayed, choose Logical Drive to display the LUN table.

< Main Menu > view and edit Logical drives view and edit logical Volumes view and edit Host luns		LUN	LU/LD	DRU	Partition	Size(MB)	RAID
v v CHL 1 ID 0 <Primary Controll v CHL 3 ID 1 <Secondary Contro v CHL 3 ID 12 <Primary Control s view system Information view and edit Event logs		0	LD	0	0	68953	RAID0
		1					
		2					
		3					
		4					
		5					
		6					
		7					

- Select the LUN you want to map the drive to.

A list of available logical drives is displayed.

- Select the logical drive (LD) that you want to map to the selected LUN.

The partition table is displayed.

< Main Menu > view and edit Logical drives view and edit logical Volumes view and edit Host luns		LUN	LU/LD	DRU	Partition	Size(MB)	RAID
v v CHL 1 ID 0 <Primary Controll		0	LD	0	0	68953	RAID0
		1					

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	456499FB	NA	RAID0	68953	GOOD				5	T	2	-	0	
P2	48B9EEC6	NA	RAID0	68953	GOOD				5	T	2	-	0	
5														
6														
7														

- Select the partition you want to map to the selected LUN.

- Choose "Map Host LUN," and then choose Yes to confirm the host LUN mapping.

< Main Menu >		LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit Logical drives		0	LD	0	0	68953	RAID0
view and edit logical Volumes							
view and edit Host luns							
v	CHL 1 ID 0 <Primary Controll	1					
v	CHL 3 ID 1 <Secondary Contro						
v	CHL 3 ID 12 <Primary Control						
s	Map Host LUN						
view system Information		3					
view and edit Event logs		4					
		5					
		6					
		7					

The partition is now mapped to the selected LUN.

< Main Menu >		LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit Logical drives		0	LD	0	0	68953	RAID0
view and edit logical Volumes							
view and edit Host luns							
v	CHL 1 ID 0 <Primary Controll	1	LD	0	0	68953	RAID0
v	CHL 3 ID 1 <Secondary Contro	2					
v	CHL 3 ID 12 <Primary Control						
s		3					
view system Information		4					
view and edit Event logs		5					
		6					
		7					

- Repeat Step 4 through Step 7 to map additional partitions to host LUNs on this channel and logical drive.
- Press Escape.
- If you are LUN mapping a redundant configuration, repeat Step 2 through Step 7 to map partitions to host LUNs with other IDs on the logical drive assigned to the primary controller.

When you map a partition to two channels in a redundant configuration, the number in the Partition column of the partition table displays an asterisk (*) to indicate that the partition is mapped to two LUNs.

Note – If you are using host-based multipathing software, map each partition to two or more host IDs so multiple paths will be available from the partition to the host.

- Repeat Step 2 through Step 10 to map hosts to the secondary controller.

12. To verify unique mapping of each LUN (unique LUN number, unique DRV number, or unique Partition number):
 - a. From the Main Menu, choose “view and edit Host luns.”
 - b. Select the appropriate controller and ID and press Return to review the LUN information.

A mapped LUN displays a number in the host LUN partition window.
13. When all host LUNs have been mapped, save the updated configuration to nonvolatile memory. See “Saving Configuration (NVRAM) to a Disk” on page 80 for more information.
14. (Solaris operating system only) For the Solaris operating system to recognize a LUN, you must first manually write the label using the `Auto configure` option of the `format (1M)` utility, as described in “To Label a LUN” on page 78.

Labeling a LUN (Solaris Operating System Only)

For the Solaris operating system to recognize a LUN, you must first manually write the label using the `Auto configure` option of the `format` command.

For additional operating system information, refer to the Installation, Operation, and Service Manual for your Sun StorEdge 3000 family array.

▼ To Label a LUN

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

```
# format
```

2. Specify the disk number when prompted.
 1. Type `Y` at the following prompt, if it is deployed, and press Return:

```
Disk not labeled. Label it now? Y
```

The Solaris operating system’s Format menu is displayed.

2. Type `type` to select a drive type.
3. Type `0` to choose the `Auto configure` menu option.
Choose the `Auto configure` menu option regardless of which drive types are displayed by the `type` option.
4. Type `label` and press `Y` when prompted to continue.

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

5. Type `quit` to finish using the `Format` menu.

Solaris Operating System Device Files

Perform the following procedure to create device files for newly mapped LUNs on hosts in the Solaris 8 and Solaris 9 operating system.

For additional operating system information, see the *Installation, Operation, and Service* manual for your Sun StorEdge 3000 family array.

▼ To Create Device Files for Newly Mapped LUNs

1. To create device files, type:

```
# /usr/sbin/devfsadm -v
```

2. To display the new LUNs, type:

```
# format
```

3. If the `format` command does not recognize the newly mapped LUNs, perform a configuration reboot:

```
# reboot -- -r
```

Saving Configuration (NVRAM) to a Disk

The controller configuration information is stored in non-volatile RAM (NVRAM). When you save it, the information is stored in the disk reserved space of all drives that have been configured into logical drives. Back up the controller configuration information whenever you change the array's configuration.

Saving NVRAM controller configuration to a file provides a backup of controller configuration information such as channel settings, host IDs, 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.

Note – A logical drive must exist for the controller to write NVRAM content onto it.

▼ To Save a Configuration to NVRAM

- **Choose “system Functions →Controller maintenance →Save nvram to disks,” and choose Yes to save the contents of NVRAM to disk.**

A prompt confirms that the NVRAM information has been successfully saved.

To restore the configuration, see “Restoring Your Configuration (NVRAM) From Disk” on page 284.

If you prefer to save and restore all configuration data, including LUN mapping information, use Sun StorEdge Configuration Service or the Sun StorEdge CLI 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 or to the *Sun StorEdge 3000 Family CLI User's Guide* for information about the `reset nvram` and `download controller-configuration` commands.

First-Time Configuration for FC or SATA Arrays

The Sun StorEdge 3510 FC array and Sun StorEdge 3511 SATA array are each preconfigured with a single RAID 0 logical drive mapped to LUN 0, and no spare drives. This is not a working configuration. Unmap and delete this logical drive, using the procedure in “To Unmap and Delete a Logical Drive” on page 83, and replace it with logical drives that suit your requirements.

This chapter shows you how to configure your array for the first time, or reconfigure it. It describes the normal sequence of events you follow to configure an array:

- “To View the Logical Drive Configuration” on page 82
- “To View the Channel Configuration” on page 83
- “To Unmap and Delete a Logical Drive” on page 83
- “To Verify the Optimization Mode” on page 85
- “To Change the Optimization Mode” on page 85
- “To Check Physical Drive Availability” on page 86
- “To Enable or Disable Mixed Drive Support” on page 88
- “To Modify a Channel Mode” on page 89
- “To Add or Delete a Unique Host ID” on page 93
- “To Assign an ID From a Different Range” on page 95
- “To Change the Fibre Connection Protocol” on page 97
- “To Change Cylinder and Head Settings” on page 98
- “To Create a Logical Drive” on page 99
- “To Change a Controller Assignment (Optional)” on page 108
- “To Assign a Logical Drive Name (Optional)” on page 109
- “To Partition a Logical Drive (Optional)” on page 110
- “To Map a Logical Drive Partition” on page 116
- “To Create Host Filters (FC and SATA Arrays Only)” on page 118
- “To Label a LUN” on page 123
- “To Create Device Files for Newly Mapped LUNs” on page 124
- “To Save a Configuration to NVRAM” on page 125

Before configuring your array, carefully read chapters 1, 2, and 3.

Note – As you perform the operations described in this and other chapters, you might periodically see event message pop up on the screen. To dismiss an event message after you’ve read it, press Escape. To prevent event messages for displaying so that you can only read them by displaying the event message log, press Ctrl-C. You can press Ctrl-C again at any time to enable pop-up displays of event messages. “Viewing Event Logs on the Screen” on page 287 for more information about event messages.

Existing Logical Drive Configuration

If you are configuring your array for the first time, there is no need to review the existing configuration before you delete it.

If you are reconfiguring logical drives, note the existing logical drive configuration to determine its status and any changes you want to make to the RAID level, logical drive size, number of physical drives that make up a selected logical drive, and spare drives. Also view the channel configuration to determine whether you want to make any changes to the channel mode and channel host IDs.

▼ To View the Logical Drive Configuration

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

The Logical Drive Status table is displayed. For a description of the parameters, see “Logical Drive Status Table” on page 293.

LG	ID	LU	RAID	Size<MB>	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	161637G1	NA	RAID5	103428	GOOD			7	B	4	0	0		
1			NONE											
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

2. Note the changes you want to make to the existing configuration.

▼ To View the Channel Configuration

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

The Channel Status table is displayed.

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Serial
1	Host	NA	42	AUTO	Serial	F	NA		
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	2 GHz	Serial
5	Host	NA	46	AUTO	Serial	F	NA		

2. Note the changes you want to make to the existing configuration.

Deleting Logical Drives

To assign a different RAID level or set of drives to a logical drive, or to change local spare drives, you must first unmap and delete the logical drive and then create a new logical drive.



Caution – This operation erases all data on the logical drive. If any data exists on the logical drive, copy it to another location or back it up before it is deleted.

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

▼ To Unmap and Delete a Logical Drive

1. From the Main Menu, choose “view and edit Host luns.”

A list of channel and host IDs is displayed. You might need to scroll through the list to display some of the channels and host IDs.

2. Choose a channel and host ID combination from the list.

A list of host LUNs that are assigned to the selected channel/host combination is displayed.

3. Select a host LUN and choose Yes to unmap the host LUN from the channel/host ID.

< Main Menu >		LUN	LU/LD	DRU	Partition	Size<MB>	RAID
view and edit Logical drives		0	LD	0	0	68850	RAID5
view and edit logical Volumes							
view and edit Host luns							
v CHL 0 ID 40 <Primary Control		Unmap Host Lun ?					
v CHL 1 ID 42 <Secondary Contr		<input checked="" type="checkbox"/> Yes No					
v Edit Host-ID/WWN Name List							
s view system Information		3					
view and edit Event logs		4					
		5					
		6					
		7					

4. Repeat Step 3 to unmap all remaining host LUNs that are mapped to the logical drive you want to delete.
5. Press Escape to return to the Main Menu.
6. From the Main Menu, choose “view and edit Logical drives.”
7. Select the logical drive that you unmapped and want to delete.
8. Choose “Delete logical drive,” and, if it is safe to delete the logical drive, choose Yes to confirm the deletion.

Cache Optimization Mode (FC and SATA)

Before creating any logical drives, determine the appropriate optimization mode for the array. The type of application accessing the array determines whether to use sequential or random optimization. See “Cache Optimization Mode and Stripe Size Guidelines” on page 30 for a detailed description of sequential and random optimization.

Note – Due to firmware improvements beginning with version 4.11, sequential optimization yields better performance than random optimization for most applications and configurations. Use sequential optimization unless real-world tests in your production environment show better results for random optimization.

If you are modifying an existing configuration and do not want to delete your existing logical drives, verify your optimization mode but do not change it.

▼ To Verify the Optimization Mode

1. **From the Main Menu, choose “view and edit Configuration parameters →Caching Parameters.”**

Sequential I/O is the default optimization mode.

2. **To accept the optimization mode that is displayed, press Escape.**

To change the optimization mode, see “To Change the Optimization Mode” on page 85.

▼ To Change the Optimization Mode

Once logical drives are created, you cannot use the RAID firmware to change the optimization mode without deleting all logical drives. You can, however, use the Sun StorEdge CLI `set cache-parameters` command to change the optimization mode while logical drives exist. Refer to the *Sun StorEdge 3000 Family CLI User's Guide* for more information.

If you have not deleted all logical drives, a notice will inform you of this requirement and you will not be able to change the optimization mode. See “Deleting Logical Drives” on page 83 for the procedure to delete logical drives.

1. **From the Main Menu, choose “view and edit Configuration parameters →Caching Parameters.”**

The Optimization mode that is currently assigned to the array is displayed.

2. **Choose “Optimization for Sequential I/O” or “Optimization for Random I/O” as appropriate.**

If you have not deleted all logical drives, a notice will inform you of this requirement and you will not be able to change the optimization mode.

3. **Choose Yes to change the Optimization mode from Sequential I/O to Random I/O, or from Random I/O to Sequential I/O.**

You are prompted to reset the controller:

```
NOTICE: Change made to this setting will NOT take effect until all
Logical Drives are deleted and then the controller is RESET. Prior
to resetting the controller, operation may not proceed normally.
```

```
Do you want to reset the controller now ?
```

4. Choose Yes to reset the controller.

If you do not reset the controller now, the optimization mode remains unchanged.

Physical Drive Status

Before configuring physical drives into a logical drive, you must determine the availability of the physical drives in your enclosure. Only drives with a status of FRMT DRV are available.

Note – A drive that does not show a status of FRMT DRV needs to have reserved space added. See “Changing Disk Reserved Space” on page 197 for more information.

▼ To Check Physical Drive Availability

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

A list of all the installed physical drives is displayed.

Ch1	ID	Size<MB>	Speed	LG_DRU	Status	Vendor and Product ID
2<3>	6	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	7	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	8	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	9	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	10	34732	200MB	GLOBAL	STAND-BY	SEAGATE ST336753FSUN36G
2<3>	12				SES	SUN StorEdge 3510F A

2. Use the arrow keys to scroll through the table and check that all installed drives are listed.

Note – If a drive is installed but is not listed, it might be defective or installed incorrectly.

When the power is initially turned on, the controller scans all physical drives that are connected through the drive channels.

To view more information about a drive:

- a. Select the drive.
- b. Choose “View drive information.”

Additional information is displayed about the drive you selected.

	Ch1	ID	Size<MB>	Speed	LG_DRU	Status	Vendor and Product ID
view	2<3>	3	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G
view	2	View drive information				N-LINE	SEAGATE ST336753FC
view	2	I c P m d	Revision Number		0205		3FSUN36G
view	2		Serial Number		3ET0N0U000007303		3FSUN36G
syst	2		Disk Capacity (blocks)		71132958		3FSUN36G
view	2		Node Name(WWNN)		20 00 00 04 CF AB 13 8F		3FSUN36G
view	2		Redundant Loop ID		3		3FSUN36G
	2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	2<3>	10	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	2<3>	12				SES	SUN StorEdge 3510F A

Enabling Support for SATA Expansion Units Attached to FC Arrays

It is possible to connect Sun StorEdge 3511 SATA expansion units to Sun StorEdge 3510 FC arrays, either alone or in combination with Sun StorEdge 3510 FC expansion units. Refer to the release notes and *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array for important information about limitations and appropriate uses of such a configuration.

If you do connect one or more Sun StorEdge 3511 SATA expansion units to a Sun StorEdge 3510 FC array, you must ensure that mixed drive support is enabled. Enabling mixed drive support means certain safeguard menu options and messages will ensure that you do not improperly mix FC and SATA drive types when performing such operations as creating logical drives and logical volumes, or assigning local or global spares to logical drives.

If you have not connected any SATA expansion units to a Sun StorEdge 3510 FC array, verify that mixed drive support is not enabled so that you do not see inappropriate and potentially confusing menu options and messages.

▼ To Enable or Disable Mixed Drive Support

1. From the Main Menu, choose “view and edit Configuration parameters →Disk Array Parameters →Mixed Drive Support -.”

```
< Main Menu >
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit channels
view and edit Configuration parameters
U
s
v
v
  Communication Parameters
  Caching Parameters
  Host-side Parameters
  Drive-side Parameters
  Disk Array Parameters
R
C
  Rebuild Priority - Low
  Verification on Writes
  Mixed Drive Support - Enable
```

Depending on whether Mixed Drive Support is currently enabled or disabled, a message describes the change you can make:

```
Disable Mixed Drive Support ?
```

2. Choose Yes to change the Mixed Drive Support setting or choose No to keep the current setting.

Channel Settings

The Sun StorEdge 3510 FC array and Sun StorEdge 3511 SATA array are preconfigured with the channel settings shown in “Default Channel Configurations” on page 22. Follow the procedures for configuring a channel mode if you plan to add a host connection or expansion unit, or to reassign redundant channel communications.

To make changes to channel host IDs, follow the procedures for adding or deleting a host ID.

Configuring Channel Mode

When configuring the channel mode, the following rules apply:

- You must have at least one host channel.

- Channels 0 and 1 must remain dedicated host channels.
- Channels 2 and 3 must remain drive channels.
- Channels 4 and 5 can be configured as host or drive channels.
- The mode for at least one channel must be assigned to RCCOM (redundant controller communications channel).

▼ To Modify a Channel Mode

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

The Channel Status Table is displayed.

<pre> < Main Menu > view and edit Logical drives view and edit logical Volumes view and edit Host luns </pre>										
Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid	
2	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Serial	
1	Host	NA	42	AUTO	Serial	F	NA			
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<C>	DRU+RCC	119	43	AUTO	Serial	F	NA	2 GHz	Serial	
5<C>	DRU+RCC	119	43	AUTO	Serial	F	NA	2 GHz	Serial	

The Chl column for channel 2 displays <3:C> to indicate that channel 3 is a redundant loop for channel 2. Similarly, the Chl column for channel 3 displays <2:C> to indicate that channel 2 is a redundant loop for channel 3.

2. Select the channel that you want to modify.
3. Choose “channel Mode” to display a menu of channel mode options.
4. Select the mode you want that channel to have, and then choose Yes to confirm the change.

This change does not take effect until the controller is reset.

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?

5. Choose Yes to reset the controller.

Redundant Communication Channels (RCCOM)

The RCCOM channel mode provides the communication channels by which two controllers in a redundant RAID array communicate with one another. This communication enables the controllers to monitor each other, and includes configuration updates and control of cache.

By default, channels 2 and 3 are configured as DRV + RCCOM, which combines drive and RCCOM functions on the same channel (Drive and RCCOM). In this configuration, RCCOM is distributed over all DRV + RCCOM channels, which leaves other channels free for host or drive functions.

If performance issues are particularly important, you can spread the combined DRV + RCCOM functions over four channels. Alternatively, you can configure two channels so that they are used exclusively for RCCOM, ensuring maximum I/O performance on the other host and drive channels. These two configurations are described below.

Using Four DRV + RCCOM Channels

If only channels 0 and 1 are used for communication with servers, channels 4 and 5 can be configured as DRV + RCCOM, thus providing four DRV + RCCOM channels (channels 2, 3, 4, and 5). An advantage of this configuration is that channels 4 and 5 are still available for connection of expansion units. The performance impact of RCCOM is reduced because it is now distributed over four channels instead of two. If at a later time you choose to add an expansion unit, it will not be necessary to interrupt service by resetting the controller after reconfiguring a channel.

▼ To Configure Channels 4 and 5 as Additional DRV + RCCOM Channels

1. From the Main Menu, choose “view and edit channels.”
2. Select channel 4.
3. Choose “channel Mode →Drive + RCCOM,” and then choose Yes to confirm the change.
4. Choose No to decline the controller reset, since you have another channel to reconfigure.
5. Press Enter to return to the menu.
6. Choose “Secondary controller scsi id.”
7. Specify a secondary ID (SID) that is not already in use.

You will specify this same SID for Channel 5, as shown below.

8. Choose **No** to decline the controller reset, since you have another channel to reconfigure.
9. Select channel 5.
10. Choose “channel Mode →Drive + RCCOM,” and then choose **Yes** to confirm the change.
11. Choose **No** to decline the controller reset, since you have another channel to reconfigure.
12. Press **Enter** to return to the menu.
13. Choose “Secondary controller scsi id.”
14. Specify the same secondary ID (SID) that you assigned to Channel 4.

This change does not take effect until the controller is reset, as described in the message that is displayed:

```
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?
```

15. Choose **Yes** to reset the controller.

Using Channels 4 and 5 as RCCOM-Only Channels

When only channels 0 and 1 are used for communication with servers, another option is to assign channels 4 and 5 as dedicated RCCOM channels, and then assign channels 2 and 3 as drive channels. This reduces the impact of RCCOM on the drive channels by removing RCCOM from drive channels 2 and 3. In this configuration, however, channels 4 and 5 cannot be used to communicate with hosts or to attach expansion units.



Caution – If later you reconfigure channels 4 and 5 as host or drive channels, you must restore channels 2 and 3 as DRV + RCCOM channels or the RAID array will no longer operate.

▼ To Configure Channels 4 and 5 as RCCOM-Only Channels

1. On the Main Menu, choose “view and edit channels.”
2. Select channel 4.

3. Choose “channel Mode →RCCOM,” and then choose Yes to confirm the change.
4. Choose No to decline the controller reset, since you have three more channels to reconfigure.
5. Select channel 5.
6. Choose “channel Mode →RCCOM,” and then choose Yes to confirm the change.
7. Choose No to decline the controller reset, since you have two more channels to reconfigure.
8. Select channel 2.
9. Choose “channel Mode →Drive.”
10. Choose Yes to confirm, and then choose Yes to confirm the change.
11. Choose No to decline the controller reset, since you have another channel to reconfigure.
12. Select channel 3.
13. Choose “channel Mode →Drive,” and then choose Yes to confirm the change.
This change does not take effect until the controller is reset.

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?

14. Choose Yes to reset the controller.

Host Channel IDs

Host channel IDs identify the controller to the host. Some applications require that specific IDs be assigned to host channels before they can recognize the array. Sun StorEdge 3510 FC array and Sun StorEdge 3511 SATA array default host channel IDs are shown in TABLE 3-2 and TABLE 3-3 under “Default Channel Configurations” on page 22.

The number of host IDs depends on the configuration mode:

- In point-to-point mode, only one ID should be assigned to each channel.
- In loop mode, up to 16 IDs can be assigned to each Fibre Channel, not to exceed a maximum of 32 IDs per array.

Each host ID can have up to 32 partitions, which are then mapped to LUNs to create a total not to exceed 128 in point-to-point mode or 1024 in loop mode. To map 1024 partitions in loop mode, you must add additional host IDs so that 32 IDs are mapped to the array's channels. Several 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 an FC or SATA Array (Optional, Loop Mode Only)" on page 163.

Each host channel has a unique primary and secondary ID available. 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. You can:

- Edit each host ID to change the ID of each controller host channel that is seen by the host. To change an ID, you must delete it first and then add the new ID.
- Add host IDs for loop configurations.

Note – Channel ID values of 0 to 125 are accessed in eight ranges of IDs. When you change a channel's mode, the channel ID might change to an ID that is not in the range you want to use. See "Channel ID Ranges" on page 94 for a description of channel ID ranges and a procedure for changing the ID range.

▼ To Add or Delete a Unique Host ID

Note – To change an ID, you must first delete it and then add the new ID.

1. From the Main Menu, choose "view and edit channels."
2. Select the host channel on which you want to add or change an ID.
3. Choose "view and edit scsi Id."
If host IDs have already been configured on the host channel, they will be displayed.
4. If no host IDs have been configured, choose Yes when the following message is displayed.

No SCSI ID Assignment - Add Channel SCSI ID?

5. If a host ID has already been assigned to that channel, select an ID.
6. To delete an ID, choose "Delete Channel SCSI ID," and then choose Yes to confirm the deletion.

7. To add an ID, choose "Add Channel SCSI ID."
8. Select a controller from the list to display a list of IDs.
9. Select an ID from the list, and then choose Yes to confirm your choice.

This change does not take effect until the controller is reset.

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?

10. Choose Yes to reset the controller.

Channel ID Ranges

The ID values of 0 to 125 are available when you assign a channel ID. These IDs are accessed in the eight ranges shown in TABLE 5-1.

TABLE 5-1 ID Values Assigned to Each ID Range

Range	Available ID Numbers
0	0 to 15
1	16 to 31
2	32 to 47
3	48 to 63
4	64 to 79
5	80 to 95
6	96 to 111
7	112 to 125

Once an ID is assigned to a channel, if you decide to add an ID, the only IDs that are initially displayed are those in the range of the first ID you assigned. For example, if you initially assign an ID of 40 to host channel 0, when you add IDs to host channel 0, only IDs in Range 2 (32 to 47) are available.

▼ To Assign an ID From a Different Range

1. Choose “view and edit channelS” to display the Channel Status table.
2. Select the channel whose ID range you want to change.
3. Choose “view and edit scsi Id.”
4. Select a controller.

Note – To change an ID, you must first delete it and then add the new ID.

5. Choose “Delete Channel SCSI ID,” and then choose **Yes** to confirm the deletion.
This change does not take effect until the controller is reset.

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?

6. If other IDs exist on the selected channel, choose **No** and repeat Step 5 to delete every ID configured on the channel.
7. After deleting the last ID, choose **Yes** to reset the controller.

When all IDs have been deleted, you can assign an ID from a different range.

No SCSI ID Assignment - Add Channel SCSI ID ?

8. Choose **Yes** to assign an ID.

9. Select the controller to which you want to assign an ID.

A list of IDs is displayed. Depending on the current range, adjoining ranges are displayed at the top and bottom of the ID list, except ranges 0 and 7, which only display one adjoining range. In the following example, range 7 is displayed.

< Main Menu >												
view and edit Logical drives												
view and edit logical Volumes												
view and edit Host luns												
U	Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	Cur	<To Range 6>	
U	0	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Ser	ID 112	
U	1	Host	NA	42	AUTO	Serial	F	NA			ID 113	
S	2< No SCSI ID Assignment - Add Channel SCSI ID ?									GHz	Ser	ID 114
U	3< Yes No									GHz	Ser	ID 115
U	4	Host	112	NA	AUTO	Serial	F	NA	2 GHz	Ser	ID 116	
U	5	Host	112	NA	AUTO	Serial	F	NA	2 GHz	Ser	ID 117	
											ID 118	
											ID 119	
											ID 120	
											ID 121	
											ID 122	
											ID 123	
											ID 124	
											ID 125	

10. To change to a different range, select an adjoining range.

IDs in the newly selected range are displayed.

11. Repeat Step 10 until the desired range is displayed.

12. Select an ID from the desired range, and then choose Yes to confirm the assignment.

This change does not take effect until the controller is reset.

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?

13. Choose Yes to reset the controller.

Fibre Connection Protocol

See “Fibre Connection Protocol Guidelines” on page 33 for a detailed description of Loop and Point-to-Point operation.

▼ To Change the Fibre Connection Protocol

1. **From the Main Menu, choose “view and edit Configuration parameters →Host-side Parameters →Fibre Connection Option.”**

The fibre connection that is currently assigned to the array is displayed.

2. **Choose “Loop only” or “Point to point only” as appropriate.**

Note – Do not use the command, “**Loop preferred, otherwise point to point.**” This command is reserved for special use and should be used only if you are directed to do so by technical support.

This change does not take effect until the controller is reset.

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?

3. **Choose Yes to reset the controller.**

Creating Logical Drives

The RAID array is preconfigured with one RAID 0 logical drive as described in “Default Logical Drive Configuration” on page 22. Each logical drive consists of a single partition by default.

This section describes how to modify the RAID level or add more logical drives. In these procedures, you configure a logical drive to contain one or more physical drives based on the desired RAID level, and divide the logical drive into additional partitions.

Note – Depending on the size and RAID level, it can take up to several hours to build a logical drive. Online initialization, however, enables you to begin configuring and using the logical drive before initialization is complete.

Creating a 2-Tbyte RAID 5 logical drive can take up to:

- 2.25 hours for Sun StorEdge 3310 SCSI arrays and Sun StorEdge 3510 FC arrays
- 10.3 hours for Sun StorEdge 3511 SATA arrays

Preparing for Logical Drives Larger Than 253 Gbyte (Solaris Operating System Only)

The Solaris operating system requires drive geometry for various operations, including `newfs`. For the appropriate drive geometry to be presented to the Solaris operating system for logical drives larger than 253 Gbyte, change the default settings to cover all logical drives over 253 Gbyte. These settings work for smaller configurations as well. The controller automatically adjusts the sector count so the operating system can read the correct drive capacity.

For Solaris operating system configurations, use the values in the following table.

TABLE 5-2 Cylinder and Head Mapping for the Solaris Operating System

Logical Drive Capacity	Cylinder	Head	Sector
< 253 GB	< 65536 (default)	Variable	Variable (default)
253 GB–1 TB	< 65536 (default)	64 (default)	variable (default)

After settings are changed, they apply to all logical drives in the chassis.

Note – Refer to your operating system documentation for limitations on device sizes.

▼ To Change Cylinder and Head Settings

1. Choose “view and edit Configuration parameters →Host-side Parameters →Host Cylinder/Head/Sector Mapping Configuration →Sector Ranges - Æ Variable,” and then choose Yes to confirm your choice.
2. Choose “Head Ranges - →64 Heads,” and then choose Yes to confirm your choice.

3. Choose “Cylinder Ranges - →× 65536,” and then choose Yes to confirm your choice.

▼ To Create a Logical Drive

Note – To reassign drives and add local or global spare drives on your preconfigured array, you must first unmap and then delete the existing logical drives. For more information about deleting a logical drive, see “Deleting Logical Drives” on page 83.

1. From the Main Menu, choose “view and edit Logical drives.”
Unassigned logical drives show a RAID level of NONE.
2. Select the first available unassigned logical drive (LG).

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	594EB542	NA	RAID1	34476	GOOD				7	B	2	0	0	
1			NONE											
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

You can create as many as 32 logical drives using physical drives on any loop.

If mixed drive support is enabled, a menu of drive types is displayed. If mixed drive support is disabled, proceed to the next step. See “To Enable or Disable Mixed Drive Support” on page 88 for information about mixed drive support.

- If mixed drive support is enabled, select the type of logical drive to create.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	594EB542	NA	RAID1	34476	GOOD					7	B	2	0	0
1			NONE											
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

- When prompted to “Create Logical Drive?” choose Yes to confirm your choice and display a pull-down list of supported RAID levels.
- Select a RAID level from the list to assign to the logical drive.

Note – RAID 5 is used as an example in the following steps.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	RAID 5	#FL	NAME
P0	594EB542	NA	RAID1	34476	GOOD					7	0	
1			NONE							RAID 3 RAID 1 RAID 0 NRAID		
2			NONE									
3			NONE									
4			NONE									
5			NONE									
6			NONE									
7			NONE									

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

For more information about RAID levels, see “RAID Levels” on page 316.

- Select the drives you want to include in the logical drive from the list of available physical drives, using the steps below.
 - You must select at least the minimum number of drives required for the selected RAID level.
 - RAID 3 and RAID 5 require a minimum of three physical drives.
 - RAID 0 requires a minimum of two physical drives.

- RAID 1 requires a minimum of two physical drives. Additional drives must be added in increments of two.

For redundancy, you can create a logical drive containing drives distributed over separate channels. You can then create several partitions on each logical drive. In a RAID 1 or RAID 0+1 configuration, the order in which you select the physical drives for a logical drive determines the channels to which the physical drives are assigned. If you want drives to be mirrored over two channels, select them in the appropriate order. For example:

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

Note – Logical drives that include both Fibre Channel drives and SATA drives are not supported. If you have enabled mixed drive support, only the appropriate drive types are displayed.

- Use the up and down arrow keys and press Return to select the drives you want to include in the logical drive.**

An asterisk mark (*) is displayed in the Chl (Channel) column of each selected physical drive.

Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
*2<3>	6	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
*2<3>	7	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
*2<3>	8	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
2<3>	10	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G

- To deselect a drive, press Return again on the selected drive.**

The asterisk marking that drive disappears.

- After all physical drives have been selected for the logical drive, press Escape.**

Several optional menu options are displayed. You can choose these menu options to define aspects of the logical drive you are creating:

- “**Maximum Drive Capacity**” enables you to specify the size of the logical drive.
- “**Assign Spare Drives**” enables you to specify a local spare drive to be used if an existing physical drive in the logical drive you are creating becomes defective.

- **“Disk Reserved Space”** displays the size of the reserved space used to store logical drive metadata. While it is possible to delete or change reserved space size, do not do so. See **“To Specify Disk Reserved Space”** on page 198 for more information.
- **“Logical Drive Assignments”** enables you to assign the logical drive you are creating to either the primary or secondary controller.
- **“Write Policy”** enables you to set the cache write policy for the logical drive you are creating.
- **“Initialize Mode:”** enables you specify whether the logical drive you are creating is initialized online or offline.
- **“Stripe Size:”** enables you to specify the stripe size for the logical drive you are creating.

These menu options are described in the remainder of this section.

7. (Optional) Set the maximum logical drive capacity, using the following procedure:

- a. Choose “Maximum Drive Capacity.”**

Note – Changing the maximum drive capacity reduces the size of the logical drive and leaves some disk space unused.

- b. Specify the maximum capacity of each physical drive that makes up the logical drive you are creating.**

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	594EB542	NA	RAID1	34476	GOOD				7	B	2	0	0	
1			NONE											
<div style="border: 1px solid black; padding: 5px;"> <p>Maximum Drive Capacity : 34476MB</p> <p>Maximum Available Drive Capacity(MB): 34476</p> <p>Maximum Drive Capacity(MB) : 1024</p> </div>														
6			NONE											
7			NONE											

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.

8. (Optional) Add a local spare drive from the list of unused physical drives, by following these steps:

- a. Choose “Assign Spare Drives” to display a list of available physical drives you can use as a local spare.**

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

Note – An NRAID or RAID 0 logical drive has no data redundancy or parity and does not support spare drive rebuilding.

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

b. Select a physical drive from the list to use as a local spare.

Slot	Ch1	ID	Size<MB>	Speed	LG_DRU	Status	Vendor and Product ID
*	2<3>	8	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	2<3>	10	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G

c. Press Escape to return to the menu of logical drive options.

Note – The Disk Reserved Space option is not supported while you are creating a logical drive.

If you use two controllers for a redundant configuration, you can assign a logical drive to either of the controllers to balance the workload. By default, all logical drives are assigned to the primary controller.

Logical drive assignments can be changed later, but that operation requires that you unmap host LUNS and reset the controller.

9. (Optional) For dual-controller configurations, you can assign this logical drive to the secondary controller by following these steps:



Caution – In single-controller configurations, assign logical drives only to the primary controller.

a. Choose “Logical Drive Assignments.”

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

b. Choose Yes to assign the logical drive to the redundant controller.

10. (Optional) Configure the logical drive's write policy.

Write-back cache is the preconfigured global logical drive write policy, which is specified on the Caching Parameters submenu. (See "Enabling and Disabling Write-Back Cache" on page 231 for the procedure on setting the global caching parameter.) This option enables you to assign a write policy per logical drive that is either the same as or different than the global setting. Write policy is discussed in more detail in "Cache Write Policy Guidelines" on page 32.

a. Choose "Write Policy -."

Note – The Default write policy displayed is the global write policy assigned to all logical drives.

The following write policy options are displayed:

- **Default (default).** Assigns the global write policy. If the global setting for write policy is changed, changing this setting automatically changes the write policy for this logical drive.

As described in "Cache Write Policy Guidelines" on page 32, the array can be configured to dynamically switch write policy from write-back cache to write-through cache if specified events occur. Write policy is only automatically switched for logical drives with write policy configured to Default. See "Event Trigger Operations" on page 267 for more information.

- **Write-Back.** Assigns write-back cache regardless of any changes to the global write policy.
- **Write-Through.** Assigns write-through cache regardless of any changes to the global write policy.

b. Choose a write policy option.



Note – You can change the logical drive logical drive's write policy at any time, as explained in "Changing Write Policy for a Logical Drive" on page 148.

11. (Optional) Set the logical drive initialization mode by choosing "Initialize Mode" from the list of logical drive options, and then choosing Yes to change the initialization mode.

The assigned initialization mode is displayed in the list of logical drive options.

You can choose between these two logical drive initialization options:

- On-Line (default)

This option enables you to configure and use the logical drive before initialization is complete. Because the controller is building the logical drive while performing I/O operations, initializing a logical drive online requires more time than off-line initialization.

- Off-Line

This menu option enables you to configure and use the drive only after initialization is complete. Because the controller is building the logical drive without having to also perform I/O operations, offline initialization requires less time than on-line initialization.

Because logical drive initialization can take a considerable amount of time, depending on the size of your physical disks and logical drives, you can choose online initialization so that you can use the logical drive before initialization is complete.

12. (Optional) Configure the logical drive stripe size.

Depending on the optimization mode selected, the array is configured with the default stripe sizes shown in “Cache Optimization Mode and Stripe Size Guidelines” on page 30. When you create a logical drive, however, you can assign a different stripe size to that logical drive.

Note – Default stripe sizes result in optimal performance for most applications. Selecting a stripe size that is inappropriate for your optimization mode and RAID level can decrease performance significantly. For example, smaller stripe sizes are ideal for I/Os that are transaction-based and randomly accessed. But when a logical drive configured with a 4-Kbyte stripe size receives files of 128 Kbyte, each physical drive has to write many more times to store it in 4-Kbyte data fragments. Change stripe size only when you are sure it will result in performance improvements for your particular applications.

See “Cache Optimization Mode and Stripe Size Guidelines” on page 30 for more information.

Note – Once a logical drive is created, its stripe size cannot be changed. To change the stripe size, you must delete the logical drive, and then recreate it using the new stripe size.

a. Choose Stripe Size.

A menu of stripe size options is displayed.

- If you specified random optimization, in addition to Default, you can choose among disk stripe sizes of 4 Kbyte, 8 Kbyte, 16 Kbyte, 32 Kbyte, 64 Kbyte, 128 Kbyte, and 256 Kbyte.

- If you specified sequential optimization, in addition to Default, you can choose among disk stripe sizes of 16 Kbyte, 32 Kbyte, 64 Kbyte, 128 Kbyte, and 256 Kbyte.

b. Choose Default to assign the stripe size per optimization mode, or choose a different stripe size from the menu.

Default stripe size per optimization mode is shown in “Cache Optimization Mode and Stripe Size Guidelines” on page 30.

The selected stripe size is displayed in the list of logical drive options.

13. Once all logical drive options have been assigned, press Escape to display the settings you have chosen.

```
Raid Level           : RAID 5
Online SCSI Drives  : 3
Maximum Drive Capacity : 20000 MB
Disk Reserved Space  : 256 MB
Spare SCSI Drives   : 0
Logical Drive Assignment: Secondary Controller
Write Policy        : Default<Write-Back>
Initialize Mode     : On-Line
Stripe Size         : 128K Bytes

                Create Logical Drive ?
                 Yes       No
```

14. Verify that all information is correct, and then choose Yes to create the logical drive.

Note – If the logical drive has not been configured correctly, select No to return to the logical drive status table so that you can configure the drive correctly.

Messages indicate that the logical drive initialization has begun, and then that it has completed.

15. Press Escape to close the drive initialization message.

A progress bar displays the progress of initialization as it occurs.

You can press Escape to remove the initialization progress bar and continue working with menu options to create additional logical drives. The percentage of completion for each initialization in progress is displayed in the upper left corner of the window.

i2:0%														
BAT:++++														
LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	13843684	NA	RAID5	103428	GOOD				7	B	4	1	0	
1					On-Line Initializing Notification									
2				LG:2 Logical Drive NOTICE: Starting Creation										
3														
4			NONE											
5			NONE											
6			NONE											
7			NONE											

The following message is displayed when the initialization is completed:

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	594EB542	NA	RAID1	34476	GOOD				7	B	2	0	0	
1					On-Line Initializing Notification									
2				On-Line Initialization of Logical Drive 1 Completed										
3														
4			NONE											
5			NONE											
6			NONE											
7			NONE											

16. Press Escape to dismiss the notification.

The newly created logical drive is displayed in the status window.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	594EB542	NA	RAID1	34476	GOOD				7	B	2	0	0	
\$1	4F342FDA	NA	RAID5	40000	GOOD				7	B	3	0	0	
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

Controller Assignment

By default, logical drives are automatically assigned to the primary controller. If you assign half of the logical drives to the secondary controller in a dual controller array, 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).



Caution – In single-controller configurations, 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. In a single-controller configuration, if you disable the Redundant Controller function and reconfigure the controller with the Autoconfigure option or as a secondary controller, the controller module becomes inoperable and will need to be replaced.

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 a Partition to a Host LUN” on page 112).

▼ To Change a Controller Assignment (Optional)



Caution – Assign logical drives only to primary controllers in single-controller configurations.

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the logical drive you want to reassign.
3. Choose “logical drive Assignments,” and then choose Yes to confirm the reassignment.

The reassignment is evident from the “**view and edit Logical drives**” screen. A “P” in front of the LG number, such as “P0,” 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 the secondary controller.

Logical Drive Name

You can assign a name to each logical drive. These logical drive names are used only in RAID firmware administration and monitoring and do not appear anywhere on the host. You can also edit this drive name.

▼ To Assign a Logical Drive Name (Optional)

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select a logical drive.
3. Choose “logical drive Name.”
4. Type the name you want to give the logical drive in the New Logical Drive Name field and press Return to save the name.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	594EB542	NA	RAID1	34476	GOOD				7	B	2	0	0	
S1	4F342FDA	NA	RAID5	40000	GOOD				7	B	3	0	0	

View scsi drives														
Delete logical drive														
Partition logical drive														
logical drive Name														
l														
E														
a														
r														
c														
N														
s														
Shutdown logical drive														
Write policy														

Partitions

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 and 1024 LUN assignments (loop mode only). For guidelines on setting up 1024 LUNs, see “Planning for 1024 LUNs on an FC or SATA Array (Optional, Loop Mode Only)” on page 163.



Caution – If you modify the size of a partition or logical drive, all data on the drive is lost.

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

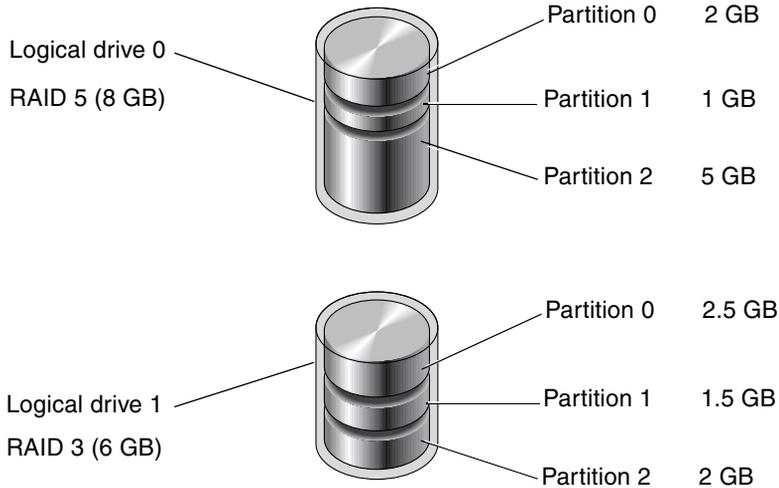


FIGURE 5-1 Partitions in Logical Drives

▼ To Partition a Logical Drive (Optional)



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

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the logical drive you want to partition.
3. Choose “Partition logical drive.”

If the logical drive has not already been partitioned, the following warning is displayed:

```
This operation may result in the LOSS OF ALL DATA on the Logical Disk.
```

```
Partition Logical Drive?
```

4. Choose Yes to continue.

A list of the partitions on 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 a partition.

6. Type the desired size of the selected partition.

The following warning is displayed:

This operation will result in the LOSS OF ALL DATA on the partition.
Partition Logical Drive?

7. Choose Yes to partition the drive.

The remaining capacity of the logical drive is automatically allocated to the next partition. In the following example, a partition size of 20000 Mbyte was entered; the remaining storage of 20000 Mbyte is allocated to the partition below the newly created partition.

Partition	Offset(MB)	Size(MB)
0	0	20000
1	20000	20000
2		
3		
4		
5		
6		
7		

8. Repeat Step 5 through Step 7 to partition the remaining capacity of your logical drive.

For information on deleting a partition, see “Deleting a Logical Drive Partition” on page 130.

Mapping a Partition to a Host LUN

A partition is a division of the logical drive that appears as a physical drive to any host that has access to that partition. You can create a maximum of 32 partitions per logical drive. So that host bus adapters (HBAs) recognize the partitions when the host bus is reinitialized, each partition must be mapped to a host LUN (logical unit number). Two methods can be used to map a partition to a host:

- LUN mapping
- LUN filtering

Note – When you modify a partition, you must first unmap the LUN.

Note – If you plan to map 128 or more LUNs, the process is easier if you use Sun StorEdge Configuration Service. Refer to the *Sun StorEdge 3000 Family Configuration Service User's Guide* for more information.

LUN Mapping

Map a partition to a LUN on a host channel to create a connection between that host channel and the partition. Note that with LUN mapping, all hosts on the mapped host channel have full access to all partitions mapped to LUNs on that channel. To provide redundant connections between a host and a partition, map the partition to a LUN on both of the host channels that connect with that host.

With LUN mapping, only one partition can be mapped to each LUN. To assign multiple partitions to the same LUN, use LUN filtering rather than LUN mapping. LUN mapping is most effective when only one host is connected to a host channel.

Channel IDs represent the physical connection between the HBA and the array. The host ID is an identifier assigned to the channel so that the host can identify LUNs. The following figure shows the relationship between a host ID and a LUN.

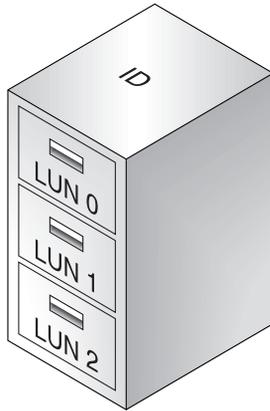


FIGURE 5-2 LUNs Resemble Drawers in a File Cabinet

The ID is like a cabinet and the drawers are like the LUNs.

- The maximum number of LUNs that can be created for a Sun StorEdge 3510 FC array or a Sun StorEdge 3511 SATA array in a point-to-point configuration is 128.
- The maximum number of LUNs that can be created for a Sun StorEdge 3510 FC array or a Sun StorEdge 3511 SATA array in a loop configuration is 1024. To create a total of 1024 LUNs, see “Planning for 1024 LUNs on an FC or SATA Array (Optional, Loop Mode Only)” on page 163.

The following figure illustrates mapping partitions to host ID/LUNs.

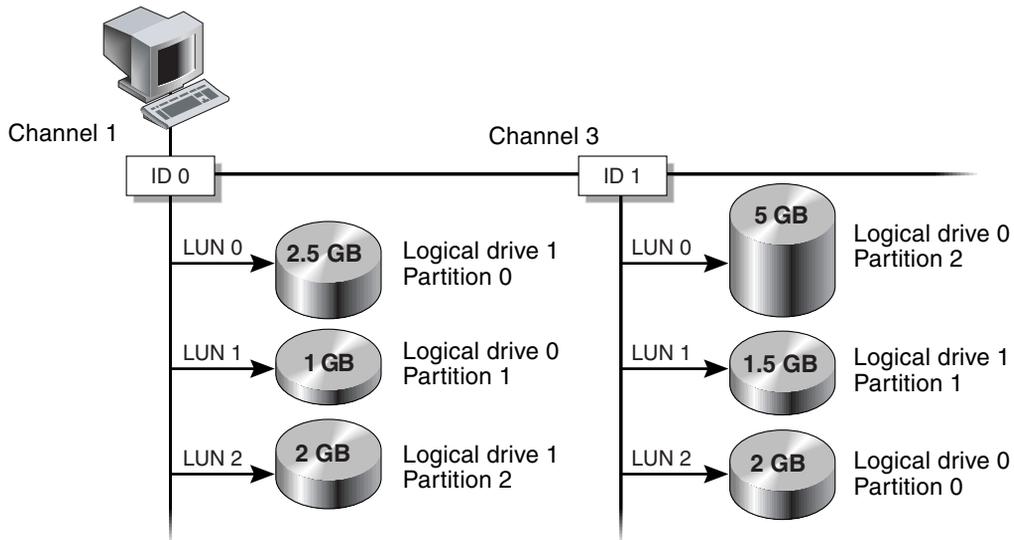


FIGURE 5-3 Mapping Partitions to Host ID/LUNs

For detailed instructions for LUN mapping, see “To Map a Logical Drive Partition” on page 116.

LUN Filtering (FC and SATA Only)

For multiple servers connected to the same FC array, LUN filtering provides an exclusive path from a server to a logical drive and essentially hides or excludes the other connected servers from seeing or accessing the same logical drive. That is, the LUN filter organizes how the array devices are accessed and viewed from host devices, and typically maps an array device to only one host so that other hosts do not access and use the same array device.

LUN filtering also enables multiple hosts to be mapped to the same LUN, allowing different servers to have their own LUN 0 to boot from, if needed. Even though host filters are created on the same LUN, each host filter can provide individual hosts exclusive access to a different partition, and even access to partitions on different logical drives. Host filters can also grant different levels of access to different hosts. LUN filtering is also valuable in clarifying mapping when each HBA typically sees twice the number of logical drives when viewed through a hub.

Each Fibre Channel device is assigned a unique identifier called a worldwide name (WWN). A WWN is assigned by the IEEE and is similar to a MAC address in IP or a URL on the Internet. These WWNs stay with the device for its lifetime. LUN filtering uses this WWN to specify which server is to have exclusive use of a specific logical drive.

As shown in the following example, when you map LUN 01 to host channel 0 and select WWN1, server A has a proprietary path to that logical drive. All servers continue to see and access LUN 02 and LUN 03 unless filters are created on them.

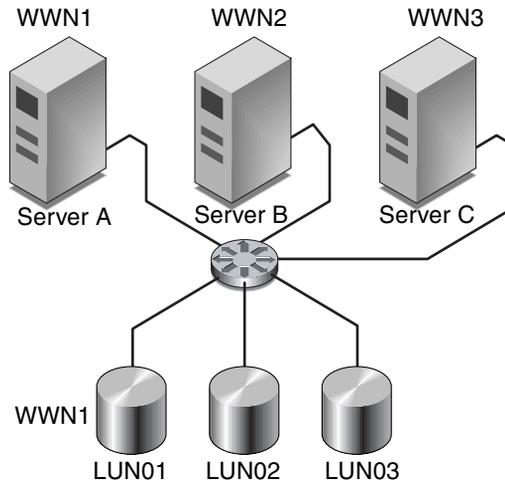


FIGURE 5-4 Example of LUN Filtering

Note – It is possible to see differing information when a fabric switch queries the WWN of an array. When the RAID controller does a Fibre Channel fabric login to a switch, during the fabric login process the switch obtains the WWN of the RAID controller. In this case, the switch displays the company name. When the switch issues an inquiry command to a mapped LUN on the array, the switch obtains the company name from the inquiry data of the LUN. In this case, the switch displays Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array, which is the inquiry data returned by the RAID controller.

Prior to using the LUN filter feature, identify which array is connected to which HBA card, and the WWN assigned to each card. This procedure varies according to the HBA you are using. Refer to the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array for instructions on identifying the WWN for your host.

For detailed instructions for creating host filters, see, “LUN Filtering (FC and SATA Only)” on page 114.

Note – You can create a maximum of 128 host filters. You can create a maximum of 64 WWNs.

Note – The process of creating host filters is easier using Sun StorEdge Configuration Service.

▼ To Map a Logical Drive Partition

1. From the Main Menu, choose “view and edit Host luns.”
A list of available channels, IDs, and their associated controllers is displayed.
2. Select a channel and host ID on the primary controller.
3. If the Logical Drive and Logical Volume menu options are displayed, choose “Logical Drive” to display the LUN table.
4. Select the LUN you want to map the drive to.
A list of available logical drives is displayed.
5. Select the logical drive (LD) that you want to map to the selected LUN.
The partition table is displayed.
6. Select the partition you want to map to the selected LUN.

< Main Menu >														
view and edit Logical drives				LUN	LU/LD	DRU	Partition	Size<MB>	RAID					
view and edit logical Volumes				0										
view and edit Host luns														
LG	ID	LU	RAID	Size<MB>	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	56AEFA18	NA	RAID5	68850	GOOD				5	B	3	0	0	
Partition				Offset<MB>	Size<MB>	3								
0				0	68850	4								
						5								
						6								
						7								

7. Choose “Map Host LUN,” and then choose Yes to confirm the host LUN mapping.

LUN	LU/LD	DRU	Partition	Size<MB>	RAID
0					
Map Logical Drive: 0 Partition : 0 To Channel : 0 ID : 40 Lun : 0 ? <input checked="" type="checkbox"/> Yes <input type="checkbox"/> No					
5					
6					
7					

The partition is now mapped to the selected LUN.

< Main Menu > view and edit Logical drives view and edit logical Volumes view and edit Host luns		LUN	LU/LD	DRV	Partition	Size<MB>	RAID
		0	LD	0	0	68850	RAID5
v	CHL 0 ID 40 <Primary Control	1					
v	CHL 1 ID 42 <Secondary Contr						
v	CHL 4 ID 50 <Primary Control	2					
s	CHL 5 ID 51 <Secondary Contr						
v	Edit Host-ID/WWN Name List	3					
v		4					
		5					
		6					
		7					

8. Repeat Step 4 through Step 7 to map additional partitions to host LUNs on this channel and logical drive.
9. Press Escape.
10. If you are LUN mapping a redundant configuration, repeat Step 2 through Step 7 to map partitions to host LUNs with other IDs on the logical drive assigned to the primary controller.

When you map a partition to two channels in a redundant configuration, the number in the Partition column of the partition table displays an asterisk (*) to indicate that the partition is mapped to two LUNs.

Note – If you are using host-based multipathing software, map each partition to two or more host IDs so multiple paths will be available from the partition to the host.

11. Repeat Step 2 through Step 10 to map hosts to the secondary controller.
12. To verify unique mapping of each LUN (unique LUN number, unique DRV number, or unique Partition number):
 - a. From the Main Menu, choose “view and edit Host luns.”
 - b. Select the appropriate controller and ID and press Return to review the LUN information.

A mapped LUN displays a number and a filtered LUN displays an “M” for masked LUN in the host LUN partition window.
13. When all host LUNs have been mapped, save the updated configuration to nonvolatile memory. See “Saving Configuration (NVRAM) to a Disk” on page 125 for more information.

- (Solaris operating system only) For the Solaris operating system to recognize a LUN, you must first manually write the label using the `Auto configure` option of the `format (1M)` utility, as described in “To Label a LUN” on page 123.

▼ To Create Host Filters (FC and SATA Arrays Only)

- From the Main Menu, choose “view and edit Host luns.”
A list of available channels and their associated controllers is displayed.
- Select a channel and host ID.
- If the Logical Drive and Logical Volume menu options are displayed, choose Logical Drive.
- Select the LUN for which you want to create the host filter.
A list of available logical drives is displayed.
- Select the logical drive (LD) for which you want to create a host filter.
- Select the partition for which you want to create a host filter.
- Choose “Create Host Filter Entry →Add from current device list.”

< Main Menu >						
view and edit Logical drives	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit logical Volumes	0	LD	0	0	68850	RAID5
view and edit host luns						
v						
v	CHL 0 ID 40 <Primary Control	1				
v	CHL 1 ID 42 <Secondary Contr					
v	CHL 4 ID 50 <Primary Control					
s	CHL 5 ID 51 <Secondary Contr					
v	Edit Host-ID/WWN Name List					
				Map Host LUN		
				Create Host Filter Entry		
				add from current device list		
				Manual add host filter entry		
	4					
	5					
	6					
	7					

This step automatically performs a discovery of the attached HBAs and displays a list of WWNs. This list includes:

- All worldwide names that HBAs on the selected channel have propagated to the array.
- All WWNs on the selected channel that have been manually entered using the “Manual add host filter entry” option.

- All WWNs that have been manually added to the “**Edit Host-ID/WWN Name List**” (“**view and edit Host luns** →**Edit Host-ID/WWN Name List**”), including worldwide names from HBAs that are not connected to the selected channel. See “**Manually Adding WWN Entries Using the Host-ID/WWN Name List**” on page 168 for more information about this menu option.

When you select a worldwide name from this list, ensure that the worldwide name you select is from an HBA on the channel where you are creating the filter.

Alternatively, you can add a worldwide name manually by choosing “**Manually add host filter entry**” rather than “**Add from current device list.**” Then type the Host-ID/WWN in the text area provided and press Return. When you manually enter a worldwide name using the “**Manually add host filter entry**” menu option, that WWN only appears in the list of WWNs when you are creating a filter on a channel where the WWN was initially added.

8. From the device list, select the WWN number of the server for which you are creating a filter, and choose Yes to confirm your choice.

A filter configuration screen displays the filter you are creating.

< Main Menu >						
view and edit Logical drives	LUN	LU/LD	DRU	Partition	Size<MB>	RAID
view and edit logical Volumes	0	LD	0	0	68850	RAID5
view and edit Host luns	1					
u CHL 0 ID 40 <Primary Control	Map Host LUN					
u CHL 1 ID 42 <Secondary Contr	Create Host Filter Entry					
u CHL 4 ID 50 <Primary Control	Host-ID/WWN					
s CHL 5 ID 51 <Secondary Contr	WWPN:0x210100E08B2139EA					
v Edit Host-ID/WWN Name List	Host-ID/WWN:0x7473657468636574 <sscsMgr>					
	Host-ID/WWN:0x216000C0FF804DE2 <server08>					
	Host-ID/WWN:0x0000000000000000 <sscsMgr>					
	6					
	7					

9. Review the filter configuration screen. Make any changes necessary by selecting the setting you want to change.

< Main Menu >						
view and edit Logical drives	LUN	LU/LD	DRU	Partition	Size<MB>	RAID
view and edit logical Volumes	0	LD	0	0	68850	RAID5
view and edit Host luns	1					
u CHL 0 ID 40 <Primary Control	Map Host LUN					
u CHL 1 ID 42 <Secondary Contr	Create Host Filter Entry					
u CHL 4 ID 50 <Primary Control	Logical Drive 1 Partition 0					
s CHL 5 ID 51 <Secondary Contr	4				Host-ID/WWN - 0x210100E08B2139EA	
v Edit Host-ID/WWN Name List	5				Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFFF	
	6				Filter Type - Include	
					Access Mode - Read/Write	
					Name - Not Set	
	7					

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

< Main Menu >						
view and edit Logical drives	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit logical Volumes	0	LD	0	0	68850	RAID5
view and edit host luns	1					
v CHL 0 ID 40 <Primary Control	Map Host LUN					
v CHL 1 ID 42 <Secondary Contr	Create Host Filter Entry					
v CHL 4 ID 50 <Primary Control	Logical Drive 1 Partition 0					
s CHL 5 ID 51 <Secondary Contr	4	Host-ID/WWN - 0x210100E08B2139EA				F
v Edit Host-ID/WWN Name List	5	Host-ID/WWN:				
	6					
	7					

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

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

< Main Menu >						
view and edit Logical drives	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit logical Volumes	0	LD	0	0	68850	RAID5
view and edit host luns	1					
v CHL 0 ID 40 <Primary Control	Map Host LUN					
v CHL 1 ID 42 <Secondary Contr	Create Host Filter Entry					
v CHL 4 ID 50 <Primary Control	Logical Drive 1 Partition 0					
s CHL 5 ID 51 <Secondary Contr	4	Host-ID/WWN - 0x210100E08B2139EA				
v Edit Host-ID/WWN Name List	5	Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFFF				
	6	Host-ID/WWN Mask:				
	7					

- c. To change the filter setting, select “Filter Type -,” and choose Yes to exclude or include the Host-ID/WWN selection.

Choose “Filter Type to Include” to grant LUN access to the host identified by the WWN and WWN Mask. Choose “Filter Type to Exclude” to deny the identified host LUN access.

< Main Menu >						
view and edit Logical drives	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit logical Volumes	0	LD	0	0	68850	RAID5
view and edit host luns	1					
v CHL 0 ID 40 <Primary Control	Map Host LUN					
v CHL 1 ID 42 <Secondary Contr	Create Host Filter Entry					
v CHL 4 ID 50 <Primary Control	Logical Drive 1 Partition 0					
s CHL 5 ID 51 <Secondary Contr	4	Host-ID/WWN - 0x210100E08B2139EA				
v Edit Host-ID/WWN Name List	5	Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFFF				
	6	Filter Type - Include				
	6	Set Filter Type to Exclude ?				
	7	Yes No				

Note – If no host has been granted access to the selected LUN (by having its Filter Type set to Include), all hosts can access that LUN. In this configuration, you can deny specific hosts access to that LUN by configuring their Filter Type to Exclude. Once any host is granted access to a LUN, only hosts with explicit access (Filter Type set to Include) can access that LUN.

- d. To change the access mode, which assigns Read-Only or Read/Write privileges, select "Access mode -," and choose Yes to confirm the assignment.

< Main Menu >						
view and edit Logical drives	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit logical Volumes	0	LD	0	0	68850	RAID5
view and edit host luns	1					
u CHL 0 ID 40 <Primary Control	Map Host LUN					
u CHL 1 ID 42 <Secondary Contr	Create Host Filter Entry					
u CHL 4 ID 50 <Primary Control	4	L	Set Access Mode to Read-Only ?			
u CHL 5 ID 51 <Secondary Contr	H		Yes	No	F	
s Edit Host-ID/WWN Name List	5	F	Access Mode - Read/Write			
	6		Name - Not Set			
	7					

- e. To set a name for the filter, select "Name -." Type the name you want to use and press Return.

10. Verify all settings and press Escape to continue.

< Main Menu >						
view and edit Logical drives	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
view and edit logical Volumes	0	LD	0	0	68850	RAID5
view and edit host luns	1					
u CHL 0 ID 40 <Primary Control	Map Host LUN					
u CHL 1 ID 42 <Secondary Contr	Create Host Filter Entry					
u CHL 4 ID 50 <Primary Control	4		Logical Drive 1 Partition 0			
u CHL 5 ID 51 <Secondary Contr	5		Host-ID/WWN - 0x210100E08B2139Ea			
s Edit Host-ID/WWN Name List	6		Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFFF			
			Filter Type - Include			
			Access Mode - Read/Write			
			Name - Not Set			
	7					

11. Verify all filter entries and press Escape.

12. Choose Yes to add the host filter entry.

< Main Menu > view and edit Logical drives view and edit logical Volumes view and edit Host luns		LUN	LU/LD	DRU	Partition	Size<MB>	RAID
		0	LD	0	0	68850	RAID5
v		1					
v	CHL 0 ID 40 <Primary Control						
v	CHL 1 ID 42 <Secondary Contr						
v	CHL 4 ID 50 <Primary Control						
s	CHL 5 ID 51 <Secondary Contr						
v	Edit Host-ID/WWN Name List						
		4					
		5					
		6					
		7					

Map Host LUN	
Create Host Filter Entry	
Add Host Filter Entry ?	
Yes	No

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 complete the host filter entry in Step 14.

13. At the server list, repeat the previous steps to create additional filters, or press Escape to continue.

< Main Menu > view and edit Logical drives view and edit logical Volumes view and edit Host luns		LUN	LU/LD	DRU	Partition	Size<MB>	RAID
		0	LD	0	0	68850	RAID5
v		1					
v	CHL 0 ID 40 <Primary Control						
v	CHL 1 ID 42 <Secondary Contr						
v	CHL 4 ID 50 <Primary Control						
s	CHL 5 ID 51 <Secondary Contr						
v	Edit Host-ID/WWN Name List						
		6					
		7					

Map Host LUN	
Create Host Filter Entry	
Host-ID/WWN	
WWPN:0x216100E08B2137EA	
Host-ID/WWN:0x7473657468636574 <sscsMgr>	
Host-ID/WWN:0x216000C0FF804DE2 <server00>	
Host-ID/WWN:0x0000000000000000 <sscsMgr>	

14. Choose Yes to complete the host LUN filter entry.

< Main Menu > view and edit Logical drives view and edit logical Volumes view and edit Host luns		LUN	LU/LD	DRU	Partition	Size<MB>	RAID
		0	LD	0	0	68850	RAID5
v		1					
v	CHL 0 ID 40 <Primary Control						
v	CHL 1 ID 42 <Secondary Contr						
v	CHL 4 ID 50 <Primary Control						
s	CHL 5 ID 51 <Secondary Contr						
v	Edit Host-ID/WWN Name List						
		6					
		7					

Map	Logical Drive :	1
To	Partition :	0
	Channel :	0
	ID :	40
	Lun :	1 ?
Yes	No	

A mapped LUN displays a number. A filtered LUN displays an “M” for “masked LUN” in the LUN column.

< Main Menu >		LUN	LU/LD	DRU	Partition	Size<MB>	RAID
view and edit Logical drives							
view and edit logical Volumes							
view and edit host luns		M 0	LD	1	0	68952	RAID5
v							
v	CHL 0 ID 40 <Primary Control	1					
v	CHL 1 ID 42 <Secondary Contr						
v	CHL 4 ID 50 <Primary Control	2					
s	CHL 5 ID 51 <Secondary Contr						
v	Edit Host-ID/WWN Name List						
v							
		4					
		5					
		6					
		7					

Labeling a LUN (Solaris Operating System Only)

For the Solaris operating system to recognize a LUN, you must first manually write the label using the `Auto configure` option of the `format (1M)` command.

▼ To Label a LUN

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

```
# format
```

2. Specify the disk number when prompted.

1. Type `Y` at the following prompt, if it is deployed, and press `Return`:

```
Disk not labeled. Label it now? Y
```

The Solaris operating system’s `Format` menu is displayed.

2. Type `type` to select a drive type.

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

Choose the Auto configure menu option regardless of which drive types are displayed by the type option.

4. Type label and press Y when prompted to continue.

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

5. Type quit to finish using the Format menu.

Creating Solaris Operating System Device Files for Newly Mapped LUNs

Perform the following procedure to create device files for newly mapped LUNs on hosts in the Solaris 8 and Solaris 9 operating system.

For additional operating system information, see the Installation, Operation, and Service manual for your Sun StorEdge 3000 family array.

▼ To Create Device Files for Newly Mapped LUNs

1. To create device files, type:

```
# /usr/sbin/devfsadm -v
```

2. To display the new LUNs, type:

```
# format
```

3. If the format command does not recognize the newly mapped LUNs, perform a configuration reboot on the host:

```
# reboot -- -r
```

Saving Configuration (NVRAM) to a Disk

The controller configuration information is stored in non-volatile RAM (NVRAM). When you save it, the information is stored in the disk reserved space of all drives that have been configured into logical drives. Back up the controller configuration information whenever you change the array's configuration.

Saving NVRAM controller configuration to a file provides a backup of controller configuration information such as channel settings, host IDs, 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.

Note – A logical drive must exist for the controller to write NVRAM content onto it.

▼ To Save a Configuration to NVRAM

- **Choose “system Functions →Controller maintenance →Save nvrnm to disks,” and choose Yes to save the contents of NVRAM to disk.**

A prompt confirms that the NVRAM information has been successfully saved.

To restore the configuration, see “Restoring Your Configuration (NVRAM) From Disk” on page 284.

If you want to save and restore all configuration data, including LUN mapping information, use Sun StorEdge Configuration Service or the Sun StorEdge CLI 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 or to the *Sun StorEdge 3000 Family CLI User's Guide* for information about the `reset nvrnm` and `download controller-configuration` commands.

Logical Drives

This chapter explains how to create and manage logical drives using the “**view and edit Logical drives**” menu option.

Note – Procedures that have already been described in other chapters are cross-referenced instead of repeated.

Topics covered include:

- “Viewing the Logical Drive Status Table” on page 128
- “Viewing Physical Drives” on page 128
- “Creating Logical Drives” on page 129
- “Deleting a Logical Drive” on page 129
- “Partitioning a Logical Drive” on page 130
- “Deleting a Logical Drive Partition” on page 130
- “Changing a Logical Drive Name” on page 132
- “Rebuilding a Logical Drive” on page 132
- “Changing a Logical Drive Controller Assignment” on page 133
- “Expanding the Capacity of a Logical Drive” on page 133
- “Adding Physical Drives” on page 137
- “Performing a Parity Check” on page 139
- “Clearing a Fatal Fail Logical Drive Condition” on page 142
- “Copying and Replacing Drives With Drives of Larger Capacity” on page 143
- “Scanning Drives for Bad Blocks” on page 145
- “Shutting Down a Logical Drive” on page 147
- “Changing Write Policy for a Logical Drive” on page 148

Ch1	ID	Size<MB>	Speed	LG_DRU	Status	Vendor and Product ID
2<3>	8	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	9	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	10	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G

Creating Logical Drives

See “Creating Logical Drives” on page 59 for the procedure for creating logical drives on a SCSI array.

See “Creating Logical Drives” on page 97 for the procedure for creating logical drives on FC and SATA arrays.

See “Deleting Logical Drives” on page 51 for the procedure for deleting logical drives on SCSI arrays.

Deleting a Logical Drive

See “Deleting Logical Drives” on page 83 for the procedure for deleting logical drives on FC and SATA arrays.

See “Deleting Logical Drives” on page 51 for the procedure for deleting logical drives on a SCSI arrays.



Caution – This operation erases all data on the logical drive. Therefore, if any data exists on the logical drive, copy it to another location or back it up before it is deleted.

Partitioning a Logical Drive

See “Partitions” on page 72 for a description of logical drive partitions and a procedure for partitioning a logical drive on a SCSI array. See “Partitions” on page 109 for a description of logical drive partitions and a procedure for an FC or SATA array.

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. For example, if a logical drive contains partitions 0 through 3 and you delete partition 3, all of the capacity of partition 3 is added to partition 2.

▼ To Delete a Logical Drive Partition

1. From the Main Menu, choose “view and edit Logical drives.”
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 to delete this partition.

LG	ID	LU	RAID	Size<MB>	Partition	Offset<MB>	Size<MB>	NAME
P0	249B125B	NA	RAID1	34476	0	0	30000	
S1	183AC015	NA	RAID5	68952	1	30000	38952	
2			NONE		Partition Size <MB>: 0			
3			NONE					
4			NONE		4			
5			NONE		5			
6			NONE		6			
7			NONE		7			

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 6-1, the capacity of the deleted partition 1 (200 MB) is added into the partition above the deleted partition (partition 0), which combines 100 MB and 200 MB for a total of 300 MB.

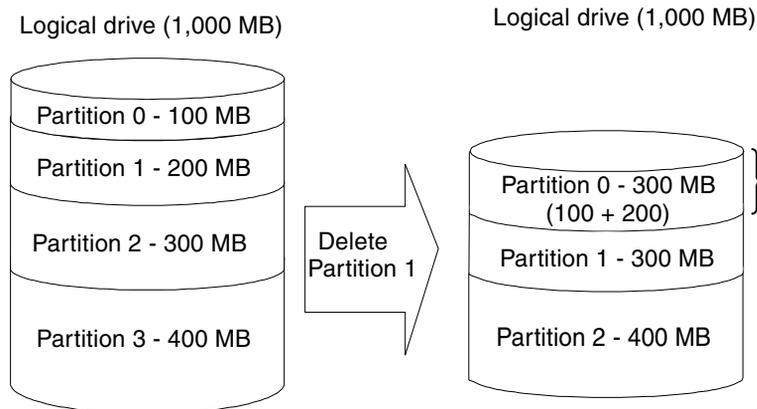


FIGURE 6-1 Example of Deleted Partitions



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.

Changing a Logical Drive Name

See “To Assign a Logical Drive Name (Optional)” on page 71 for the procedure for assigning a name to a logical drive on a SCSI array. See “To Assign a Logical Drive Name (Optional)” on page 109 for the procedure for assigning a name to a logical drive on a FC or SATA array.

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.

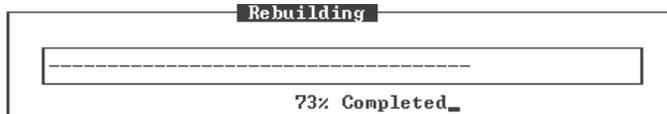
For more information about manually and automatically rebuilding logical drives, see “Rebuilding Logical Drives” on page 306.

▼ To Rebuild a Logical Drive

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the logical drive that has a failed member drive.
3. Choose “Rebuild logical drive,” and then choose Yes to rebuild the logical drive.

Note – The “Rebuild logical drive” option only displays when the selected logical drive (with RAID level 1, 3, or 5) includes a failed physical drive (Status DRV FAILED). NRAID and RAID 0 configurations provide no data redundancy.

The rebuilding progress is displayed on the screen.



When rebuilding has already started or the logical drive has been automatically rebuilt by a local spare drive or global spare drive, choose “Rebuild progress” to view the rebuilding progress. To cancel a rebuild operation, choose “Abort rebuild.”

A notification message informs you when the process is complete.

```
Rebuild of Logical Drive x Completed.
```

Changing a Logical Drive Controller Assignment

See “To Change a Controller Assignment (Optional)” on page 70 for the procedure for changing a logical drive controller assignment on the SCSI array. See “To Assign a Logical Drive Name (Optional)” on page 109 for the procedure for changing a logical drive controller assignment on the FC array or SATA array.

Expanding the Capacity of a Logical Drive

Use this option to expand the capacity of an existing logical drive. For example, you might originally have had several 18-Gbyte physical drives, of which only 9 Gbyte each were selected to build a logical drive. To use the remaining 9 Gbyte in each physical drive, you need to expand the logical drive. After a logical drive is expanded, the total additional capacity is displayed as a new partition. The new partition must be mapped to a host LUN in order for an HBA to recognize it. RAID levels 0, 1, 3, and 5 support expansion.

Note – To expand a logical volume, you must first expand the logical drives that make up the logical volume.

▼ To Expand a Logical Drive

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

3. Delete any local spare drives assigned to this logical drive.

It is not necessary to delete any global spare drives. See “Deleting a Spare Drive” on page 178 for more information about how to delete a local spare drive.

4. Choose “Expand logical drive.”

A menu displays the following options:

- Drive Expand Capacity

Choose this option to determine how much of the available capacity to add to the logical drive.

- Initialize Mode

Choose this option to expand the logical drive on-line or off-line.

- On-Line

Choose this option to use the logical drive before initialization is complete. Because the controller is building the logical drive while performing I/O operations, initializing a logical drive on-line requires more time than off-line initialization.

- Off-Line

Choose this option to use the drive only after initialization is complete. Because the controller is building the logical drive without having to also perform I/O operations, off-line initialization requires less time than on-line initialization.

5. Choose “Drive Expand Capacity.”

The maximum available drive free capacity information is displayed.

Note – The free capacity shown refers to the maximum available free capacity per physical drive, based on the smallest physical drive in the logical drive.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	1873F567	NA	RAID5	40000	GOOD				5	B	3	0	0	
				Drive Expand Capacity : 14476MB										
				Maximum Available Drive Free Capacity: 14476MB Maximum Drive Expand Capacity(MB) : 14476_										
3														
4														
5			NONE											
6			NONE											
7			NONE											

6. Press Return to expand the logical drive using all available capacity, or enter a value up to the maximum drive expand capacity.

The capacity shown in the Maximum Available Drive Free Capacity field is the maximum available free disk space per physical drive, based on the smallest physical drive in the logical drive. The capacity you specify is added to *each* physical drive in the logical drive.

As described in the following examples, the total amount of capacity that is added to the logical drive is automatically calculated based on the RAID level.

- RAID 0 – Multiply the amount entered in the Maximum Drive Expand Capacity field by the total number of physical drives contained in the logical drive. For example, 100 Mbyte \times 3 = 300 Mbyte total capacity added to the logical drive.
- RAID 1 – Multiply the amount entered in the Maximum Drive Expand Capacity field by the total number of physical drives contained in the logical drive, and then divide by two to account for mirroring. For example, 100 Mbyte \times 4 = 400 Mbyte. $400/2 = 200$ Mbyte total capacity added to the logical drive.
- RAID 3 and 5 – Multiply the amount entered in the Maximum Drive Expand Capacity field (n) by the total number of physical drives contained in the logical drive, and then subtract n to account for parity. For example, if $n = 100$, 100 Mbyte \times 3 = 300 Mbyte. $300 \text{ Mbyte} - 100 \text{ Mbyte} = 200 \text{ Mbyte}$ total capacity added to the logical drive.

If you know the total maximum drive capacity by which you want to expand a logical drive, perform the following calculations based on the RAID level to determine the amount to enter in the Maximum Drive Expand Capacity field:

- RAID 0 – Divide the total maximum drive capacity by the total number of physical drives contained in the logical drive. For example, if you want to add a total of 100 Mbyte to a logical drive that contains four physical drives, $100 \text{ Mbyte}/4 = 25 \text{ Mbyte}$ maximum drive expand capacity.
- RAID 1 – Divide the total number of physical drives contained in the logical drive by two to get n . Then divide the maximum drive capacity by n . For example, if you want to add a total of 100 Mbyte to a logical drive that contains four physical drives, $4/2 = 2$. $100/2 = 50 \text{ Mbyte}$ maximum drive expand capacity.
- RAID 3 and 5 – Subtract a single drive from the total number of physical drives to get n . Then divide the total maximum drive capacity by n . For example, if you want to add a total of 100 Mbyte to a logical drive that contains five physical drives, $5-1 = 4$. $100/4 = 50 \text{ Mbyte}$ maximum drive expand capacity.

Note – The Maximum Drive Expand Capacity cannot exceed the Maximum Available Drive Free Capacity.

7. (Optional) If you want to expand the logical drive offline instead of the default online mode, choose "Initialize mode," and then choose Yes to confirm your choice.

- **On-Line (default).** Enables you to configure and use the logical drive before initialization is complete. Because the controller is building the logical drive while performing I/O operations, initializing a logical drive on-line requires more time than off-line initialization.
- **Off-Line.** Enables you to configure and use the drive only after initialization is complete. Because the controller is building the logical drive without having to also perform I/O operations, off-line initialization requires less time than on-line initialization.

You can repeat Step 7 if you want to change logical drive expansion back on to on-line.

See “Creating Logical Drives” on page 59

8. When the logical drive capacity and initialize mode are selected, press Escape, and then choose Yes to expand the logical drive.

A notification message informs you when the process is complete:

On-Line Expansion of Logical Drive 0 Completed

9. Press Escape to return to the previous screen.

Drive capacity information is displayed. The total capacity of the logical drive has been expanded to 68952 Gbyte.

LG	ID	LU	RAID	Size<MB>	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	1873F567	NA	RAID5	68952	GOOD						5	3	0	0
1			NONE											
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

10. Map the new partition:

- For a SCSI array, map the new partition as described in “To Map a Logical Drive Partition” on page 75.
- For a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array, map the new partition as described in “To Map a Logical Drive Partition” on page 116.

11. (Solaris operating system only) For the Solaris operating system to recognize a LUN, you must first manually write the label using the `Auto configure` option of the `format (1M)` utility.

For the SCSI array, see “To Label a LUN” on page 78. For the FC array or SATA array, see “To Label a LUN” on page 123.

Adding Physical Drives

When you add a physical drive to a logical drive, the capacity of the original logical drive remains the same and additional capacity is displayed as a new partition. For example, if you have a single 200-Gbyte logical drive and add a 36-Gbyte drive, the total logical drive is 236 Gbyte 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 it.

Physical drives can be added to RAID 0, 1, 3, and 5 logical drives. For RAID 1 configurations, physical drives must be added in pairs.

The operation of adding physical drives cannot be canceled once it has started. If a power failure occurs, the add operation pauses. When the power comes back on, the controller automatically continues the operation.

If a physical drive in a logical drive fails during the add operation, the add operation pauses and resumes automatically after a logical drive rebuild has completed.

▼ To Add a Physical Drive to a Logical Drive

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

Note – For logical drives configured as RAID 1, you must add physical drives in pairs.

3. Choose “add Scsi drives” to display a list of available drives.

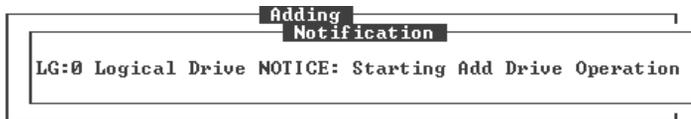
4. Select one or more drives to add to the target logical drive.

The physical drive you choose should have a capacity no less than the original physical drive. If possible, use drives of the same capacity. All drives in the array are treated as though they have the capacity of the smallest physical drive in the logical drive. The selected drive is indicated by an asterisk (*) mark.

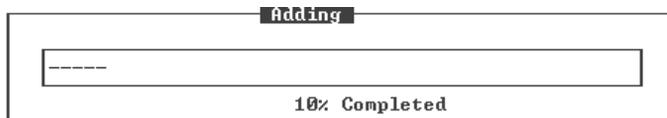
Ch1	ID	Size<MB>	Speed	LG_DRU	Status	Vendor and Product ID
*2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
2<3>	10	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G

5. Press Escape when all physical drives have been selected, and then choose Yes to add the drives.

A notification is displayed.



6. Press Escape to display a status bar that indicates the percentage of progress.



A notification message informs you when the process is complete.



When the add operation is completed, data is re-striped across the original and newly added drives.

7. Map the partition:

- For a SCSI array, map the new partition as described in "To Map a Logical Drive Partition" on page 75.
- For a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array, map the new partition as described in "To Map a Logical Drive Partition" on page 116.

8. (Solaris operating system only) For the Solaris operating system to recognize a LUN, you must first manually write the label using the `Auto configure` option of the `format (1M)` utility.

For the SCSI array, see “To Label a LUN” on page 78. For the FC array or SATA array, see “To Label a LUN” on page 123.

Performing a Parity Check

For RAID 3 and RAID 5 configurations, the primary purpose of the parity checking operation is to ensure that all sectors of the media can be successfully read and to provide alerts if a drive is starting to experience read or write errors.

A RAID 3 and RAID 5 parity checking operation recalculates the parity of data stripes in each of the logical drive’s RAID stripe sets and compares it with the stored parity. Depending on which options are enabled, if a discrepancy is found, an error is reported and the new correct parity is substituted for the stored parity. For RAID 1 configurations, if an inconsistency is encountered, data is copied from the master disk to the slave disk. If a bad block is encountered when parity is regenerated, data is copied from the other disk, master or slave, to the reporting disk drive reallocating the bad block.

You can check the integrity of your data before overwriting the parity drive. In this case you can manually regenerate the parity after you have reloaded the data from backup. See “Manually Regenerating the Parity of a Logical Drive” on page 139 for more information. Or you can enable the “**Overwrite Inconsistent Parity**” menu option, which causes the RAID controller to correct the data on the parity drive automatically whenever a parity check detects an inconsistency as explained in “Overwriting Inconsistent Parity” on page 140. To specify whether to report inconsistent parity errors as system events, see “Generating a Check Parity Error Event” on page 141.

Manually Regenerating the Parity of a Logical Drive

If you have disabled the “**Overwrite Inconsistent Parity**” menu option, the controller reports any inconsistency found in a parity check without overwriting the parity drive, which enables you to check your data and determine if it is intact or if the error occurred on a data drive. 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 Regenerate the Parity of a Logical Drive

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the logical drive that you want to regenerate the parity on.
3. Choose “reGenerate parity →Execute Regenerate Logical Drive Parity,” and then choose Yes to regenerate parity.
A notification is displayed that parity regeneration has begun.
4. Press Escape to view the progress indicator.

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

Overwriting Inconsistent Parity

Enabling the “**Overwrite Inconsistent Parity**” menu option causes the RAID controller to correct the data on the parity drive automatically whenever a parity check detects an inconsistency. In most cases, it is important to correct the data on the parity drive as soon as an inconsistency is detected to avoid the potential for data loss in the event of a drive failure.

However, you might prefer to check the integrity of your data before overwriting the parity drive. Disabling the “**Overwrite Inconsistent Parity**” menu option causes the controller to report any inconsistency found in a parity check without overwriting the parity drive. In this case, you can check your data and determine if it is intact or if the error occurred on a data drive. 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.



Caution – Restoring data by automatically overwriting the original data might cause data loss. To enable you to check your data and determine if it is intact or if the error occurred on a data drive, disable the “**Overwrite Inconsistent parity**” menu option.

▼ To Enable or Disable Overwrite Inconsistent Parity

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the logical drive whose automatic parity area overwriting you want to enable or disable.
The default value is Enabled.
3. To toggle between enabling and disabling this menu option, choose “reGenerate parity →Overwrite Inconsistent Parity -,” and then choose Yes to confirm the change.
4. Choose “Regenerate Logical Drive Parity,” and then choose Yes to regenerate parity.

Generating a Check Parity Error Event

When a parity check is performed, you can specify whether to report inconsistent parity errors as system events.

▼ To Enable or Disable Reporting Inconsistent Parity Errors as System Events

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the logical drive whose inconsistent parity you want to enable or disable as a system event.
The default value is Enabled.
3. To toggle between enabling and disabling this menu option, choose “reGenerate parity →Generate Check Parity Error Event -,” and then choose Yes to confirm the change.

Clearing a Fatal Fail Logical Drive Condition

With a RAID array system, your system is protected with the RAID parity drive and a global spare or spares. A FATAL FAIL condition occurs when more drives fail than your RAID redundancy can accommodate. The redundancy of your RAID array depends on your configuration. In a RAID 3 or RAID 5 configuration, two or more drives must fail for a FATAL FAIL status. In a RAID 1 configuration, you can lose multiple drives without fatal failure if all the failed drives reside on one side of a mirrored pair.

It is sometimes possible to recover the RAID array from a FATAL FAIL. For the detailed procedure to use, refer to the “Recovering From Fatal Drive Failure” section in the “Troubleshooting Your Array” chapter of the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for Sun StorEdge 3510 FC and 3511 SATA Arrays.

The Fatal Fail condition remains after the drive problem is fixed and the controller is reset. A “Clear state” menu option enables you to wait until you are sure the drive problem has been corrected before clearing the Fatal Fail logical drive condition and rebuilding the logical drive, if necessary.

1. **From the controller firmware’s Main Menu, choose “view and edit Logical drives.”**
2. **Select the logical drive with the FATAL FAIL status and press Enter.**
3. **Select “Clear state.”**
4. **Choose Yes when the “Back to degraded?” or “Back to normal?” prompt is displayed.**

After clearing the FATAL FAIL, the status changes to DRV FAILED.

Copying and Replacing Drives With Drives of Larger Capacity

For logical drives configured for RAID levels 0, 3, and 5, you can copy and replace existing physical drives with drives of the same or higher capacity. Because the logical drive uses the capacity size of its smallest size, all drives must be replaced with drives of the same or higher capacity. For example, as shown in FIGURE 6-2, a logical drive that originally contains three 36-Gbyte physical drives can be replaced with three new 73-Gbyte physical drives.

Note – To use the additional capacity provided by drives with higher capacity, you must expand the capacity as explained in “To Expand a Logical Drive” on page 133.

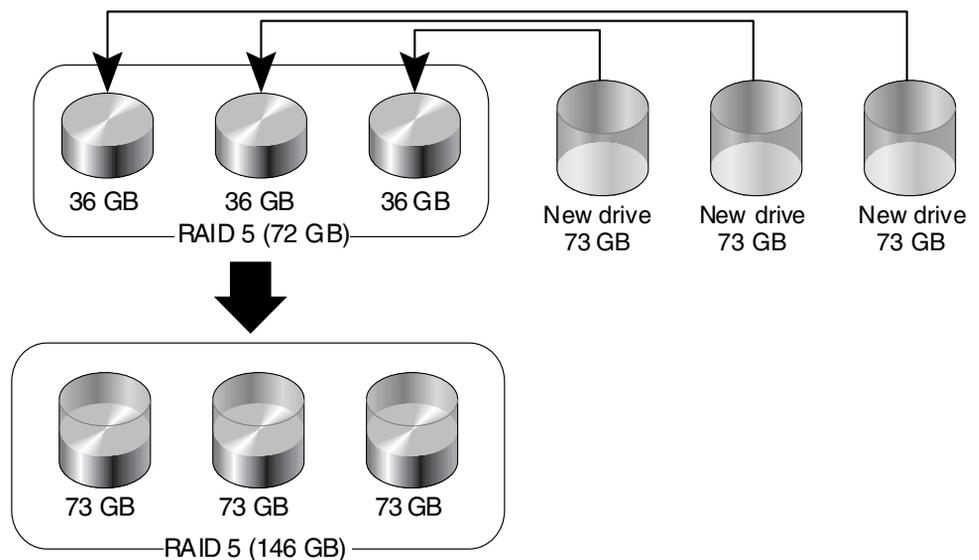


FIGURE 6-2 Expansion by Copy and Replace

Additional capacity is displayed as a new partition. The new partition must be mapped to a host LUN for the HBA to recognize it.

▼ To Copy and Replace a Drive

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

2. Select a target logical drive.

3. Choose “cOpY and replace drive.”

The physical drives belonging to the selected logical drive are listed.

4. Select the member drive (the source drive) that you want to replace with a larger drive.

A table of available physical drives is displayed.

5. Select a new drive onto which you will copy the content of the source drive.

The channel number and ID number of both the source drive and the destination drive are displayed in a confirmation message.

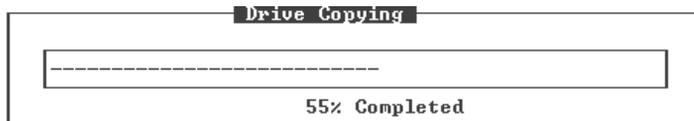
LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	1F7D914F	NA	RAID5	68952	GOOD					5	B	3	0	0
View scsi drives														
Del	Slot	Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID						
Par	2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G							
log	Source Drive: Channel=2 ID=6													
log	Destination Drive: Channel=2 ID=9													
Ex	Copy and Replace Drive ?													
ad	Yes No													
re	ONE	FRMT DRU	SEAGATE ST336753FSUN36G											
Me	0	ON-LINE	SEAGATE ST336753FSUN36G											
sh	0	ON-LINE	SEAGATE ST336753FSUN36G											
Wr	0	ON-LINE	SEAGATE ST336753FSUN36G											

Choose Yes to confirm.

6. A notification message is displayed.

```
LG:0 Logical Drive NOTICE:CHL:2 ID:6 Starting Clone
```

7. Press Escape to view the progress.



A notification message informs you when the process is complete.

```
LG:0 Logical Drive NOTICE:CHL:2 ID:6 Copy and Replace Completed
```

8. Repeat these steps to copy and replace all member drives with drives of higher capacity as needed.

You can now choose “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  
LG:0 Logical Drive NOTICE:CHL:2 ID:6 Copy and Replace Completed
```

Scanning Drives for Bad Blocks

The media scan feature sequentially checks each physical drive in a selected logical drive, block by block, for bad blocks. If a bad block is encountered, the controller rebuilds the data from the bad block onto a good block if one is available on the physical drive. If no good blocks are available on the physical drive, the controller designates the physical drive BAD and generates an event message. If a spare drive is available, the controller begins rebuilding data from the bad physical drive onto the spare.

You can specify whether media scanning begins automatically whenever the controller is powered up or reset. See “Media Scan at Power-Up” on page 246 for more information.

See “Using Media Scan on Individual Drives” on page 193 for more information about performing media scans on individual drives.

By default, assigned global spares are not scanned, nor are unassigned drives.

Whenever media scan is running on a drive, its front-panel LED flashes green.

You can change the priority of a particular media scan operation to specify the frequency of drive scanning.

▼ To Terminate a Media Scan

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select the logical drive that is currently being scanned.
3. Choose “Media Scan →Abort Media Scan,” and then choose Yes to terminate the media scan.

▼ To Perform a Media Scan

1. From the Main Menu, choose “view and edit Logical drives.”
2. Select a logical drive.
3. Choose “Media scan” to display a menu of media scan options.
4. (Optional) You can determine the priority of media scanning as related to other CPU tasks.

a. Choose “Media Scan Priority.”

The Media Scan Priority menu is displayed.

- Low
Media scan is not performed until other tasks have been completed.
- Normal.
Media scan is typically performed within three seconds.
- Improved.
Media scan is typically performed within one second.
- High.
Media scan is performed immediately.

b. Select a priority.

5. (Optional) Configure the media scan iteration count to specify whether the physical drives that make up the selected logical drive are checked one time or continuously, by choosing “Iteration Count -” and then choosing Yes to confirm the change.
6. When media scan is satisfactorily configured, press Escape, and then choose Yes to begin the media scan.
The front-panel LEDs for the drives being scanned will blink until the media scan is terminated.

Shutting Down a Logical Drive

LG:x NOTICE: CHL:x ID:x Starting Media Scan

Use the “**Shutdown logical drive**” menu option to:

- terminate I/O to a logical drive
- write data from cache to the logical drive
- place the logical drive in an offline state

▼ To Shut Down a Logical Drive

1. From the Main Menu, choose “view and edit Logical drive” to display a list of logical drives in the Logical Drive Status table.
2. Select the logical drive that you want to shut down.
3. Choose “sHutdown logical drive,” and then choose Yes to shut down the logical drive.

The Status column of the Logical Drive Status window changes to SHUTDOWN.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME	
P0	2EED5EB0	NA	RAID1	34476	GOOD					5	B	2	0	0	
S1	5EDFAB69	NA	RAID1	34476	SHUTDOWN					5	B	2	0	0	
2			NONE												
3			NONE												
4			NONE												
5			NONE												
6			NONE												
7			NONE												

4. Reset the array controller (“system Functions →Reset Controller”) to restore the logical drive to an online state (GOOD status).

Changing Write Policy for a Logical Drive

The global write policy for all logical drives is configured to write-back cache (default) or write-through cache as described in “Enabling and Disabling Write-Back Cache” on page 231. You can configure a different write policy for individual logical drives using the “**Write policy**” menu option. A logical drive’s write policy can be changed at any time.

▼ To Configure a Logical Drive’s Write Policy

1. From the Main Menu, choose “view and edit Logical drive” to display a list of logical drives in the Logical Drive Status table.
2. Select the logical drive that you want to configure.
3. Choose “Write policy -.”

The following write policy options are displayed:

- Default (default).

This menu option assigns the global write policy to the selected logical drive. If the global setting for write policy is changed, automatically changes the write policy for this logical drive.

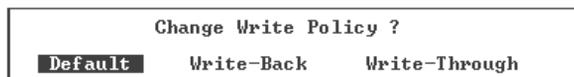
As described in “Cache Write Policy Guidelines” on page 32, the array can be configured to dynamically switch write policy from write-back cache to write-through cache if specified events occur. Write policy is only automatically switched for logical drives with write policy configured to Default. See “Event Trigger Operations” on page 267 for more information.

- Write-Back

This menu option assigns write-back cache regardless of any changes to the global write policy.

- Write-Through

This menu option assigns write-through cache regardless of any changes to the global write policy.



4. Choose a write policy option



The logical drive's write policy is changed.

Logical Volumes

This chapter explains how to creating and using logical volumes using the “**view and edit logical Volumes**” menu option.

While the ability to create and manage logical volumes remains a feature of Sun StorEdge 3000 family 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 and do not work in those configurations. In particular, the use of logical volumes is not supported in Sun Cluster environments.

Note – Avoid using logical volumes and use logical drives instead. For more information about logical drives, see Chapter 6.

Topics covered in this chapter include:

- “Understanding Logical Volumes (Multilevel RAID)” on page 152
 - “Logical Volume Limitations” on page 152
 - “Logical Drive and Logical Volume Partitions” on page 153
 - “RAID Expansion” on page 153
 - “Multilevel RAID Arrays” on page 154
 - “Spare Drives” on page 154
- “Viewing the Logical Volume Status Table” on page 155
- “Creating a Logical Volume” on page 155
- “Deleting a Logical Volume” on page 158
- “Expanding a Logical Volume” on page 158

Understanding Logical Volumes (Multilevel RAID)

A logical volume (LV) 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 physical drives according to its mirroring, parity, or striping schema.

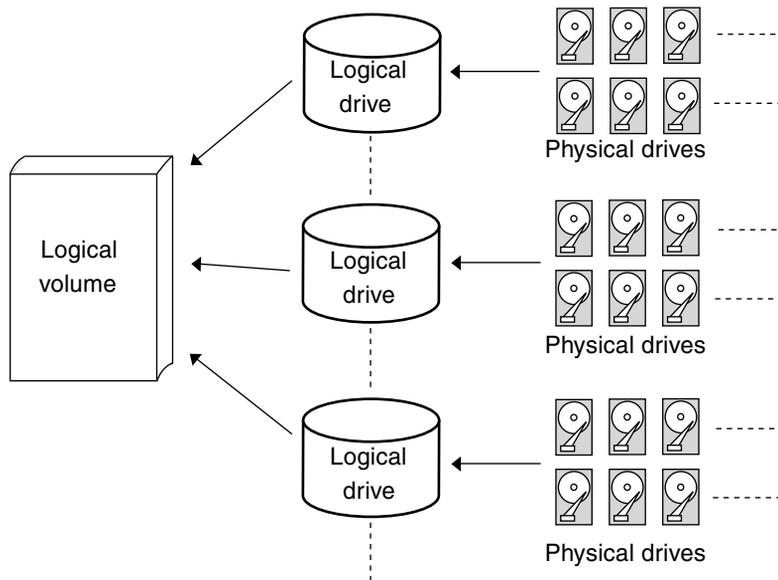


FIGURE 7-1 Logical Volume Composed of Multiple Drives

A logical volume can be divided into a maximum of 32 partitions. 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

- The use of logical volumes is not supported in Sun Cluster environments.
- A logical volume cannot be created using any logical drives that have a status of FATAL FAIL.

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.

Logical Drive and Logical Volume Partitions

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 all partitions in the logical drive until only one partition remains that includes 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 of a logical volume, that logical drive can no longer be partitioned using the “**View and Edit Logical Drives**” menu option. Instead, partition the logical volume 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 host computers to access 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 access the newly added capacity. To perform RAID expansion on a logical volume, first 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 a 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 50.** 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.

Viewing the Logical Volume Status Table

To check and configure logical drives, from the Main Menu choose “**view and edit logical Volumes**” to display the status of all logical volumes is displayed. The following table describes the categories of information shown in the status table.

TABLE 7-1 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 Mbyte
#LD	The number of logical drives in this logical volume

Creating a Logical Volume

A logical volume consists of one or several logical drives.

▼ To Create a Logical Volume

1. **From the Main Menu, choose “view and edit logical Volumes.”**

The current logical volume configuration and status are displayed.

If mixed drive support is enabled, a menu of drive types is displayed. If mixed drive support is disabled, proceed with Step 3.

See “Enabling Support for SATA Expansion Units Attached to FC Arrays” on page 87 for information about mixed drive support.

2. **If mixed drive support is enabled, select the type of logical drive you want to include in your logical volume.**



3. Select a logical volume number (0–7) that has not yet been defined, and then choose **Yes** to confirm your choice.

The following notice is displayed.

Only logical drives that have single partition and are not mapped to host luns will be listed.

4. Press **Escape** to clear the notice and display a list of logical drives that are available for inclusion in the logical volume you are creating.
5. Select one or more available logical drives from the list and, for each logical drive selected, press **Return** to tag it to be included in the volume.

An asterisk (*) is displayed in the LG field of the tagged drive.

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
*P0	2EED5EB0	NA	RAID1	34476	GOOD					5 B	2	0	0	
*P1	5EDFAB69	NA	RAID1	34476	GOOD					5 B	2	0	0	

Note – Logical drives must be assigned to either the primary controller or the secondary controller.

Note – Press **Return** again to deselect a tagged logical drive.

6. When all logical drives to be included in the logical volume have been selected, press **Escape** to display a menu of logical volume options.

Write Policy - Default
Logical Volume Assignment - Primary

- **Write Policy.** The RAID array write policy is preconfigured to enable write-back cache. Use the Write Policy menu option to configure a write policy for the logical volume that is different from the global write policy (see “Enabling and Disabling Write-Back Cache” on page 231).
 - **Logical Volume Assignment.** The logical volume will be assigned to the Primary Controller by default. Use the Logical Volume Assignment option to assign the logical volume to the Secondary controller.
7. (Optional) Assign a write policy for the logical volume that is different from the write policy assigned to the RAID array.

The write policy assigned to the logical volume is displayed in the list of logical volume options. You can change the logical volume’s write policy at any time.

Note – The Default write policy displayed is the global write policy.

a. Select “Write Policy -.”

The following write policy options are displayed.

- **Default (default)**

This menu option assigns the global write policy. If the global setting for write policy is changed, automatically changes the write policy for this logical volume.

As described in “Cache Write Policy Guidelines” on page 32, the array can be configured to dynamically switch write policy from write-back cache to write-through cache if specified events occur. Write policy is only automatically switched for logical volumes with write policy configured to Default. See “Event Trigger Operations” on page 267 for more information.

- **Write-Back**

This menu option assigns write-back cache regardless of any changes to the global write policy.

- **Write-Through**

Assigns write-through cache regardless of any changes to the global write policy.

b. Choose a write policy option.

8. (Optional) Change the logical volumes assignment from Primary controller to secondary controller by choosing “Logical Volume Assignment” and then choosing Yes to confirm the change.

9. Press Escape to display the configuration of the logical volume you are creating, and then choose Yes to create the logical volume.

Deleting a Logical Volume

This section describes how to delete a logical volume.

▼ To Delete a Logical Volume

1. **From the Main Menu, choose “view and edit logical Volumes.”**

The current logical volume configuration and status are displayed.

2. **Select a logical volume to delete.**

3. **Choose “Delete logical volume.”**

The following warning message is displayed.

```
This operation will result in the LOSS OF ALL DATA on the logical
Volume.
Delete Logical Volume ?
```

4. **Choose Yes to delete the logical volume.**

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.**

For more information, see “Expanding the Capacity of a Logical Drive” on page 133.

2. **Choose “view and edit logical Volumes” to display a list of logical volumes.**

3. **Select a logical volume to expand.**

4. **Choose “Expand logical volume,” and then choose Yes to expand the logical volume.**

Host LUNs

The “**view and edit Host luns**” menu allows you to map a partition, logical drive, or logical volume to a host channel. Every partition, logical drive, or logical volume can be mapped more than once to achieve a redundant data path, using multipathing software.

Note – Procedures that have already been described other chapters are cross-referenced instead of repeated.

Topics covered in this chapter include:

- “Mapping Logical Drive Partitions to Host LUNs” on page 160
- “Planning for 128 LUNs on a SCSI Array (Optional)” on page 162
- “Planning for 1024 LUNs on an FC or SATA Array (Optional, Loop Mode Only)” on page 163
- “Planning for 64 LUNs in Redundant FC or SATA Point-to-Point Configurations” on page 164
- “Mapping a Partition to a LUN” on page 165
- “Deleting a Host LUN Mapping” on page 165
- “Creating Host Filter Entries (FC and SATA Only)” on page 166
- “Determining a Host World Wide Name” on page 166
- “Manually Adding WWN Entries Using the Host-ID/WWN Name List” on page 168
- “Viewing and Modifying Host Filter Information” on page 168

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.

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 8-1 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 3310 SCSI array or Sun StorEdge 3320 SCSI array is 128. To create a total of 128 LUNs, see “Planning for 128 LUNs on a SCSI Array (Optional)” on page 162.

The maximum number of LUNs that can be created for a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array in loop configurations is 1024.

The maximum number of LUNs that can be created for a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array in redundant point-to-point configurations is 64. Refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array to see more information about point-to-point configurations.

Note – Some versions of operating system software or utilities might not display all mapped LUNs if there is no partition or logical drive mapped to LUN 0. Map a partition or logical drive to LUN 0 if you are in doubt, or refer to your operating system documentation.

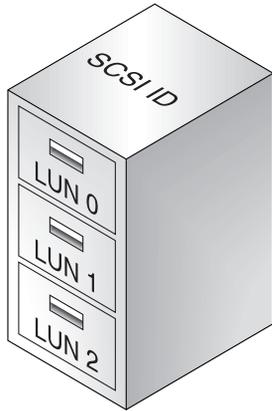


FIGURE 8-1 Filing Cabinet Represents the SCSI or FC ID

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

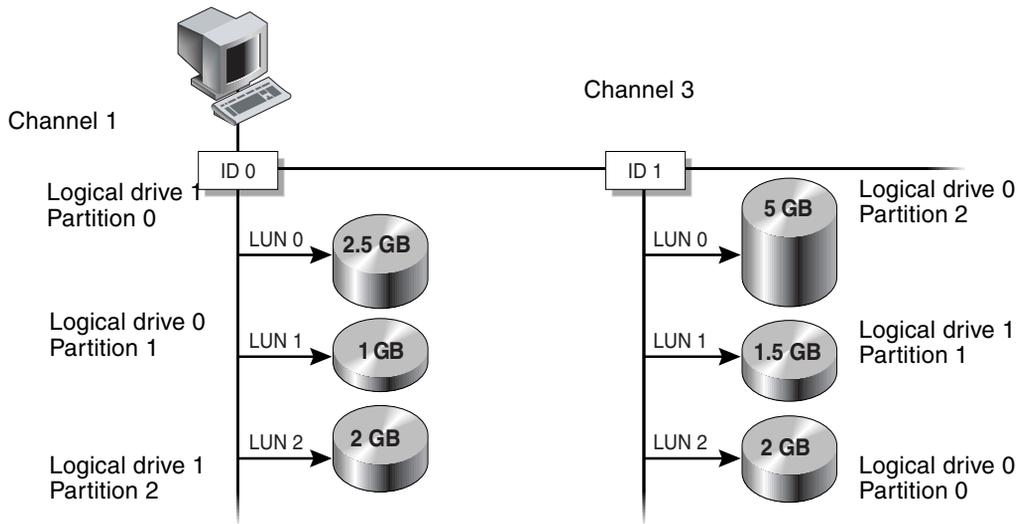


FIGURE 8-2 Mapping Partitions to Host ID/LUNs

Planning for 128 LUNs on a SCSI Array (Optional)

You can create up to 128 LUNs, which is the maximum number of storage partitions that can be mapped on a Sun StorEdge 3310 SCSI array or Sun StorEdge 3320 SCSI array. There are several ways you can meet this requirement. For example, you can 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.

or

- Create six host IDs (this requires three host drives), perform one of the following steps, and then map the 128 partitions to the 6 host IDs:
 - Create 4 logical drives of 32 partitions each.
 - Create 5 logical drives with the total number of partitions equaling 128 (four logical drives with 25 partitions each and one with 28 partitions).
 - Create 6 logical drives (5 logical drives with 21 partitions each and 1 with 23 partitions).

For details about how to add host IDs, see “To Add or Delete a Unique Host ID” on page 58.

Note – For an overview of how partitions, LUNs, and host IDs work, refer to “Mapping a Partition to a Host LUN” on page 74.

▼ To Create 128 LUNs

1. Create a minimum of four host IDs.

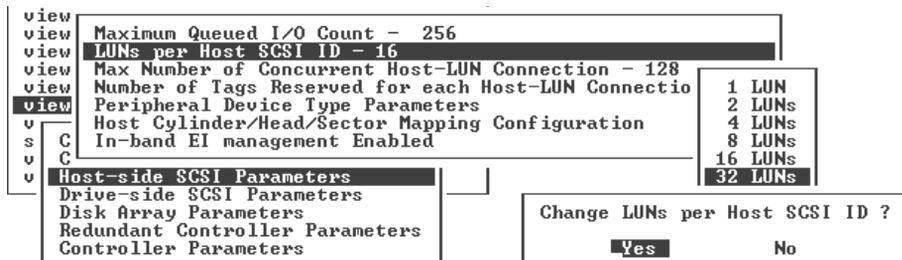
By default, you have two host IDs: Channel 1 ID 0 (primary controller) and Channel 3 ID 1 (secondary controller). You can have a total of two IDs per channel, one for the primary controller and one for the secondary controller.

See “To Add or Delete a Unique Host ID” on page 58 for more information.

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

3. Choose “view and edit Configuration parameters →Host-side Parameters.”

4. If the “LUNs per Host SCSI ID” setting is not 32, choose “LUNs per Host SCSI ID” and select 32. Then choose Yes to confirm.



5. Create at least four logical drives.
See “Creating Logical Drives” on page 59 for more information.
6. Create partitions on each logical drive until you reach a total of 128 partitions.
7. Map those partitions created in Step 6 to the host IDs created in Step 1.
See “Partitions” on page 72 and “Planning for 128 LUNs on a SCSI Array (Optional)” on page 162 for more information.

Planning for 1024 LUNs on an FC or SATA Array (Optional, Loop Mode Only)

If you want to create 1024 LUNs on a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array, which represent the maximum number of storage partitions that can be mapped for these arrays, you must map 32 IDs to the array’s host channels. There are several ways you can meet this requirement. For example, you can set up the following configuration described below.

▼ To Create 1024 LUNs

1. If necessary, edit “Host-side Parameters” so “LUNs Per Host SCSI ID” is 32.
For more information, see Step 1 of “Planning for 128 LUNs on a SCSI Array (Optional)” on page 162.

2. **Ensure that four default host channels (CH 0, 1, 4, and 5) are configured as host channels.**
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.**
See “To Add or Delete a Unique Host ID” on page 93.
4. **Create 32 logical drives.**
See “Creating Logical Drives” on page 97.
5. **Partition each logical drive into 32 partitions**
32 (logical drives) x 32 (partitions per logical drive) = 1024 (partitions).
6. **Link the 1024 partitions to the 32 host IDs.**
See “Partitions” on page 72 and “LUN Filtering (FC and SATA Only)” on page 114 for more information.

TABLE 8-1 Configuration for 1024 LUNs

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

Planning for 64 LUNs in Redundant FC or SATA 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, multipathing software is required on each host accessing the array.

Note – Multipathing for FC arrays is provided by Sun StorEdge Traffic Manager software. Refer to the release notes for your array 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 8-2 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 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

See “Mapping a Partition to a Host LUN” on page 74 for information about how to map a partition to a LUN on a SCSI array.

See “To Map a Logical Drive Partition” on page 116 for information about how to map a partition to a LUN on an FC or SATA array.

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. Select the channel and ID that are mapped to the host LUN.
3. Select the specific host LUN.

4. Choose Yes to delete the host LUN.

This option deletes the logical drive or logical volume mapping to a host channel. It 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 and SATA Only)

See “LUN Filtering (FC and SATA Only)” on page 114 for a description and procedure for creating host filter entries.

Determining a Host World Wide Name

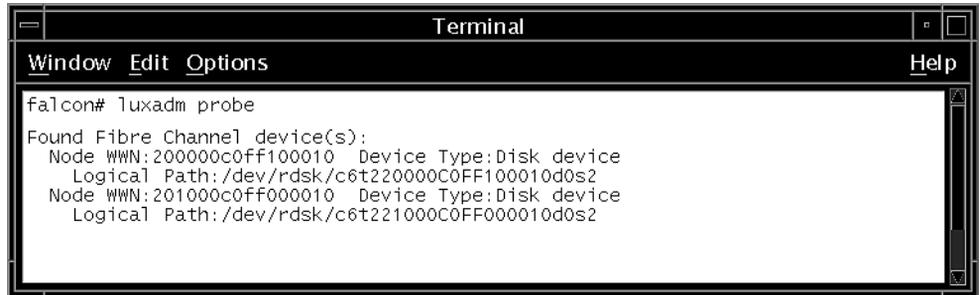
Prior to using LUN filtering, you need to identify which FC array is connected to which HBA card, and the WWN assigned to each card.

▼ To Determine a WWN for the Solaris Operating System

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.



```
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 a WWN for the Linux, Windows 2000, or Windows 2003 Operating System

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, such as Alt-Q or Ctrl-A. If the host has multiple HBA cards, select the card that is connected to the storage.

For more information about accessing your BIOS, refer to your system or HBA documentation.

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

Note – Refer to appendixes in the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your FC or SATA 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.

Manually Adding WWN Entries Using the Host-ID/WWN Name List

Use the “**view and edit Host luns** →**Edit Host-ID/WWN Name List**” menu option to manually add to the current device list any HBA worldwide names that have not been automatically propagated to the array. You can add up to 64 WWNs per array.

Note that when you add a worldwide name using this menu option, that worldwide name appears to be available to all channels. Ensure that, when you create a filter for a worldwide name that has been added using the “**Edit Host-ID/WWN Name List**” menu option, you also create the filter on a channel where the HBA is connected.

Note – You can also manually add worldwide names using the “**Manual add host filter entry**” menu option when you are configuring a host filter. When you use the “**Manual add host filter entry**” option, the worldwide name you add is only displayed in the worldwide name list when you create a filter on the channel where that worldwide name was entered.

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 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 “LUN Filtering (FC and SATA Only)” on page 114 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.

Note – You can add up to 64 WWNs for an array.

See “LUN Filtering (FC and SATA Only)” on page 114 for the steps to follow once you choose this menu option.

Physical Drives

This chapter provides information about viewing and editing physical drive parameters, assigning spare drives, recognizing good and bad drives, identifying failing drives using SMART detection and media scanning, and cloning failing drives.

Topics covered include:

- “Viewing the Status of a Physical Drive” on page 172
- “SCSI Drive IDs (SCSI Only)” on page 173
- “FC Drive IDs (FC and SATA Only)” on page 175
- “Viewing Physical Drive Information” on page 176
- “Assigning a Local Spare Drive” on page 177
- “Assigning a Global Spare” on page 178
- “Deleting a Spare Drive” on page 178
- “Scanning Drives (SCSI Only)” on page 179
- “Adding or Deleting Drive Entries (SCSI Only)” on page 179
- “Identifying a Failed Drive for Replacement” on page 180
 - “Flashing a Selected Physical Drive” on page 181
 - “Flashing All SCSI Drives” on page 182
 - “Flashing All But Selected Drive” on page 183
- “Fault Protection Measures” on page 183
 - “Cloning a Failing Drive” on page 184
 - “Terminating a Perpetual Clone” on page 188
 - “Viewing the Status of a Cloning Operation” on page 189
 - “Using SMART Functionality” on page 190
- “Using Media Scan on Individual Drives” on page 193
- “SCSI Drive Utilities (Reserved)” on page 195
 - “SCSI Drive Low-Level Format Utilities” on page 195
 - “Read/Write Test” on page 196
- “Changing Disk Reserved Space” on page 197

Viewing the Status of a Physical Drive

The Physical Drive Status table shows you the status of all physical drives in your array.

▼ To View the Physical Drive Status Table

1. From the Main Menu choose “view and edit Drives” to view your array’s physical drives, and to edit physical drive parameters.

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

Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
2<3>	6	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	7	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	8	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	9	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	10	34732	200MB	GLOBAL	STAND-BY	SEAGATE ST336753FSUN36G
2<3>	12				SES	SUN StorEdge 3510F A

2. Select a drive from the physical drive table if you want to modify its configuration information or view additional information.

A menu of available options is then displayed, as described in “Deleting a Spare Drive” on page 178.

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

Unlike Sun StorEdge 3510 FC arrays or Sun StorEdge 3511 SATA arrays, when a Sun StorEdge 3310 SCSI array or Sun StorEdge 3320 SCSI array is powered up, the controller scans all physical drives that are connected through the drive channels. If you install a physical drive after the SCSI controller completes initialization, select the drive and then choose the “Scan scsi drive” menu option to force the controller to recognize the newly added drive. Then you can configure it as a member of a logical drive.

Note – Newly added drives are automatically scanned on Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA arrays.

When a Sun StorEdge 3310 SCSI array or Sun StorEdge 3320 SCSI array is powered up, it can take approximately 30–40 seconds before the drive speed is displayed correctly. Before that happens, the drive speed can display as ASYNC.

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 with 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 because the drive still has reserved space data from a logical drive.

When the logical drive 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 256 Mbyte 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 Drives**” menu, the drive status changes to NEW.

To handle BAD drives, see “Scanning Drives for Bad Blocks” on page 145. If two drives show BAD and MISSING status, refer to the Troubleshooting chapter of the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array.

For detailed information about the Physical Drive Status Table, see “Physical Drive Status Table” on page 296.

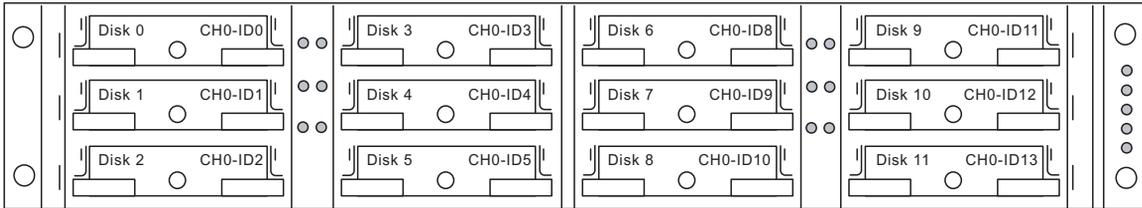
SCSI Drive IDs (SCSI Only)

Each SCSI array must be configured as a single-bus configuration or a split-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 SCSI array.

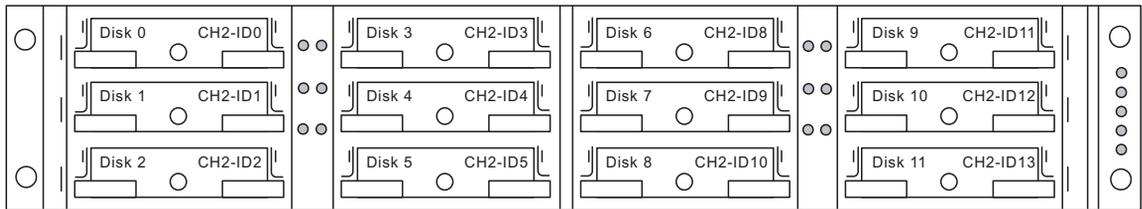
The drive-bus configuration determines how drives and drive IDs are assigned to drive channels on the controller.

- 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

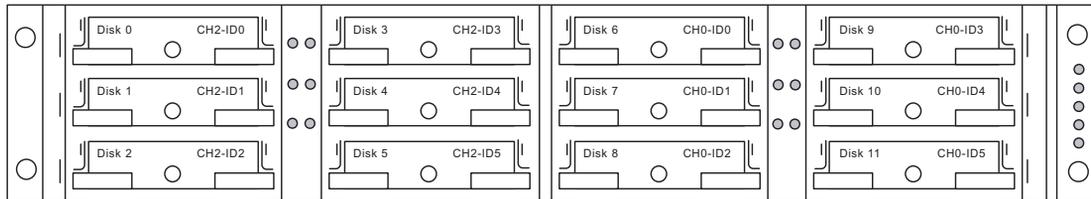


Expansion unit - single-bus configuration - default IDs

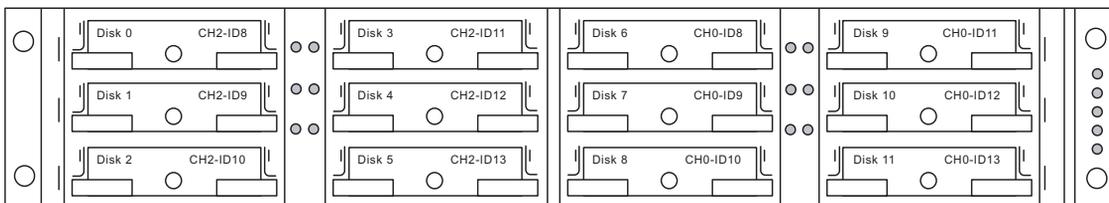


- A split-bus configuration assigns six disk drive IDs to CH 0 and six disk drive IDs to CH 2 in the RAID array, and then typically adds an additional six disk drive IDs to both CH 0 and CH2 when it is connected to an expansion unit.

RAID array - split-bus configuration - default IDs



Expansion unit - split-bus configuration - default IDs



FC Drive IDs (FC and SATA Only)

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

Use the ID switch on the left front side of the expansion unit to set the loop IDs for the disk drives to different ranges of ID values so that the same disk IDs are not repeated on the same Fibre loop.

Caution – An ID conflict can occur if the ID is changed while the expansion unit is in use or if it is not power-cycled after the setting has been changed. Change the switch ID only when the expansion unit is powered off or not in use. Power the expansion unit on after the switch setting is changed (or power-cycle it if it is already powered up).

By default, the ID switch on all RAID arrays is set to 0, so that the range of IDs is 0 to 11 for 12 drives (IDs 12–15 are not used for drives). By default, the ID switch on all RAID expansion units is set to 1.

The ID switch offers eight ID ranges. Each set contains 16 IDs (the last four IDs in each range are not used for drives), except for the last set, which contains 15 IDs, the last three of which are not used).

These ranges are shown in TABLE 9-1. Refer to the Installation, Operation and Service manual for your array to see the procedure for configuring the loop ID on your array.

TABLE 9-1 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

Viewing Physical Drive Information

The drive revision number, serial number, and disk capacity of each individual drive can be displayed. If you select a SCSI Enclosure Services (SES) chip rather than a drive, the “**View drive information**” menu option also displays the worldwide node name (WWN). From this information you can determine the serial number of the chassis, as described in the following procedure.

▼ To View Physical Drive Information

1. From the Main Menu, choose “view and edit Drives.”
2. Select the physical drive whose information you want to view.
3. Choose “View drive information” to display the drive’s revision number, serial number, and disk capacity in blocks of 512 Kbyte.

Slot	Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
	2<3>	6	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
			View drive information		0	ON-LINE	SEAGATE ST336753FSUN36G
			Revision Number		0349		36753FSUN36G
			Serial Number		3HX0YEJT00007349		36753FSUN36G
			Disk Capacity (blocks)		71132958		36753FSUN36G
			Node Name(WWNN)		20 00 00 0C 50 33 2B FD		36753FSUN36G
			Redundant Loop ID		6		36753FSUN36G
	2<3>	12				SES SUN	StorEdge 3510F A

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

Slot	Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
	2<3>	6	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
			Revision Number		1040		36753FSUN36G
			Serial Number		N/A		36753FSUN36G
			Disk Capacity (blocks)		20 40 00 C0 FF 00 2F 18		36753FSUN36G
			Node Name(WWNN)		12		36753FSUN36G
			Redundant Loop ID				36753FSUN36G
			View drive information		NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	2<3>	12				SES SUN	StorEdge 3510F A

The world wide name is shown in hexadecimal format; for example:

```
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.

Assigning a Local Spare Drive

A local spare drive is a standby drive assigned to serve a single 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 the logical drive.

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

▼ To Assign a Local Spare Drive

1. From the Main Menu, choose “view and edit Drives.”
2. Select the drive you want to designate as a spare.
3. Choose “add Local spare drive.”
4. Select the logical drive to which you want to assign the local spare, and then choose Yes to assign the local spare to that logical drive.

Assigning a Global Spare

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

▼ To Assign a Global Spare

1. From the Main Menu, choose “view and edit Drives.”
2. Select the drive you want to designate as a spare.
3. Choose “add Global spare drive,” and then choose Yes to assign the global spare.

Deleting a Spare Drive

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

▼ To Delete a Spare Drive

1. From the Main Menu, choose “view and edit Drives.”
2. Select the local spare drive or global spare drive you want to delete.
3. Choose “Delete global/local spare drive,” and then choose Yes to delete the spare drive.

The status of the spare drive you deleted, or any drive you replaced from a logical unit, is now indicated as FRMT DRV. This drive can now be assigned to another logical device.

Scanning Drives (SCSI Only)

When a SCSI array is powered on, the controller scans all physical drives that are connected through drive channels.

Unlike Fibre Channel and SATA arrays, if a SCSI array has completed initialization and then a physical drive is connected, the controller does not recognize the new drive until the next controller reset. This difference in behavior is due to differences between Fibre Channel and SCSI architectures and protocols.

Use the “**Scan scsi drive**” menu option to force the controller to scan a drive that has been added to a SCSI array.

Note – Neither the “**Periodic Auto-Detect Failure Drive Swap Check Time**” menu option nor the “**Periodic Drive Check Time**” menu option force the scanning of a SCSI drive.

▼ To Scan a New SCSI Drive

1. From the Main Menu, choose “view and edit Drives.”
2. Select the drive you wish to scan.
3. Choose “Scan scsi drive” to display the available SCSI channels.
4. Select a SCSI channel to display the available drive IDs.
5. Select the ID of the drive you want to scan, and then choose Yes to scan the drive.

Adding or Deleting Drive Entries (SCSI Only)

Use the “**add drive Entry**” menu option to add an additional record to the drive table of a SCSI array 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.

▼ To Add a Drive Entry

1. From the Main Menu, choose “view and edit Drives.”
2. Select the drive to which you want to add a drive entry.
3. Choose “add drive Entry” to display a list of available channels.
4. Select a channel to display a list of available IDs on that channel.
5. Select an ID, and then choose Yes to create a drive entry.
The drive entry is created. Its status is displayed as ABSENT.

▼ To Remove an Empty Drive Entry

1. From the Main Menu, choose “view and edit Drives.”
2. Select an empty SCSI drive whose status is ABSENT.
3. Choose “Clear drive status.”
The empty drive entry is removed from the SCSI drive table.

Identifying a Failed Drive for Replacement

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



Caution – If a drive in a logical drive fails, it is important to remove the failed drive. If you mistakenly remove the wrong drive from the same logical drive, you will have failed a second drive and may cause a critical failure of the logical drive.

To locate a failed drive, identify a single drive, or test all drive activity LEDs, you can flash the LEDs of any or all drives in an array. Since a defective drive does not flash, this provides a good way for you to visually identify a failed drive before replacing it.

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

▼ To Identify a Drive

1. **Terminate media scan for the logical drive that contains the drive you wish to identify, or terminate media scan for all logical drives.**

Terminating media scan prevents the drive LEDs from flashing green, which otherwise would make it difficult to identify the drive. See “To Terminate a Media Scan” on page 146 for more information.

2. **From the Main Menu, choose “view and edit Drives.”**
3. **Select the drive you want to identify.**
4. **Choose “Identify scsi drive →flash All drives” to flash the activity LEDs of all of the drives in the drive channel.**
5. **(Optional) Delete the current Flash Drive Time and type in a new Flash Drive Time.**
6. **Press Return, and then choose Yes to confirm.**

The read/write LEDs flash for all drives except failed hard drives. The absence of a flashing 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 a Selected Physical Drive

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

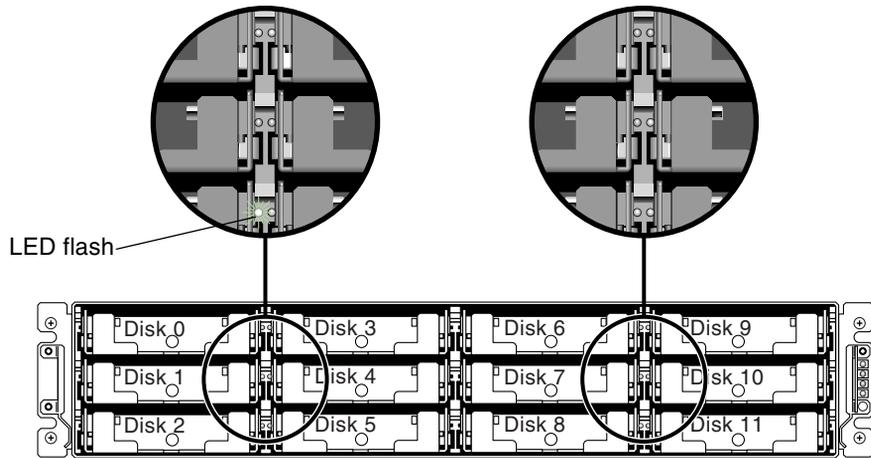


FIGURE 9-1 Flashing the Drive LED of a Selected Drive

Flashing All SCSI Drives

The “Flash All SCSI Drives” menu option flashes the LEDs of all good drives but does not flash the LEDs of any defective drives.

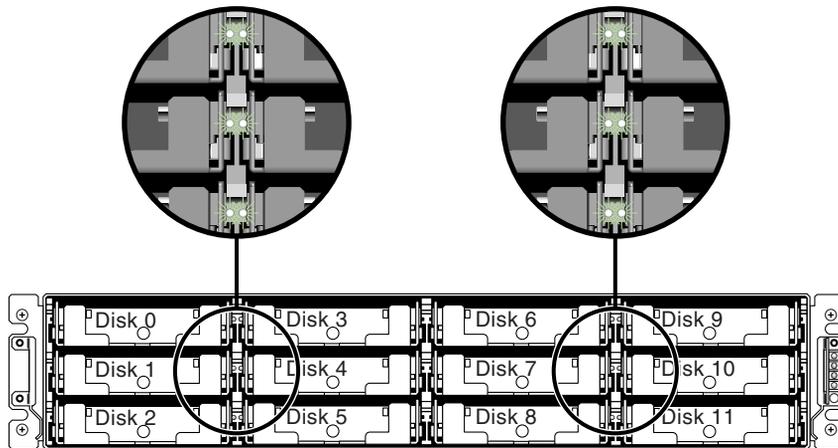


FIGURE 9-2 Flashing All Drive LEDs to Detect a Defective Non-Flashing Drive

Flashing All But Selected Drive

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

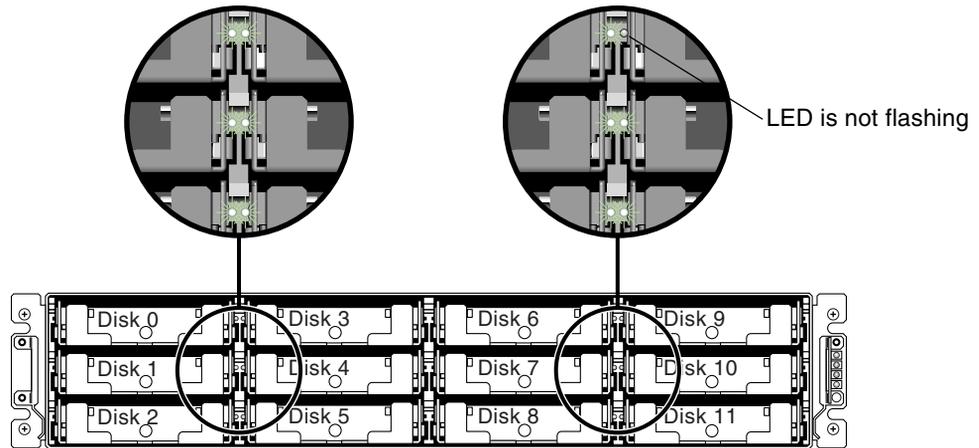


FIGURE 9-3 Flashing All Drive LEDs Except a Selected Drive LED

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 with media scan is one common predictor of a drive that is about to fail. For more information about media scan, see:

- “Scanning Drives for Bad Blocks” on page 145
- “Using Media Scan on Individual Drives” on page 193

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

This section covers the following topics:

- “Cloning a Failing Drive” on page 184
- “Terminating a Perpetual Clone” on page 188
- “Viewing the Status of a Cloning Operation” on page 189
- “Using SMART Functionality” on page 190

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 NRAID logical drives; its use is not supported in those 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 drive 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 no standby drive (local or global spare drive) is available, you must add a new drive and configure it as a local or global spare drive. If no standby drive is available, the “**clone Failing drive**” option is not displayed.

▼ To Replace After Clone

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

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

4. Choose "Replace After Clone," and then choose Yes to clone the drive.

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). A notification message is displayed:

LG:0 Logical Drive NOTICE: CHL:2 ID:6 Starting Clone

5. Press Escape to clear the message and display a progress bar.

Drive Copying										
S1	-----								uct ID	
	18% Completed_								3FSUN36G	
								3FSUN36G		
	2<3>	8	34732	200MB	0	ON-LINE	SEAGATE	ST336753FSUN36G		
	2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE	ST336753FSUN36G		
	2<3>	10	34732	200MB	GLOBAL	STAND-BY	SEAGATE	ST336753FSUN36G		
	2<3>	12				SES	SUN	StorEdge 3510F A		

6. (Optional) To close the progress bar, press Escape to return to the table of SCSI drives.

If you close the progress bar and want to return to it so you can view the cloning progress or abort the drive clone operation, you can perform the following steps:

a. Select the drive indicated as CLONING.

Slot	Chl	ID	Size<MB>	Speed	LG_DRU	Status	Vendor and Product ID
	2<3>	6	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
	2<3>	7	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
					0	ON-LINE	SEAGATE ST336753FSUN36G
					NONE	FRMT DRU	SEAGATE ST336753FSUN36G
						CLONING	SEAGATE ST336753FSUN36G
						SES	SUN StorEdge 3510F A

b. Choose "clone Failing drive" 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.

```
LG:0 Logical Drive NOTICE: CHL:2 ID:6 Copy and Replace
Completed
```

7. **Press Escape to clear the message and display the SCSI drives' status after the cloning process.**

Perpetual Clone

Data on the source drive (the drive with a predicted error or any selected member drive) is cloned to the standby spare, but the spare does not become the new source drive. The standby spare drive clones the source drive without substituting it.

The status of the spare drive is displayed as a CLONE drive as soon as the cloning process is complete. The source drive remains a member of the logical drive.

▼ To Enable Perpetual Clone

1. **From the Main Menu, choose "view and edit Drives."**
2. **Select the member drive with a predicted error.**
3. **Choose "clone Failing drive →Perpetual Clone," and then choose Yes to clone the drive.**

The controller automatically starts the cloning process, using the existing standby (local or global spare drive) to clone the source drive.

Note – If no standby drive (local or global spare drive) is available, you must add a new drive and configure it as a global spare or local spare drive.

A notification message is displayed when the cloning process begins:

```
LG:0 Logical Drive NOTICE: CHL:2 ID:10 Starting Clone
```

4. Press Escape to clear the message and display a progress bar.

Drive Copying								
S1	-----						uct ID	
	18% Completed						3FSUN36G	
	2<3>	8	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G	
	2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G	
	2<3>	10	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G	
	2<3>	12				SES	SUN StorEdge 3510F A	

5. (Optional) To close the progress bar, press Escape to return to the table of SCSI drives.

If you close the progress bar and want to return to it so you can view the cloning progress or abort the drive clone operation, you can perform the following steps:

- a. Select the drive indicated as CLONING.

Slot	Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
	2<3>	6	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
	2<3>	7	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
	View drive information				0	ON-LINE	SEAGATE ST336753FSUN36G
	Identify scsi drive				NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	clone failing drive						
	Source Drive: Channel 2 ID 6				CLONING		SEAGATE ST336753FSUN36G
	View clone progress						
	Abort clone				SES	SUN	StorEdge 3510F A

- b. Choose “clone Failing drive” to view the current status.

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

A notification message informs you when the process is complete.

LG:0 Logical Drive NOTICE: CHL:2 ID:10 Clone Completed
--

- Press **Escape** to clear the notification message and display the SCSI drives' status after the cloning process.

The source drive (Channel 2 ID 10) remains as a member of logical drive 0, and the standby drive (Channel 2 ID 6, the local or global spare drive) becomes a CLONE drive.

Slot	Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
	2<3>	6	34732	200MB	0	CLONE	SEAGATE ST336753FSUN36G
	2<3>	7	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
	2<3>	8	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
	2<3>	9	34732	200MB	NONE	FRMT DRU	SEAGATE ST336753FSUN36G
	2<3>	10	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
	2<3>	12				SES	SUN StorEdge 3510F A

Terminating a Perpetual Clone

Once data from a failing drive has been cloned onto a spare drive during a perpetual clone operation, the failing drive remains part of the logical drive and the spare drive remains a clone drive until the perpetual clone is manually terminated.

▼ To Terminate a Perpetual Clone

- From the Main Menu, choose "view and edit Drives."
- Select the drive whose status is CLONING.
- Choose "clone Failing drive" to identify the source drive that is being cloned.

The channel and ID of the source drive are displayed, along with the following options:

- Replace original with clone

Choose this option to take the failing (source) drive offline and replace it in the logical drive with the spare drive on which data from the source drive has been reconstructed.

- Delete clone

Choose this option to terminate the clone, leaving the source drive as part of the logical drive and reestablishing the clone drive as part of the logical drive.

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

See “Terminating a Perpetual Clone” on page 188 for instructions on how to disable perpetual cloning.

Using SMART Functionality

SMART is an industry-standard technology that provides near-term failure prediction for disk drives. When SMART is enabled, as it is in the Sun StorEdge 3000 family RAID controllers, 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 from 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

The Sun StorEdge 3000 family arrays implement the ANSI-SCSI X3T10/94-190 standard. The Detect and Clone+Replace menu option is the default setting.

Note – SMART functionality is not supported on the SATA drives used in Sun StorEdge 3511 RAID controllers or expansion units.

Procedures for the SMART prediction feature are:

- “To Enable and Use SMART Detection” on page 191
- “To Test a Drive for SMART Functionality” on page 192
- “To Disable SMART Detection” on page 193

▼ To Enable and Use SMART Detection

1. Choose “view and edit Configuration parameters →Drive-side Parameters →Periodic Drive Check Time” to display a list of time intervals.
2. Select a time interval, and then choose Yes to confirm your choice.
3. From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Drive Predictable Failure Mode (SMART).”
4. Keep the default “Detect and Clone+Replace” setting, or choose another SMART menu option and choose Yes to confirm your choice.

■ Detect and Clone + Replace

This is the default setting which enables the SMART monitoring.

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 immediately clones the drive whose failure has been predicted to a standby 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. To replace a drive, refer to the *Sun StorEdge 3000 Family FRU Installation Guide*.

Note – To see the progress of cloning, press Escape to clear the notification message and see the status bar.

■ Detect and 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 for the logical drive that contains the predicted failure. If the clone drive was a global spare, it now functions as a local spare and has lost its global spare references.

If the drive whose failure has been predicted does fail subsequently, the clone drive takes over immediately. To view the status and drive identities or cancel the cloning process, see “To View the Status of a Cloning Operation” on page 189.

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 starts to rebuild the failed drive immediately. This helps prevent a fatal drive error if yet another drive fails.

However, if the drive whose failure has been predicted continues to work successfully and another drive in a *different logical drive* fails, the clone drive does not act as the standby spare for that logical drive. This happens because once a global spare becomes associated to a logical drive as a perpetual clone, it replaces its global references with local references.

5. Choose Yes to confirm your choice.

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

6. Assign at least one spare drive to the logical drive (either a local spare or global spare).

See “To Assign a Local Spare Drive” on page 177 or “To Assign a Global Spare” on page 178.

▼ To Test a Drive for SMART Functionality

1. Choose “view and edit Configuration parameters → Drive-side Parameters → Periodic Drive Check Time” to display a list of time intervals.

2. Select a time interval, and then choose Yes to confirm your choice.

3. From the Main Menu, choose “View and edit Drives.”

4. Select a drive to test that is an active part of a logical drive.

A “Predictable Failure Test” menu option is displayed in the SCSI drive menu.

Note – If the SMART feature is not properly enabled, this menu option is not displayed.

5. Choose “Predictable Failure Test,” and then choose Yes to begin the test.

The drive simulates a predictable drive error.

The next time the controller performs the periodic drive check, the controller detects the error simulated by the selected drive and displays an error message:

```
SMART-CH:2 ID:6 Predictable Failure Detected<TEST>
```

The “<TEST>” component of the message indicates that no predictable failure was actually detected and that no action is necessary.

▼ To Disable SMART Detection

1. From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Drive Predictable Failure Mode (SMART) →Disabled.”
2. Choose Yes to confirm your choice.

Using Media Scan on Individual Drives

The media scan feature sequentially checks a physical drive in a logical drive, block by block, for bad blocks. If a bad block is encountered, the controller rebuilds the data from the bad block onto a good block if one is available on the physical drive. If no good blocks are available on the physical drive, the controller designates the physical drive BAD, generates an event message, and if a spare drive is available, will begin rebuilding data from the bad physical drive onto the spare.

If a spare drive is not immediately available, you can add a physical drive, assign it as a global spare, and then manually clone the failing drive to the spare. See:

- “Assigning a Global Spare” on page 178
- “Cloning a Failing Drive” on page 184.

In addition to performing media scan on individual drives, you can select a logical drive and change media scan settings for all physical drives in that logical drive. See “Scanning Drives for Bad Blocks” on page 145 for more information.

▼ To Perform a Media Scan

You can perform media scans on an individual physical drive assigned to a logical drive.

1. From the Main Menu, choose “view and edit Drives.”
2. Select a physical drive that is part of a logical drive.

3. Choose “mediA scan” to display a menu of media scan options.
4. (Optional) You can determine the priority of media scan as related to other CPU tasks.
 - a. Choose “Media scan priority -.”

The Media Scan Priority menu is displayed.

 - Low
Media scan is not performed until other tasks have been completed.
 - Normal.
Media scan is typically performed within three seconds.
 - Improved.
Media scan is typically performed within one second.
 - High.
Media scan is performed immediately.
 - b. Select a priority.
5. (Optional) Configure the media scan iteration count to specify whether the physical drive is checked one time or continuously, by choosing “Iteration Count -,” and then choosing Yes to confirm the change.
6. When media scan is satisfactorily configured, press Escape, and then choose Yes to begin the media scan.

A notification is displayed.

LG:x NOTICE: CHL:x ID:x Starting Media Scan

▼ To Terminate a Media Scan

You can select an individual logical drive and abort the media scan of a specific physical drive in that logical drive.

1. From the Main Menu, choose “view and edit Drives.”
2. Select a physical drive that is currently being scanned.
3. Choose “mediA scan →Abort Media scan,” and then choose Yes to terminate the media scan.

SCSI Drive Utilities (Reserved)

Do not use the “**scsi drive Utilities**” menu options unless you are directed to use them by technical support.

Note – This menu option is only displayed if disk reserved space has been deleted (see “Changing Disk Reserved Space” on page 197).

SCSI Drive Low-Level Format Utilities

Use these menu options only when a disk has been taken out of service and is unusable unless it is reformatted.



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

The low-level disk format utility cannot be used on a spare drive (local or global) or a member drive of a logical drive. “**Disk Reserved space**” must be removed before this menu option is available. See “Changing Disk Reserved Space” on page 197 for more information.

▼ To Low-Level Format a Physical Drive

1. From the Main Menu, choose “view and edit Drives.”
2. Select a drive you want to format.
3. Choose “scsi drive Utilities →SCSI Drive Low-level Format,” and then choose Yes to begin formatting the drive and to display the formatting progress indicator.
4. (Optional) Press Escape to close the progress indicator and return to the menu.
5. (Optional) Select the same drive and choose “scsi drive Utilities →SCSI Drive Low-Level Format →View Drive Format Progress” to again view the progress of the format operation.

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

A notification message informs you when the process is complete.

CHL:*n* ID:*n* Drive NOTICE: Scan Drive Successful

6. Choose “scsi drive Utilities →SCSI Drive Low-Level Format →Clear Format Completed Status,” and then choose **Yes** to clear the completed status and make the drive available for logical device operations such as adding reserved space and then adding the drive to a logical device.

Read/Write Test

This section describes how to perform read/write test operations. The physical drive on which this test will be performed cannot be a spare drive (local or global) or a member drive of a logical drive. Disk reserved space must be removed before this menu option is available. See “Changing Disk Reserved Space” on page 197 for more information.

▼ To Perform a Read/Write Test

1. From the Main Menu, choose “view and edit Drives.”
2. Select a drive on which the read/write test is to be performed.
3. Choose “scsi drive Utilities →Read/Write Test.”
4. (Optional) Enable or disable the following options and press **Return** after each change.
 - “Auto Reassign Bad Block”

When this option is enabled, if a bad block is encountered during the Read/Write test, the controller reassigns that block to an unused good block and writes a message to the event log.
 - “Abort When Error Occurs”

When this option is enabled, if an error occurs during the Read/Write test, the test is aborted.

- **“Drive Test for - Read Only” or “Drive Test for - Read and Write”**
Use this option to configure the Read/Write test to perform only read operations on the disk, or to perform both read and write operations on the disk.
5. **When configuration is complete, choose “Execute Drive Testing,” and then choose Yes to test the drive and to display the test progress indicator.**
 6. **(Optional) Press Escape if you want to close the progress indicator and return to the menu.**
 7. **(Optional) Select the same drive and choose “scsi drive Utilities →Read/Write Test →View Read/Write Testing Progress” to again view the progress of the read/write test.**
 8. **(Optional) To view other aspects of this test, select the drive being tested and choose “scsi drive Utilities →Read/Write Test.”**
 - Choose **“List Current Bad Block Table”** to display a table of the bad blocks located so far.
 - Choose **“Abort Drive Testing”** if you want to stop testing the drive.
 9. **Choose “scsi drive Utilities →SCSI Drive Low-Level Format →Clear R/W Test Completed Status,” and then choose Yes to clear the completed status and make the drive available for logical device operations such as adding reserved space and then adding the drive to a logical device.**

Changing Disk 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.

Note – You can change disk reserved space only for drives that are not configured as spare drives or as currently part of a logical drive. If you attempt to change disk reserved space on a drive that is a member of a logical drive, the controller displays an error message. Since disk reserved space is a feature of the physical drive rather than of a logical drive, the information in disk reserved space is not RAID-protected.

▼ To Remove Reserved Space from a Drive

1. **Choose “view and edit Drives” from the Main Menu.**
2. **Select the drive whose reserve space you want to change.**
Ensure that the drive is not part of a logical drive.
3. **Choose “disk Reserved space -,” and then choose Yes to remove the reserve space.**
The “disk Reserved space -” menu option now indicates that the reserved space is unformatted.

▼ To Specify Disk Reserved Space

1. **Choose “view and edit Drives” from the Main Menu.**
The SCSI drive status table is displayed.
2. **Select the drive whose reserve space you want to restore.**
The drive’s status will be NEW DRV.
3. **Choose “disk Reserved space - →256 MB” to allocate reserved space, and then choose Yes to confirm.**
The drive’s status changes to FRMT DRV.

Host and Drive Channels

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

Note – Where procedures vary between platforms, the headings are marked appropriately.

Topics covered in this chapter include:

- “Host and Drive Channel Status Table” on page 200
- “Configuring Channels as Host or Drive” on page 201
- “Creating Additional Host IDs” on page 201
- “Deleting a Host Channel SCSI ID” on page 201
- “Drive Channel SCSI IDs” on page 202
- “Setting SCSI Channel Termination (SCSI Only) (Reserved)” on page 202
- “Setting Transfer Clock Speed (SCSI Only)” on page 203
- “Setting the SCSI Transfer Width (SCSI Only)” on page 204
- “Enabling SCSI Bus Byte Parity Checking (SCSI Only)” on page 204
- “Viewing Chip Information” on page 205
- “Viewing Channel Host-ID WWN Information (FC and SATA Only)” on page 206
- “Viewing Device Port Names (WWPN) (FC and SATA Only)” on page 206
- “Setting a Channel’s Data Rate (FC and SATA Only)” on page 207
- “Issuing a Loop Initialization Primitive (FC and SATA Only)” on page 209

Host and Drive Channel Status Table

To view and configure channels choose “**view and edit channels**” from the Main Menu. See “Channel Status Table” on page 298 for a description of the Channel Status Table.

▼ To Check and Configure Host and Drive Channels

1. From the Main Menu, choose “**view and edit channels**” to display the status of all host and drive channels for this controller.

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	40	NA	AUTO	Serial	F	NA	2 GHz	Serial
1	Host	NA	42	AUTO	Serial	F	NA		
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	2 GHz	Serial
5	Host	NA	46	AUTO	Serial	F	NA		



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

Each controller has a separate RS-232 port as well as an Ethernet port. This architecture ensures continuous communication in case one controller fails. Since the connection is established to only one controller at a time (even when the array is in redundant mode), the CurSyncClk and CurWid settings are displayed for the primary controller. If you map one LUN to the primary controller and another LUN to a secondary controller, only the established connection to the primary controller is displayed. If a primary ID is not mapped to a channel, and a secondary ID is mapped, either ASYNC displays in the CurSyncClk field or the field is left blank.

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 negotiated rate on certain errors (predominantly parity errors). There is little or no performance change.

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

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

Configuring Channels as Host or Drive

This menu option is described in Chapter 4 and Chapter 5.

- For SCSI arrays, see “Channel Settings” on page 56.
- For FC or SATA arrays, see “Channel Settings” on page 88.

Creating Additional Host IDs

This menu option is described in Chapter 4 and Chapter 5.

- For SCSI arrays, see “To Add or Delete a Unique Host ID” on page 58.
- For FC or SATA arrays, see “To Add or Delete a Unique Host ID” on page 93.

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. From the Main Menu, choose “view and edit channels.”
2. Select the host channel with the SCSI ID you want to delete.
3. Choose “view and edit scsi Id” to display existing IDs.
4. Select the ID you want to delete.

5. Choose “Delete Channel SCSI ID,” and then choose Yes to confirm the deletion.

This change does not take effect until the controller is reset.

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?

6. Choose Yes to reset the controller.

Drive Channel SCSI IDs

When a drive channel or DRV + RCCOM channel is selected, the “**view and edit channels**” menu provides two menu options to change the default drive SCSI IDs:

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

These menu options are normally not used. However, when reassigning drive channels as Drive+RCCOM for a Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array, it is necessary to assign a new secondary controller ID. See “To Configure Channels 4 and 5 as Additional DRV + RCCOM Channels” on page 90 for more information.



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 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. From the Main Menu, choose “view and edit channels.”
2. Select the channel for which you want the terminator enabled or disabled.
3. Choose “scsi Terminator,” and choose Yes to change the setting.
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 host or drive channels. Do not use this menu option. It is reserved and should be used only by qualified technicians.

▼ To Change the Sync Transfer Clock Speed (SCSI Only)

1. From the Main Menu, choose “view and edit channels.”
2. Select the drive channel or host channel whose sync transfer clock speed you want to change.
3. Choose “sync transfer Clock” to display a menu of sync transfer clock speeds.
4. Select a clock speed, and choose Yes to confirm your choice.
This change does not take effect until the controller is reset.

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?

5. Choose Yes to reset the controller.

Setting the SCSI Transfer Width (SCSI Only)

Typically the default setting for transfer width is not changed for the host or drive channel. Do not use this menu option. It is reserved and should be used only by qualified technicians.

▼ To Change the Transfer Width Options (SCSI Only)

1. From the Main Menu, choose “view and edit channels.”
2. Select the drive channel or host channel whose transfer width you want to change.
3. Choose “Wide transfer,” and choose Yes to enable or disable wide transfer.

This change does not take effect until the controller is reset.

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?

4. Choose Yes to reset the controller.

Enabling SCSI Bus Byte Parity Checking (SCSI Only)

The SCSI Bus Byte parity check menu option must be enabled if the SCSI HBA also has this option enabled, which enables a handshake on both the host and the array to check parity and verify that there are no bit errors when transmitting data.

Use the following procedure to enable and disable parity checking.

▼ To Enable or Disable Parity Checking

1. From the Main Menu, choose “view and edit channels.”
2. Select the channel whose parity check setting you wish to change.
3. Choose “parity check,” and then choose Yes to enable parity checking if it is currently disabled, or to disable parity checking if it is currently enabled.

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 a firmware identifier, which might include version information.

▼ To View Chip Information

1. From the Main Menu, choose “view and edit channels.”
2. Select a host or drive channel.
3. Choose “view chip inFormation” to display the channel’s chip information.

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

Viewing Channel Host-ID WWN Information (FC and SATA Only)

Use the “**view channel host-id/Wwn**” menu option to view the I/O processor’s worldwide node name (WWNN) and worldwide port name (WWPN) for a selected host channel. Some host-based management software requires these names to address storage devices.

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

1. From the Main Menu, choose “**view and edit channels.**”
2. Select a host channel.
3. Choose “**view channel host-id/Wwn**” to display the channel’s worldwide node name and worldwide port name.

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	40	NA	AUTO	Serial	F	NA	2 GHz	Serial
1	channel Mode					1	F	NA	
2	view and edit scsi Id					1	F	NA	2 GHz
2<	view chip inFormation								
3	view channel host-id/Wwn					1	F	NA	2 GHz
3<	U wpn>								
4	D WWNN:0x206000C0FF004DE2					1	F	NA	2 GHz
4	i WWPN:0x266000C0FF004DE2								
5	Host	NA	46	AUTO	Serial	F	NA		

Viewing Device Port Names (WWPN) (FC and SATA 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 “**Host-ID WWN name list**” from the “**view and edit Host luns**” 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 for ease of identification. This becomes especially useful when you want to set up multiple filtering entries so you can grant or deny specific hosts access to logical drives. See “LUN Filtering (FC and SATA Only)” on page 114 for more information.

▼ To View a Channel’s Device Port Name List

1. From the Main Menu, choose “view and edit 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 is displayed.

```
view and edit Logical drives
view and edit logical Volumes
```

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Serial
1	WWPN:0x210100E08B2139EA	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 and SATA 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.

Note – For the Sun StorEdge 3511 SATA array, channels 2, 3, 4, and 5 are 2-GHz only.

▼ To Set a Channel's Data Rate

1. From the Main Menu, choose "view and edit channels."
2. Select a host or drive channel.
3. Choose "Data rate" to display a menu of data rate choices:
 - Auto
 - 1 GHz
 - 2 GHz

Note – If the channel being configured is connected to a 1 GHz HBA that does not support auto-negotiation protocols, choose 1 GHz instead of Auto. If you select Auto when the channel is connected to one of these older HBAs, the host may initially be able to access the array, but if that host is powered off and back on, the host will lose access to the array. Refer to the release notes for your array for a list of supported HBAs, with identification of HBAs that must be connected to a 1-GHz channel because they do not support auto-negotiation.

4. Choose a data rate from the menu, and choose Yes to confirm your choice.

Note – Ensure that you set a host channel's data rate to a rate that is compatible with the HBA or data switch that is attached to that channel. If you assign a data rate that is not compatible with an attached HBA or network switch, hosts connected through that channel will be unable to access storage on the array.

Note – Channels 2, 3, 4 and 5 on a Sun StorEdge 3511 SATA array only support 2 GHz. If you try to configure one of these channels to 1 GHz or Auto, the array will appear to change the data rate, and the rate you assign will be displayed in the Channel Status table. However, if you assign a data rate of 1 GHz to one of these channels, the port for that channel is disabled. If you assign a data rate of Auto to one of these channels, the port for that channel is configured at 2 GHz.

This change does not take effect until the controller is reset.

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?

5. Choose Yes to reset the controller.

Issuing a Loop Initialization Primitive (FC and SATA Only)

Whenever a device is added to or removed from a Fibre Channel loop, issue a Loop Initialization Primitive (LIP) so that devices on the loop repeat the discovery process.

▼ To Issue a LIP

1. From the Main Menu, choose “view and edit channels” to display the Channel Status Table.
2. Select a channel on which you want to issue the LIP.
3. Choose “issUe lip,” and then choose Yes to issue a LIP.

Configuration Parameters

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

- “Communication Parameters” on page 212
 - “RS-232 Port Configuration (Reserved)” on page 212
 - “Setting an IP Address” on page 213
 - “Network Protocol Support” on page 215
 - “Setting Telnet Inactivity Timeout Time” on page 216
- “SNMP” on page 217
 - “SNMP Traps” on page 217
 - “SNMP Trap Objects” on page 219
 - “System Group Objects” on page 220
 - “Enterprise Objects” on page 224
- “Configuring the Notification Processing Center (NPC)” on page 226
 - “A Simple Sample agent . ini File” on page 228
 - “A Complete Sample agent . ini File” on page 228
 - “agent . ini File Parameters” on page 229
- “Caching Parameters” on page 231
 - “Enabling and Disabling Write-Back Cache” on page 231
 - “Setting Optimization” on page 232
 - “Setting Periodic Cache Flush Time” on page 232
- “Host-Side Parameters Menu Options” on page 233
 - “Maximum Queued I/O Count” on page 233
 - “LUNs Per Host SCSI ID” on page 234
 - “Maximum Number of Concurrent Host-LUN Connections” on page 234
 - “Number of Tags Reserved for Each Host LUN Connection” on page 235
 - “Peripheral Device Type Parameters (Reserved)” on page 236
 - “Host Cylinder/Head/Sector Mapping Configuration” on page 236
 - “Preparing for Logical Drives Larger Than 253 Gbyte on Solaris Systems” on page 237
 - “Configuring In-Band EI Management” on page 238
 - “Fibre Connection Options (FC and SATA Only)” on page 238
- “Drive-Side Parameters Menu” on page 240
 - “Configuring Drive Motor Spin-Up (Reserved)” on page 240

- “Configuring Disk Access Delay Time” on page 241
- “Configuring Drive I/O Timeout” on page 241
- “Configuring the Maximum Tag Count (Tag Command Queuing)” on page 242
- “Configuring Periodic Drive Check Time” on page 243
- “Configuring Periodic SAF-TE and SES Device Check Time” on page 243
- “Configuring Periodic Auto-Detect Failure Drive Swap Check Time” on page 244
- “Drive Predictable Failure Mode (SMART)” on page 245
- “Auto-Assign Global Spare Drive (FC and SATA Only)” on page 245
- “Disk Array Parameters Menu” on page 246
 - “Setting Rebuild Priority” on page 246
 - “Verification on Writes” on page 247
- “Redundant Controller Parameters” on page 247
 - “Enabling and Disabling Cache Synchronization” on page 248
- “Controller Parameters” on page 248
 - “Configuring a Controller Name” on page 248
 - “Password Validation Timeout” on page 249
 - “Controller Unique Identifier (Reserved)” on page 250
 - “Setting the Controller’s Date and Time” on page 251

Communication Parameters

Use the “**Communication parameters**” menu option to view and change communication settings. Use the “**Internet Protocol (TCP/IP)**” menu option to set or change your array’s IP address.

RS-232 Port Configuration (Reserved)

The RS-232 Port parameters should not be changed. They are reserved for use by technical support personnel.

The RAID controller includes one serial port (COM1).

▼ To Configure the COM Port Baud Rate

1. Choose “view and edit Configuration parameters →Communication Parameters →RS-232 Port Configuration →COM1 Configuration →Baud rate” to display a list of available baud rates.

A list of available baud rates is displayed.

2. Select a baud rate, and choose Yes to confirm.

▼ To Enable or Disable Terminal Emulation Through a Serial Port



Caution – Do not use the “Terminal Emulation” menu option. This menu option is reserved and should be used only if directed by technical support personnel.

- To enable terminal emulation, choose “view and edit Configuration parameters → Communication Parameters → RS-232 Port Configuration → COM1 Configuration → Terminal Emulation,” and choose Yes to confirm.

Setting an IP Address

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

- Sun StorEdge Configuration Service. Refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide* for details.
- The Sun StorEdge CLI. Refer to the *Sun StorEdge 3000 Family CLI 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.

Note – If you assign an IP address to an array to manage it out of band, for security reasons consider using an IP address on a private network rather than a publicly routable network. Using the controller firmware to set a password for the controller limits unauthorized access to the array. Changing the firmware’s Network Protocol Support settings can provide further security by disabling the ability to remotely connect to the array using individual protocols such as HTTP, HTTPS, telnet, FTP, and SSH.

▼ 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 →Address:."
5. Configure the Ethernet port.

You can specify an IP address and the related Netmask and Gateway addresses manually. If your network is configured so that system addresses can be automatically provided by a DHCP server or RARP server, you can enable this automatic provisioning by typing either DHCP or RARP in place of the IP address.

To configure the port to be assigned an IP address by a DHCP server, type DHCP in the text box and press Return.

To configure the port to be assigned an IP address by a RARP client, type RARP in the text box and press Return.

Note – If you prefer to disable the LAN port so that the array cannot be reached by an IP address, delete the content of the Address field and press Return to set all three of the selected LAN port's fields to Not Set.

If you are manually assigning the IP address you will need to know the IP address your system administrator has assigned for this array, as well as the netmask and gateway addresses to use.

- a. Type an IP address and press Return.
- b. Choose "NetMask."
- c. Type the appropriate netmask and press Return.
- d. Choose "Gateway."
- e. Type the appropriate gateway address and press Return.

6. Press **Escape** to continue, and then choose **Yes** to change the IP address.

This change does not take effect until the controller is reset.

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?

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

Network Protocol Support

For security reasons, you can enable only the network protocols you wish to support, which limits the ways in which security can be breached.

▼ To Enable and Disable Network Protocols

- From the Main Menu, choose “view and edit Configuration parameters → Communication Parameters → Network Protocol Support” to display the list of network protocols you can enable or disable and show the current status of each protocol.

The following sample configuration is appropriate for most situations:

- TELNET - Enabled means that telnet access to the IP address is enabled.

Note – If you enable telnet access, users connecting to the array using telnet are prompted for a password. If a password has been set for the controller, type it at that time. If no password has been set, press Return.

- HTTP - Disabled means that the Hypertext Transport Protocol access is disabled.
- HTTPS - Disabled means that the Secure Hypertext Transport Protocol access is disabled.
- FTP - Disabled means that File Transfer Protocol access is disabled.

Note – If you enable FTP access, users connecting to the array using FTP are prompted for a password. If a password has been set for the controller, type it at that time. If no password has been set, press Return.

- SSH - Disabled means that Secure Shell protocol access is disabled.

- `PriAgentAll` - Enabled means that an important internal communication protocol used by the controller is enabled. This protocol must remain enabled for Sun StorEdge Configuration Service and the Sun StorEdge CLI to receive information from the controller firmware.

Note – Do not disable `PriAgentAll`.

- `SNMP` - Enabled means that Simple Network Management Protocol access is enabled. SNMP can be used to communicate with external management software.
- `DHCP` - Enabled means that Dynamic Host Configuration Protocol access is enabled. DHCP is used in some networks to dynamically assign IP addresses to systems on the network. See “Setting an IP Address” on page 213 for more information about DHCP.
- `ping` - Enabled means that ping access is enabled, making it possible for hosts on the network to ping the array to see if it is available.

Setting Telnet Inactivity Timeout Time

Set this security measure so that any telnet connection automatically disconnects after the connection has been idle for a configurable period of time. The current setting is displayed with the menu option.

▼ To Set the Telnet Inactivity Timeout

1. From the Main Menu, choose “view and edit Configuration parameters → Communication Parameters →Telnet Inactivity Timeout Time -” to display the current timeout setting as well as a menu of alternate choices.
2. Select a time interval or select “Disable,” and then choose Yes to confirm your choice.

The new time interval is displayed with the “Set Telnet Inactivity Timeout Time” menu option.

SNMP

Simple Network Management Protocol (SNMP) is a widely used network monitoring and control protocol. Data is passed from SNMP agents reporting activity on each network device to the workstation console used to oversee the network. The agents return information contained in a Management Information Base (MIB), which is a data structure that defines what is obtainable from the device and what can be controlled (turned on and off, etc.).

The Sun StorEdge MIB for your array is included with the latest controller firmware patch. Refer to the patch README for the file name and location. Refer to the documentation for your SNMP manager for details on installing MIB files. Instructions for loading MIBs in HP OpenView Network Node Manager are provided in Appendix F.

An SNMP Object Identifier (OID) is a number assigned to devices in a network for identification purposes. OID numbering is hierarchical. Using the IETF notation of digits and dots resembling very long IP addresses, various registries such as ANSI assign high-level numbers to vendors and organizations. They, in turn, append digits to the number to identify individual devices or software processes.

For details on sending SNMP traps via the RAID firmware, see “Configuring the Notification Processing Center (NPC)” on page 226. For information on using Sun StorEdge Configuration Service to send SNMP traps, refer to the “Email and SNMP” appendix in the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

Note – Sun StorEdge 3000 family arrays can use either SNMPv1 or community based SNMPv2 (v2c) protocols for SET and GET messages, but they only send SNMPv1 trap messages. Regardless of protocol, the security mechanism is community based and the community name is “public” for traps as well as SET and GET messages.

SNMP Traps

The controller agent sends an SNMPv1 trap message for each controller event. SNMPv1 traps are identified by the value of three fields:

- enterprise
- generic-trap
- specific-trap

The enterprise value is given as follows:

```
SNMPv2-SMI::enterprises.42.2.180.product-ID
```

where *product-ID* is the Sun StorEdge array product number (3310, 3320, 3510, or 3511). The enterprises OID is defined by SNMPv2-SMI as:

```
.iso.org.dod.internet.private.enterprises
```

where the preceding string has the following numeric value:

```
.1.3.6.1.4.2
```

For example, a Sun StorEdge 3510 FC array has the following enterprise value:

```
SNMPv2-SMI::enterprises.42.2.180.3510
```

This can be described numerically by the following value:

```
.1.3.6.1.4.1.42.2.180.3510
```

The generic-trap field has the following value:

```
enterpriseSpecific
```

where `enterpriseSpecific` has the numeric value, 6.

The `specific-trap` field has the numeric value, 1.

Note – While all products have a `specific-trap` value of 1, the name for this value varies based on the product-ID specified by the MIB.

Note – The `specific-trap` field always has the value 1, regardless of the event severity.

SNMP Trap Objects

The information carried by the trap is given by a list of variables or object/value pairs. The array sends one object/value pair. The Object ID (OID) for the variable has the form:

```
SNMPv2-SMI::enterprises.42.2.180.product-ID.1.2.3.1.1
```

where *product-ID* is the Sun StorEdge array product number (3310, 3320, 3510, or 3511).

For example, a Sun StorEdge 3510 FC array has the following trap OID:

```
SNMPv2-SMI::enterprises.42.2.180.3510.1.2.3.1.1
```

The value for the object is a string and contains the following information:

- Source IP address
- Controller unique identifier
- Event severity
- Time stamp
- Event text string
- Primary |secondary controller

Examples:

```
RAID Event: Host=192.168.0.1, Serial/Unique=5BF, Severity=
Notification, Time=10/31/2005 15:37:19 -- LG:3 Logical Drive
NOTICE: Starting Creation(Primary)
```

```
RAID Event: Host=192.168.0.2, Serial/Unique=4504, Severity=Alert,
Time=10/31/2005 15:34:35 -- Controller BBU Absent or Failed!
!(Secondary)
```

Note – The format and contents of this string are subject to change.

All SNMPv1 messages include the SNMP version and community name. The community name for Sun StorEdge 3000 products is “public”.

The SNMP trap listener (`snmptrapd`) is used to show the enterprise, `specific-trap`, and the variable OID and value.

Example:

```
# snmptrapd -f -Lo -n
2005-10-31 15:12:58 NET-SNMP version 5.2.1.2 Started.
2005-10-31 15:13:00 192.168.0.1(via UDP: [192.168.0.1]:2766) TRAP,
SNMP v1, community public
        SNMPv2-SMI::enterprises.42.2.180.3510 Enterprise Specific
Trap (1) Uptime: 4:56:30.00
        SNMPv2-SMI::enterprises.42.2.180.3510.1.2.3.1.1 = STRING:
"RAID Event: Host=192.168.0.1, Serial/Unique=4504, Severity=
Notification, Time=10/31/2005 15:55:25 -- Controller BBU is
Charging !(Secondary) "
2005-10-31 15:13:06 NET-SNMP version 5.2.1.2 Stopped.
```

Note – NET-SNMP is a popular and freely available software package that includes command line tools for performing SNMP requests such as GET and SET and a daemon that listens for traps. For more information, see <http://www.net-snmp.com>.

Note – The agent .ini file must be configured with trap destinations in order to receive traps. See “To Create and Save the NPC Configuration File (agent.ini)” on page 227.

System Group Objects

Sun StorEdge products support standard system group objects under the `iso.org.dod.internet.mgmt.mib-2` subtree, defined numerically as follows:

```
1.3.6.1.2.1
```

The system group defines information about the system and includes the objects described in TABLE 11-1.

TABLE 11-1 System Group Objects

Object	Access	Description
sysDescr	Read Only	A description of the entity
sysObjectID	Read Only	Authoritative identification of the management subsystem
sysUpTime	Read Only	Time since the network management subsystem was initialized
sysContact	Read-Write	Identification and contact information for the person responsible for this node
sysName	Read-Write	An administrative name assigned to the node
sysLocation	Read-Write	Physical location of the node
sysServices	Read Only	A value indicating services offered

The values of these objects can be obtained with an SNMP `GET` message on the OID. A community name of “public” and a protocol version must be specified with the `GET` message. Protocol versions 1 and 2c are supported. The objects designated with access Read-Write can be set with an SNMP `SET` message. The community name “public” must be specified along with a protocol version for SNMP `SET` messages. All other objects are Read Only.

SNMP must be enabled for `GET` and `SET` messages to work. From the Main Menu, choose “view and edit Configuration parameters →Communication Parameters →Network Protocol Support →SNMP” and verify that SNMP is enabled.

sysDescr – The value of this object is a string that describes the product type. For example, using the `NET-SNMP` command, `snmpget`:

```
# snmpget -c public -v 2c 192.168.0.2 sysDescr.0
SNMPv2-MIB::sysDescr.0 = STRING: Sun StorEdge 3310
```

In the preceding example, `public` specifies the community name `public`, and `2c` specifies the protocol version.

sysObjectID – This object is used for administrative purposes to identify the hardware. SNMP managers, such as HP Openview Network Node Manager, can detect this object and value during the discovery process.

Example:

```
# snmpget -c public -v 2c 192.168.0.2 sysObjectID.0
SNMPv2-MIB::sysObjectID.0 =
OID: SNMPv2-SMI::enterprises.42.2.180.3310.1
```

where `public` specifies the community name `public`, `2c` specifies the protocol version, and `3310` indicates a Sun StorEdge 3310 SCSI array.

sysUpTime – This value specifies the time since the last controller reset or power on.

Example:

```
# snmpget -c public -v 2c 192.168.0.2 sysUpTime.0
SNMPv2-MIB::sysUpTime.0 = Timeticks: (1302400) 3:37:04.00
```

sysContact – This object is used to specify a person responsible for the array and their contact information. The object can be set via the NET-SNMP `snmpset` command.

Example:

```
# snmpset -c public -v 2c 192.168.0.2 sysContact.0 s "John Doe,
212 555 1212"
SNMPv2-MIB::sysContact.0 = STRING: John Doe, 212 555 1212
```

where `s` indicates the value is of type `STRING`.

```
# snmpget -c public -v 2c 192.168.0.2 sysContact.0
SNMPv2-MIB::sysContact.0 = STRING: John Doe, 212 555 1212
```

sysName – This object specifies a symbolic name for the array. This value is equivalent to the `controller-name` that can be set and displayed via the CLI. This value can be set with an SNMP `SET` message.

Example:

```
# sccli 192.168.0.2 set controller-name "my array"
sccli: selected se3000://192.168.0.2:58632 [SUN StorEdge 3310
SN#0005BF]

# snmpget -c public -v 2c 192.168.0.2 sysName.0
SNMPv2-MIB::sysName.0 = STRING: my array

# snmpset -c public -v 2c 192.168.0.2 sysName.0 s "my other array"
SNMPv2-MIB::sysName.0 = STRING: my other array

# sccli 192.168.0.2 show controller-name
sccli: selected se3000://192.168.0.2:58632 [SUN StorEdge 3310
SN#0005BF]
controller-name: "my other array"
```

sysLocation – This object specifies the physical location of the array, such as a floor or room number.

Example:

```
# snmpset -c public -v 2c 192.168.0.2 sysLocation.0 s "1st floor"
SNMPv2-MIB::sysLocation.0 = STRING: 1st floor

# snmpget -c public -v 2c 192.168.0.2 sysLocation.0
SNMPv2-MIB::sysLocation.0 = STRING: 1st floor
```

sysServices – The value of this object specifies layers of the TCP/IP architecture offered. This value is hard-coded at 72.

Example:

```
# snmpget -c public -v 2c 192.168.0.2 sysServices.0
SNMPv2-MIB::sysServices.0 = INTEGER: 72
```

Enterprise Objects

Certain array information can be obtained from Sun StorEdge 3000 products via SNMP. These objects are located under the following subtree:

```
.iso.org.dod.internet.private.enterprises.sun.prod.storEdgeEL
```

or, defined numerically as follows:

```
.1.3.6.1.4.1.42.2.180
```

This node is further branched by product ID as follows:

```
.1.3.6.1.4.1.42.2.180.3310  
.1.3.6.1.4.1.42.2.180.3320  
.1.3.6.1.4.1.42.2.180.3510  
.1.3.6.1.4.1.42.2.180.3511
```

The value for these objects can be obtained via an SNMP GET message on the Object ID. A community name of "public" and a protocol version must be specified with the GET message. Protocols '1' and '2c' are supported. These objects are defined by the MIB for each product. All objects are read-only.

See the NET-SNMP documentation for the location where MIBs are installed. A typical location is `~/snmp/mibs`.

Example:

```
# cp path-to-mibs/SUN-STOREDGE-3310-MIB.txt ~/snmp/mibs/
```

Using the NET-SNMP `snmpget` command with the MIB, the array information can be retrieved as shown in the following example.

```
# snmpget -c public -v 2c -m +SUN-STOREDGE-3310-MIB 192.168.0.2  
ctlrUniqueID.0  
SUN-STOREDGE-3310-MIB::ctlrUniqueID.0 = STRING: "5BF"
```

The same example, showing the Object ID in numeric format:

```
# snmpget -On -c public -v 2c -m +SUN-STOREEDGE-3310-MIB 192.168.0.2
ctlrUniqueID.0
.1.3.6.1.4.1.42.2.180.3310.1.1.1.1.16.0 = STRING: "5BF"
```

The `-m` option specifies the MIB to use.

The `snmpwalk` command can be used to retrieve all index values of an Object ID in a sequence.

```
# snmpwalk -c public -v 2c -m +SUN-STOREEDGE-3310-MIB 192.168.0.2
ldID
SUN-STOREEDGE-3310-MIB::ldID.1 = STRING: "4FCF8CED"
SUN-STOREEDGE-3310-MIB::ldID.2 = STRING: "513A2612"
SUN-STOREEDGE-3310-MIB::ldID.3 = STRING: "38131F15"
SUN-STOREEDGE-3310-MIB::ldID.4 = STRING: "1D20B424"
SUN-STOREEDGE-3310-MIB::ldID.5 = STRING: "283EA66A"
```

Note – See the product MIB for details on OIDs and descriptions available for enterprise objects.

Using Multiple Product MIBs

Individual product MIBs contain many of the same names for OIDs. For example, the `ctlrUniqueID.0` object in the MIB for the Sun StorEdge 3510 FC array has the following OID:

```
.1.3.6.1.4.1.42.2.180.3510.1.1.1.1.16.0
```

The `ctlrUniqueID.0` object in the MIB for the Sun StorEdge 3310 SCSI array has the following OID:

```
.1.3.6.1.4.1.42.2.180.3310.1.1.1.1.16.0
```

This can lead to ambiguity if multiple products with multiple MIBs are used. In this case, the fully qualified OID or name must be used.

The fully qualified name can be specified with the MIB module name given by the DEFINITIONS ::=BEGIN token in the MIB.

```
# snmpget -c public -v 2c -On
-m +SUN-STOREDGE-3510-MIB:SUN-STOREDGE-3310-MIB
192.168.0.1 SUN-STOREDGE-3510-MIB::ctlrUniqueID.0
.1.3.6.1.4.1.42.2.180.3510.1.1.1.1.16.0 = STRING: "4504"
```

The fully qualified name can be specified numerically as follows:

```
# snmpget -c public -v 2c -On
-m +SUN-STOREDGE-3510-MIB:SUN-STOREDGE-3310-MIB
192.168.0.1 .1.3.6.1.4.1.42.2.180.3510.1.1.1.1.16.0
.1.3.6.1.4.1.42.2.180.3510.1.1.1.1.16.0 = STRING: "4504"
```

Configuring the Notification Processing Center (NPC)

The Notification Processing Center (NPC) is a sub-module on the controller for event notifications. It consists of three components:

- SNMP traps to be received by SNMP-based monitoring software
- Email messages
- Broadcast messages

The array can send SNMP traps about mass storage events to an enterprise management console that uses SNMP, such as HP OpenView or Sun Management Center. It can also send events as email messages or broadcast them to a number of servers.

You specify the SNMP traps, email messages, and broadcast messages you want to use and configure the recipients of these events by creating a text file called `agent.ini` and storing it in reserved space on your array.

Note – The `agent.ini` file must be configured with trap destinations in order to receive traps. See “To Create and Save the NPC Configuration File (`agent.ini`)” on page 227.

For information about the alternative approach of configuring Sun StorEdge Configuration Service to use SNMP, refer to the “Email and SNMP” appendix of the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

▼ To Create and Save the NPC Configuration File (agent.ini)

1. **Enable FTP on your array by choosing “view and edit Configuration parameters → Communication Parameters → Network Protocol Support → FTP -” and verifying that it is enabled.**
2. **Create a new text file using a plain text editor.**
3. **Type in the necessary information, specifying which sections you want to enable and the host and email addresses necessary.**

For a sample file enabling the sending of events as SMTP traps, see “A Simple Sample agent.ini File” on page 228.

For a sample file enabling the sending of events as SMTP traps, email, and broadcast notifications, see “A Complete Sample agent.ini File” on page 228.

For a full description of agent.ini file parameters, see “agent.ini File Parameters” on page 229.
4. **Save the file as agent.ini.**
5. **FTP to your array from your workstation.**
6. **Login as root and press Return.**
7. **When prompted for a password, press Return.**

If you have previously specified a password for the controller, you need to type that password when prompted.
8. **Set the filetype to binary (BIN).**
9. **cd to the /cfg directory.**
10. **put the file from your workstation into the /cfg directory.**
11. **Quit the FTP session.**
12. **(Optional) For security reasons, disable FTP on your array by choosing “view and edit Configuration parameters → Communication Parameters → Network Protocol Support → FTP -” and verifying that it is disabled.**
13. **Reset the controllers by choosing “system Functions → Reset controller.”**

A Simple Sample agent.ini File

A basic sample configuration is shown below where *nnn.nnn.nnn.nnn* is the IP address of the machine you are referring to, *xxx@address.com* is the sender or recipient's email address, and ENABLED=0 for OFF, or 1 for ON.

This configuration enables the sending of controller events as SNMP traps to a single host IP address (RECEIVER1) and the sending of controller events in email messages from a specified email address (SENDER_MAIL_BOX) through its SMTP mail server (SMTP_SERVER) to a single email address (also called RECEIVER1, but different from the host computer receiving the SMTP traps). Sending controller events as broadcast messages is not enabled and not specified in the file.

```
[SNMP_TRAP]
ENABLED=1
COMMUNITY=public
RECEIVER1=nnn.nnn.nnn.nnn
[EMAIL]
ENABLED=1
SUBJECT=RAID Event
SENDER_MAIL_BOX=xxx@address.com
SMTP_SERVER=123.123.123.123
RECEIVER1=xxx@address.com
```

A Complete Sample agent.ini File

This configuration enables the sending of controller events of any severity level as SNMP traps to four host IP addresses and the sending of controller events in email messages to four email addresses. It also enables sending controller events as broadcast messages to four host IP addresses at the same time.

```
[SNMP_TRAP]
ENABLED=1
SEVERITY=1
RECEIVER1=192.168.0.1
RECEIVER2=192.168.0.2
RECEIVER3=192.168.0.3
RECEIVER4=192.168.0.4
[EMAIL]
ENABLED=1
SEVERITY=1
SUBJECT=Event from StorEdge RAID Controller
SENDER_MAIL_BOX=user@example.com
SMTP_SERVER=192.168.0.1
RECEIVER1=user1@example.com
RECEIVER2=user2@example.com
RECEIVER3=user3@example.com
RECEIVER4=user4@example.com
[BROADCAST]
ENABLED=1
SEVERITY=1
RECEIVER1=192.168.0.1
RECEIVER2=192.168.0.2
RECEIVER3=192.168.0.3
RECEIVER4=192.168.0.4
```

agent.ini File Parameters

The parameters you can specify in the agent.ini file are defined below:

The configuration file is comprised of three major sections: SNMP, Email and Broadcast. Each notifying method can be separately enabled or disabled.

The SNMP_TRAP Section

[SNMP_TRAP] – section header

[ENABLED] – 1=enabled, 0=disabled (applies to this section only)

[SEVERITY] - level of severity of the messages to be received. (1 covers events of all levels. 3 sends only the most serious events.)

- 1. notification
- 2. warning
- 3. alert

[COMMUNITY] – SNMP community name of the destination/ receiver

[RECEIVER#] – The IP address of the receiver computer. Add additional lines to specify multiple receivers. Up to 4 receivers can be configured.

The EMAIL Section

[EMAIL] – section header

[ENABLED] – 1=enabled, 0=disabled (applies to this section only)

[SEVERITY] - level of severity of the messages to be received: notification, 2. warning, 3. alert. "1" covers events of all levels. "3" sends only the most serious events.)

[SUBJECT] – add a topic to email. This can be used to specify the location of the RAID system, if there are many.

[SENDER_MAIL_BOX] – a valid email address to be used as the "from" part of the email message.

[SMTP_SERVER] – SMTP server used to send email. IP address only, do not enter a host name here.

[RECEIVER#] – receiver's email address. The receiver's number followed by an "=" mark, an email address, comma, and the number to specify the message severity level.

The BROADCAST Section

[BROADCAST] – section header

[ENABLED] – 1=enabled, 0=disabled (applies to this section only)

[SEVERITY] – level of severity of the messages to be received: 1. notification, 2. warning, 3. alert. "1" covers events of all levels. "3" only the most serious events will be broadcast.)

[RECEIVER#] – The IP address of the receiver computer. Add additional lines to specify multiple receivers. Up to four receivers can be configured.

Caching Parameters

Caching parameters allow you to configure write-back cache, write-through cache, optimization modes, and periodic flushing of cache to logical drives.

Enabling and Disabling Write-Back Cache

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

In a single-controller configuration, disable the Write-Back Cache feature to avoid the possibility of data corruption in the event of a controller failure. This will have a negative effect on performance. To avoid either issue, use dual controllers.

Note – Using two single-controller arrays in a clustering environment with host-based mirroring provides some of the advantages of using a dual controller. However you still need to disable the Write-Back Cache in case one of the single controllers fails and you want to avoid the risk of data corruption. For this reason, a dual controller configuration is preferable.

The caching parameters configured using the “**view and edit Configuration parameters**” menu option are applied globally to all the logical drives. You can also configure write-back policy for individual logical drives and logical volumes that are independent of the RAID array write-back policy. For more information, see:

- “To Configure a Logical Drive’s Write Policy” on page 148 for a procedure that describes configuring write policy for a specific logical drive.
- “Creating a Logical Volume” on page 155 for a procedure for creating a logical volume that includes instructions for configuring write-back cache for a specific logical volume.
- “Event Trigger Operations” on page 267 for information about setting triggers to automatically switch from write-back-enabled to write-back-disabled (write-through) if certain hardware malfunctions occur.

▼ To Change the Write-Back Cache Option

- Choose “view and edit Configuration parameters →Caching Parameters →Write-Back Cache,” and then choose Yes to confirm that you want to change the write-back cache setting.

Setting Optimization

- For SCSI arrays, see “Cache Optimization Mode (SCSI)” on page 52.
- For Fibre Channel and SATA arrays, see “Cache Optimization Mode (FC and SATA)” on page 84.

Setting Periodic Cache Flush Time

Set Periodic Cache Flush Time so that the controller flushes cache to logical drive storage at specified intervals. This safety measure prevents the accumulation of data in cache that could be lost in the event of power loss. Note that setting this value to an interval of less than one minute (1/2 min or Continuous Sync) may diminish performance.

▼ To Set Periodic Cache Flush Time

1. From the Main Menu, choose “view and edit Configuration parameters →Caching Parameters →Periodic Cache Flush Time” to display the intervals:
 - **Disabled.** Terminate periodic cache flush so the controller flushes cache only after data in cache is written to disk.
 - **Continuous Sync.** Continuously flush data from cache to logical drive storage.
 - **1/2 min.** Flush cache to logical drive storage after each 30-second interval.
 - **1 min.** Flush cache to logical drive storage after each one-minute interval.
 - **2 min.** Flush cache to logical drive storage after each two-minute interval.
 - **5 min.** Flush cache to logical drive storage after each five-minute interval.
 - **10 min.** Flush cache to logical drive storage after each 10-minute interval.
2. Select a cache flush interval or select Disable to terminate periodic cache flush, and then choose Yes to confirm your choice.

Host-Side Parameters Menu Options

The Host-Side 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
- Peripheral Device Type Parameters (Reserved)
- Host Cylinder/Head/Sector Mapping Configuration
- Preparing for Logical Drives Larger Than 253 Gbyte on Solaris Systems
- Configuring In-Band EI Management
- Fibre Connection Options (FC and SATA Only)

Maximum Queued I/O Count

Use Maximum Queued I/O Count 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 1024 I/O operations per logical drive.

The appropriate Maximum Queued I/O Count setting depends on how many I/O operations the attached servers and the controller itself are performing. This can vary according to the amount of host memory present, the number of drives and their size, and buffer limitations.

▼ To Set the Maximum Queued I/O Count

1. **From the Main Menu, choose “view and edit Configuration parameters →Host-side Parameters →Maximum Queued I/O Count” to display a list of values.**
2. **Select a value, and then choose Yes to confirm your choice.**

This change does not take effect until the controller is reset.

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?

3. **Choose Yes to reset the controller.**

LUNs Per Host SCSI ID

Use this function to change the number of LUNs supported per host SCSI ID. Each time a host channel ID is added, it supports 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 or Sun StorEdge 3320 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 ($4 \times 32 = 128$). If you want to allocate more than four host channel IDs, you must 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. From the Main Menu, choose “view and edit Configuration parameters →Host-side Parameters →LUNs per Host SCSI ID” to display a list of values.
2. Choose a value, and then choose Yes to confirm your choice.

This change does not take effect until the controller is reset.

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?

3. Choose Yes to reset the controller.

Maximum Number of Concurrent Host-LUN Connections

Use Max Number of Concurrent Host-LUN Connection to determine the maximum number of concurrently supported host-LUN connections. Change this menu option setting only if you have more than four logical drives or partitions.

Maximum concurrent host LUN connections determines the controller internal resources that are available for use with a specific number of the concurrent connections.

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
- Host B accesses ID 1
- Host C accesses ID 2
- Host D accesses ID 3

These connections are all queued in the cache.

If there is I/O in the cache with four connections, and another host I/O operation arrives that is different from the four that are currently in the cache (for example, host A accesses ID 3), the controller returns busy. This occurs with the concurrent active connections; if the cache is cleared, the controller accepts four different connections again.

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

1. From the Main Menu, choose “view and edit Configuration parameters →Host-side Parameters →Max Number of Concurrent Host-LUN Connection” to display a list of values.
2. Choose a value, and then choose Yes to confirm your choice.

This change does not take effect until the controller is reset.

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?

3. Choose Yes to reset the controller.

Number of Tags Reserved for Each Host LUN Connection

Use this menu option 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 you are directed to change it by technical support.

Each host/LUN connection has 32 (the default setting) tags reserved. This setting ensures that the controller accepts at least 32 tags per connection. 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 connection.

▼ To Modify the Tag Command Queuing on the Host-LUN Connection

1. From the Main Menu, choose “view and edit Configuration parameters →Host-side Parameters→Number of Tags Reserved for each Host-LUN Connection” to display a list of values.

2. Choose a value, and then choose Yes to confirm.

This change does not take effect until the controller is reset.

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?

3. Choose Yes to reset the controller.

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. When you follow the instructions for creating a logical drive using a serial port connection or a telnet session, the “**Peripheral Device Type Parameters**” menu option is unnecessary.

Host Cylinder/Head/Sector Mapping Configuration

Physical 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 appropriate settings.

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 system, you can choose and 64 heads and either the < 65536 cylinder setting or the Variable cylinder setting to cover all logical drives over 253 GByte and under the maximum limit. The controller automatically adjusts the sector count, and then the operating system can read the correct drive capacity.

After changing the size of a disk in the Solaris operating system, 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. From the Main Menu, choose “view and edit Configuration parameters →Host-side Parameters →Host Cylinder/Head/Sector Mapping Configuration →Sector Ranges” to display a list of sector ranges.
2. Select a value, and then choose Yes to confirm.
3. Choose “Head Ranges” to display a list of head ranges.
4. Select a value, and then choose Yes to confirm.
5. Choose “Cylinder Ranges” to display a list of cylinder ranges.
6. Select a value, and then choose Yes to confirm.

Preparing for Logical Drives Larger Than 253 Gbyte on Solaris Systems

The Solaris operating system requires drive geometry for various operations, including `newfs`. For the appropriate drive geometry to be presented to the Solaris operating system for logical drives larger than 253 Gbyte, use the default settings

shown below to cover all logical drives over 253 Gbyte. These settings work for smaller configurations as well. The controller automatically adjusts the sector count, and then the operating system can read the correct drive capacity.

For Solaris operating system configurations, use the values in the following table.

TABLE 11-2 Cylinder and Head Mapping for the Solaris Operating System

Logical Drive Capacity	Cylinder	Head	Sector
< 253 Gbyte	< 65536 (default)	variable	variable (default)
253 Gbyte–1 Tbyte	< 65536 (default)	64 (default)	variable (default)

See “To Change Cylinder and Head Settings” on page 61 for instructions on how to apply these settings to SCSI arrays. See “To Change Cylinder and Head Settings” on page 98 for instructions on how to apply these settings to FC and SATA arrays.

Note – Refer to your operating system documentation for limitations pertaining to device sizes.

Configuring In-Band EI Management

The firmware’s external interface enables interactions between the firmware and external applications. Use In-band External Interface Management to enable or disable in-band management of the array by these applications.

▼ To Configure In-Band EI Management

- From the Main Menu, choose “view and edit Configuration parameters →Host-side Parameters →In-band EI management” to enable or disable in-band external interface communications, and then choose Yes to confirm.

Fibre Connection Options (FC and SATA Only)

Choose the “Loop only” menu option from the Fibre Connection Option menu to support a FC loop configuration. Choose the “Point to point only” menu option to support point-to-point connections. It is important that you choose the correct option for your configuration.

See “Fibre Connection Protocol” on page 97 for information about using this menu option.



Caution – An additional menu option defaults to a loop configuration 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.

For more information about point-to-point and loop configurations, refer to the *Sun StorEdge 3000 Family Best Practices Manual and Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array.

It is important for point-to-point configurations to also specify only a primary 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 “Deleting a Host Channel SCSI ID” on page 201 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. From the Main Menu, choose “view and edit Configuration parameters →Host-side Parameters →Fibre Connection Option.”



Caution – Do not choose the “Loop preferred, otherwise point to point” menu option. This option is reserved for special use and should be used only if directed by technical support.

2. Choose “Loop only” or “Point to point only,” depending on how your network is configured, and then choose Yes to confirm your choice.

This change does not take effect until the controller is reset.

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?

3. Choose Yes to reset the controller.

Drive-Side Parameters Menu

The Drive-Side Parameters menu options include:

- Configuring Drive Motor Spin-Up (Reserved)
- Configuring Disk Access Delay Time
- Configuring Drive I/O Timeout
- Configuring the Maximum Tag Count (Tag Command Queuing)
- Configuring Periodic Drive Check Time
- Configuring Periodic SAF-TE and SES Device Check Time
- Configuring Periodic Auto-Detect Failure Drive Swap Check Time
- Drive Predictable Failure Mode (SMART)
- Auto-Assign Global Spare Drive (FC and SATA Only)

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.

Configuring Drive Motor Spin-Up (Reserved)



Caution – Do not use the Drive Motor Spin-Up menu option. It is reserved and should be used only by qualified technicians.

The Drive Motor Spin-up menu option determines how the physical drives in a disk array are started. When the power supply is unable to provide sufficient current for all physical drives and controllers that are powered on at the same time, spinning up the physical drives serially requires less current.

If Drive Motor Spin-Up is enabled, the drives are powered up sequentially and 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 longer for the drive to be ready.

▼ To Spin Up SCSI Hard Drives (Reserved)

- From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Drive Motor Spin-Up,” and then choose Yes to confirm the change.

Configuring Disk Access Delay Time

This function sets the delay time that the controller waits before it tries to access the physical 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. From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Disk Access Delay Time” to display a list of delay intervals.
2. Select a delay interval, and then choose Yes to confirm your choice.

Configuring Drive I/O Timeout

The Drive I/O timeout controls the time interval that the controller waits 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 Drive I/O timeout interval, the drive will be designated a failed drive.



Caution – The correct setting for “**Drive I/O Timeout**” is 30 seconds. Do not change this setting. Setting the timeout to a lower value, or to Default, causes the controller to designate a drive as failed while a drive is still retrying or when a drive is unable to arbitrate the 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 a spare block. However, all of this takes time. The time required to perform these operations can vary between brands and models of drives.

During SCSI bus arbitration, a device with higher priority can use the bus first. A device with lower priority sometimes receives a SCSI I/O timeout when devices of higher priority continue to use the bus.

▼ To Choose Drive I/O Timeout

1. From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Drive I/O Timeout –” to display a list of timeout intervals.
2. Select a timeout interval, and then choose Yes to confirm your choice.

This change does not take effect until the controller is reset.

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?

3. Choose Yes to reset the controller.

Configuring the 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. Each drive has built-in cache that is used to sort all of the I/O requests (“tags”) that are sent to the drive, so the drive can 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. (Please see the following note pertaining to SATA drives.)

Note – The tag count default for SATA drives is 4. To prevent any performance issues, do not go above the default value for SATA.

Note – Changing the maximum tag count to Disable will disable the use of write-back cache by all hard drives.

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 for SCSI arrays and 256 for FC arrays.

▼ To Change the Maximum Tag Count Setting

1. From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Maximum Tag Count” to display a list of available tag count values.
2. Select a tag count number, and then choose Yes to confirm your choice.



Caution – Disabling the maximum tag count disables the use of internal cache by all the physical drives.

Configuring Periodic Drive Check Time

Note – The periodic drive check time setting determines how often the controller checks physical drives. At the specified interval, the controller checks all drives shown in the Drive Status table. If any drive is then removed, the controller detects the removal even if no host attempts to access that drive. The periodic drive check does not force the scanning of a drive that has been added to a SCSI array. See “Scanning Drives (SCSI Only)” on page 179 for more information.

▼ To Set the Periodic Drive Check Time

1. From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Periodic Drive Check Time -” to display a list of intervals.
2. Select an interval, and then choose Yes to confirm your choice.



Caution – Do not set this interval for less than one second. Setting Periodic Drive Check Time to less than one second can adversely impact the array’s performance.

Configuring Periodic SAF-TE and SES Device Check Time

If there are remote devices in your RAID enclosure monitored by SAF-TE or SES, including power supplies, use this function to determine the interval after which the controller checks the status of those devices.



Caution – Do not set this interval for less than one second. Setting Periodic SAF-TE and SES Device Check Time to less than one second can adversely impact the array’s reliability.

▼ To Set the Periodic SAF-TE and SES Device Check Time

1. From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Periodic SAF-TE and SES Device Check Time” to display a list of intervals.
2. Select an interval, and then choose Yes to confirm your choice.

Configuring Periodic Auto-Detect Failure Drive Swap Check Time

This menu option periodically polls the unit to detect the replacement of a bad drive. If no spare drive is present in the array, the logical drive begins an automatic rebuild of a degraded logical drive when the firmware detects replacement of the bad drive.

The drive-swap check time is the interval at which the controller checks to see whether a failed drive has been replaced. 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 is Disabled. When Disabled is selected, the controller does not auto-detect the replacement of a failed drive. The controller is then 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.

Note – The periodic auto-detect failure drive swap check does not force the scanning of a drive that has been added to a SCSI array. See “Scanning Drives (SCSI Only)” on page 179 for more information.

▼ To Set the Auto-Detect Failure Drive Swap Check Time

1. From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Periodic Auto-Detect Failure Drive Swap Check Time.”

A list of intervals is displayed.

2. Select an interval, and then choose Yes to confirm.

When you choose an interval 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 no host attempts to access data on the drive.

Drive Predictable Failure Mode (SMART)

Use this menu option to enable SMART functionality. See “Using SMART Functionality” on page 190 for information about how to configure your Drive Predictable Failure Mode setting.

Note – SMART functionality is not supported on the SATA drives used in Sun StorEdge 3511 RAID controllers or expansion units.

Auto-Assign Global Spare Drive (FC and SATA Only)

This feature is disabled by default.

When you choose Auto-Assign Global Spare Drive, the system automatically assigns global spare status to the unassigned drive with the lowest drive ID. This enables the array to use the global spare to rebuild a logical drive automatically without user intervention when a failing drive needs to be replaced.

▼ To Automatically Assign Replacements to Faulty Drives

- From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Auto-Assign Global Spare Drive,” and then choose Yes to change the setting.

Media Scan at Power-Up

This menu option specifies whether Media Scan should automatically begin whenever the controller is powered up or reset. The default setting is Disabled. The setting you choose is maintained when you power-cycle or reset the controller.

- **From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Media Scan at Power-Up,” and then choose Yes to change the setting from Enabled to Disabled or from Disabled to Enabled.**

Disk Array Parameters Menu

The menu options on the Disk Array Parameters menu are described in this section.

Setting Rebuild Priority

The RAID controller provides a background rebuilding ability. This means the controller is able to serve I/O requests while rebuilding logical drives. The time required to rebuild a logical drive 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. **From the Main Menu, choose “view and edit Configuration parameters →Disk Array Parameters →Rebuild Priority” to display a list of rebuild priority selections.**

A list of the priority selections is displayed:

- **Low.** The default priority. Low allocates the least controller resources to rebuild and most controller resources to I/O operations.
- **Normal.** This priority allocates additional controller resources to speed up the rebuilding process, but decreases I/O performance correspondingly.
- **Improved.** This priority allocates still more resources to the rebuilding process, but decreases I/O performance still more.
- **High.** This priority uses maximum controller resources to complete the rebuilding process in the shortest possible time, but heavily impacts I/O performance.

2. **Choose a rebuild priority, and then choose Yes to confirm your choice.**

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 verification 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 write performance during normal use.

▼ To Enable and Disable Verification Methods

1. From the Main Menu, choose “view and edit Configuration parameters →Disk Array Parameters →Verification on Writes” to display the verification methods that are available.
2. Choose a method that you want to enable or disable, and then choose Yes to confirm the change.

Note – Follow the same procedure to enable or disable each method.

Redundant Controller Parameters

The Redundant Controller Parameters menu allows you to enable or disable cache synchronization.

Enabling and Disabling Cache Synchronization

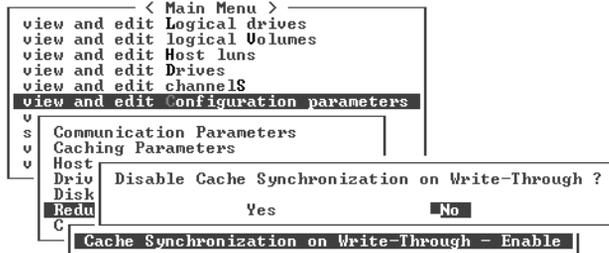
If your redundant controller system is configured with Write-Back Cache disabled, you can disable Cache Synchronization on Write-Through. Mirroring and transferring data between partner controllers will be turned off. This increases array performance for subsystems that operate without write caching.



Caution – If a controller fails and cache synchronization is disabled, all data that was in cache at the time of failure is lost.

▼ To Enable or Disable Cache Synchronization

- From the Main Menu, choose “view and edit Configuration parameters → Redundant Controller Parameters → Cache Synchronization on Write-Through” to change the current setting, and then choose Yes to confirm the change.



Controller Parameters

Procedures for viewing and displaying controller parameters are described in this section.

Configuring a Controller Name

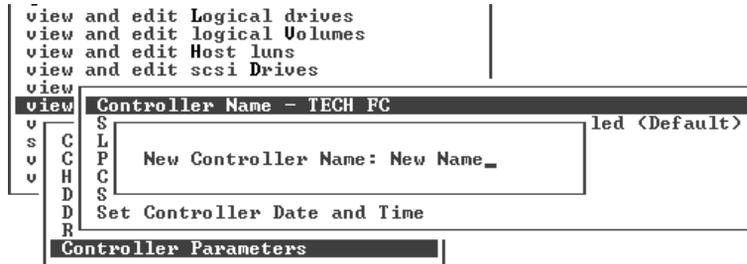
The controller name is displayed only in the firmware application and is used to identify controllers.

Note – The controller password and controller name share a 32-character space. Because the minimum length of the controller password is 1 (when the controller password is empty), the maximum length for the controller name is 31. When the controller name occupies 31 characters, there is only one character left for the controller password, and vice versa.

▼ To View and Display the Controller Name

1. From the Main Menu, choose “view and edit Configuration parameters → Controller Parameters → Controller Name” to name or rename the controller.

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.

Password Validation Timeout

This menu option sets a timeout value that determines how many minutes of inactivity elapse before the password is requested again. This security measure helps avoid unauthorized operation when the user is away from the array.

In most cases, the Always Check default should be left unchanged. If Always Check value is specified, there is no defined timeout and the operator has unlimited opportunities to enter the correct password, but each attempt is validated before access to the firmware’s functions is permitted.

The timeout has no effect if no password has been set. See “Setting and Changing the Controller Password” on page 276 for more password information.

If Disabled is selected, any entry provides immediate access to firmware menu options, regardless of whether a password has been established.

Note – Only one password can be stored.

▼ To Set a Password Validation Timeout

1. From the Main Menu, choose “view and edit Configuration parameters → Controller Parameters → Password Validation Timeout” to display a list of timeout values.
2. Select a validation timeout, and then choose Yes to confirm your choice.

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 specify a new nonzero value unless you have replaced the chassis and the original chassis serial number must be retained. It is especially important in a Sun Cluster environment to maintain the same disk device names in a cluster. Do not change the controller unique identifier unless instructed to do so by qualified service personnel.

▼ To Specify the Controller Unique Identifier

1. From the Main Menu, choose “view and edit Configuration parameters → Controller Parameters → Controller Unique Identifier <hex>” to display the current identifier value and make it possible to change the value.

2. **Type in the value 0 to automatically read the chassis serial number from the midplane, or type in the hexadecimal value for the original serial number of the chassis if the midplane has been replaced and you want to retain the previous identifier.**

The value 0 is immediately replaced with the hex value of the chassis serial number. Any other value is displayed as it was typed.

This change does not take effect until the controller is reset.

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?

3. **Choose Yes to reset the controller.**

Setting the Controller's Date and Time

You can specify the controller date and time so that event messages in the event log display the date and time of the event correctly.

▼ To Set the Controller's Time Zone

1. **From the Main Menu, choose "view and edit Configuration parameters → Controller Parameters → Set Controller Date and Time → Time Zone" to display the currently set time zone as an offset from Greenwich Mean Time (GMT).**

2. Type the appropriate offset from Greenwich Mean Time for your location in the format hour:minute (hh:mm) and press Return.

Your time zone is specified as Greenwich Mean Time (GMT) followed by a plus (+) or minus (-) sign and the number of hours earlier or later your location is from the Greenwich mean time. For instance, the time zone setting for Japan is GMT +9 and the time zone for New York is GMT -4 or -5, depending on daylight savings.

```

< Main Menu >
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit channels
view
v
s C Controller Name - TECH FC
C LCD Title Display - Controller Logo
v C Pass
v H Cont
D SDR4 Default TimeZone - GMT +08:00
D Set New TimeZone - GMT _
R
Cont Time Zone
Date and Time

```

▼ To Set the Controller Date and Time

1. From the Main Menu, choose “view and edit Configuration parameters → Controller Parameters →Set Controller Date and Time →Date and Time.”
2. Type the current date and time in the format MMDDhhmmYYYY and press Return.

For example, enter 072511052004 for 11:05 a.m., July 25, 2004.

```

view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view
view
v
s C Controller Name - TECH SCSI
C Set Telnet Inactivity Timeout Time - Disabled (Default)
v C LCD Title Display - Controller Logo
v C Password Validation Timeout - Always Check
v H Cont
D SDR4
D Set Time and Date [MMDDhhmm[YYYY]] : 072511052004
R
Cont T
Date and Time

```

Peripheral Devices

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

- “Viewing Peripheral Device Controller Status” on page 253
- “Viewing SES Status (FC and SATA Only)” on page 254
 - “Identifying Fans (FC and SATA Only)” on page 256
 - “SES Temperature Sensor Locations (FC and SATA Only)” on page 259
 - “SES Voltage Sensors (FC and SATA Only)” on page 259
 - “SES Power Supply Sensors (FC and SATA Only)” on page 261
- “Viewing Peripheral Device SAF-TE Status (SCSI Only)” on page 261
 - “Identifying Fans (SCSI Only)” on page 263
 - “SAF-TE Temperature Sensor Locations (SCSI Only)” on page 264
 - “SAF-TE Power Supply Sensors (SCSI Only)” on page 265
- “Setting Peripheral Device Entry” on page 265
 - “Redundant Controller - Primary” on page 265
 - “Event Trigger Operations” on page 267
- “Viewing Controller Voltage and Temperature Status” on page 270
- “Fibre Channel Error Statistics (FC and SATA Only)” on page 273

Viewing Peripheral Device Controller Status

To view the status of each controller, from the Main Menu, choose “**view and edit Peripheral devices** → **View Peripheral Device Status.**”

A table displays the status of the available peripheral devices.

```

view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit channels
view and edit Configuration parameters
view and edit Peripheral devices

```

```

S
v
u
View Peripheral Device Status

```

ITEM	STATUS	LOCATION
Redundant Controller	Failback Complete	Primary
SES Device	Enclosure Device	Channel 2 ID 12

FIGURE 12-1 View Peripheral Device Status

Viewing SES Status (FC and SATA Only)

A Fibre Channel array's SCSI Enclosure Services (SES) processor is located on the I/O module. The SES processor monitors chassis-based environmental conditions such as temperature sensor readings, cooling fans status, the beeper speaker condition, power supplies, and slot status. The SES processor is supported by Sun StorEdge Configuration Service and the Sun StorEdge CLI. These chassis sensors are separate from the controller sensors described in "Viewing Controller Voltage and Temperature Status" on page 270.

For Sun StorEdge 3510 FC JBOD arrays and Sun StorEdge 3511 SATA JBOD arrays, both Sun StorEdge Configuration Service and the Sun StorEdge 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/rdsk/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 and SATA Only)

1. From the Main Menu, choose “view and edit Peripheral devices →View Peripheral Device Status →SES Device” to display a list of environmental sensors and other hardware components of that SES device.

STATUS	LOCATION
Failback Complete	Primary

SES Device	Enclosure Device	Channel 2 ID 12
------------	------------------	-----------------

2. Select an item from the list and press Return to display information about it or see a list of its component attributes.

Status :OK
 Actual Speed:Fan at speed 4

STATUS	LOCATION
Failback Complete	Primary

SES Device	Enclosure Device	Channel 2 ID 12
------------	------------------	-----------------

Choosing Overall Status, as 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.

3. 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.

```

view and edit Logical drives
view an
view an Enclosure Descriptor
view an Help Text
view an Device
view an Co
view an Te Element Descriptor rameters
s v Uo Overall Status es
v v Po Element 0 us
v S Au Element 1
v A No Element 2
v C SC Element 3
R PB Element 4
SES Element 5

```

STATUS	LOCATION
Failback Complete	Primary
Enclosure Device	Channel 2 ID 12

In this case the descriptor is Disk Drives.

```

view an
view an Disk Drives
view an
view an Help Text
view an Device
view an Co
view an Te Element Descriptor rameters
s v Uo Overall Status es
v v Po Element 0 us
v S Au Element 1
v A No Element 2
v C SC Element 3
R PB Element 4
SES Element 5

```

STATUS	LOCATION
Failback Complete	Primary
Enclosure Device	Channel 2 ID 12

Identifying Fans (FC and SATA Only)

You can view the status of SES components, including the pair of fans located in each power supply module. A fan is identified in the SES Device menus as a cooling element.

▼ To View the Status of Each Fan

1. From the Main Menu, choose “view and edit Peripheral devices →View Peripheral Device Status →SES Device →Cooling element.”

In some cases you have to “drill down” to display information about components, as as shown in the following illustrations. The following series of screens provide the fan (cooling element) status for each fan.

	STATUS	LOCATION
Failback Complete		Primary

SES Device Enclosure Device Channel 2 ID 12

2. Choose one of the elements (element 0, 1, 2, or 3).

	STATUS	LOCATION
Failback Complete		Primary

SES Device Enclosure Device Channel 2 ID 12

Normal fan speeds are indicated by numbers 1 through 7, indicating speeds in the normal range of 4000 to 6000 RPM. The number 0 indicates that the fan has stopped.

TABLE 12-1 Fan Status and Fan Speeds

Fan Status	Fan RPM
0 Fan stopped	0 - 3999
1 Fan at lowest speed	4000 - 4285
2 Fan at second lowest speed	4286 - 4570
3 Fan at speed 3	4571 - 4856
4 Fan at speed 4	4857 - 5142

TABLE 12-1 Fan Status and Fan Speeds (Continued)

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 power supply module and fan.

Cooling elements in the status table can be identified for replacement as shown in TABLE 12-2. Cooling fan locations are identified in FIGURE 12-2.

TABLE 12-2 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

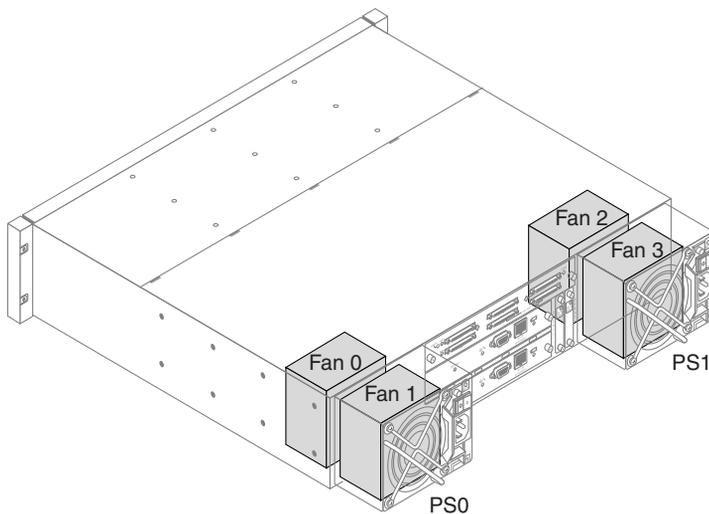


FIGURE 12-2 FC and SATA Cooling Fan Locations

SES Temperature Sensor Locations (FC and SATA 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 in 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**.”

Note – Press the down arrow to access an element ID that is not currently displayed in the displayed list of sensors.

TABLE 12-3 Temperature Sensor Locations (FC and SATA)

Element ID	Description
0	Drive Midplane Left Temperature Sensor #0
1	Drive Midplane Left Temperature Sensor #1
2	Drive Midplane Center Temperature Sensor #2
3	Drive Midplane Center Temperature Sensor #3
4	Drive Midplane Right Temperature Sensor #4
5	Drive Midplane Right Temperature Sensor #5
6	Upper I/O Module (IOM) Left Temperature Sensor #6
7	Upper I/O Module (IOM) Left Temperature Sensor #7
8	Lower I/O Module (IOM) Temperature Sensor #8
9	Lower I/O Module (IOM) Temperature Sensor #9
10	Left Power Supply Temperature Sensor #10
11	Right Power Supply Temperature Sensor #11

SES Voltage Sensors (FC and SATA Only)

Voltage sensors make sure that the array’s voltages are within normal ranges. The voltage components differ for the Sun StorEdge 3510 FC array and the Sun StorEdge 3511 SATA array.

The following tables describe each voltage sensor. The element ID corresponds to the identifier shown when you choose “**view and edit Peripheral devices →View Peripheral Device Status →SES Device →Voltage Sensor.**”

TABLE 12-4 Voltage Sensors for Sun StorEdge 3510 FC Arrays

Element ID	Description	Location	Alarm Condition
0	Voltage Sensor #0	Left Power Supply (5V)	< 4.00V or > 6.00V
1	Voltage Sensor #1	Left Power Supply (12V)	< 11.00V or > 13.00V
2	Voltage Sensor #2	Right Power Supply (5V)	< 4.00V or > 6.00V
3	Voltage Sensor #3	Right Power Supply (12V)	< 11.00V or > 13.00V
4	Voltage Sensor #4	Upper I/O Module (2.5V Local)	< 2.25V or > 2.75V
5	Voltage Sensor #5	Upper I/O Module (3.3V Local)	< 3.00V or > 3.60V
6	Voltage Sensor #6	Upper I/O Module (Midplane 5V)	< 4.00V or > 6.00V
7	Voltage Sensor #7	Upper I/O Module (Midplane 12V)	< 11.00V or > 13.00V
8	Voltage Sensor #8	Lower I/O Module (2.5V Local)	< 2.25V or > 2.75V
9	Voltage Sensor #9	Lower I/O Module (3.3V Local)	< 3.00V or > 3.60V
10	Voltage Sensor #10	Lower I/O Module (Midplane 5V)	< 4.00V or > 6.00V
11	Voltage Sensor #11	Lower I/O Module (Midplane 12V)	< 11.00V or > 13.00V

TABLE 12-5 Voltage Sensors for Sun StorEdge 3511 SATA Arrays

Element ID	Description	Location	Alarm Condition
0	Voltage Sensor #0	Left Power Supply (5V)	< 4.86V or > 6.60V
1	Voltage Sensor #1	Left Power Supply (12V)	< 11.20V or > 15.07V
2	Voltage Sensor #2	Right Power Supply (5V)	< 4.86V or > 6.60V
3	Voltage Sensor #3	Right Power Supply (12V)	< 11.20V or > 15.07V
4	Voltage Sensor #4	Upper I/O Module (1.8V)	< 1.71V or > 1.89V
5	Voltage Sensor #5	Upper I/O Module (2.5V)	< 2.25V or > 2.75V
6	Voltage Sensor #6	Upper I/O Module (3.3V)	< 3.00V or > 3.60V
7	Voltage Sensor #7	Upper I/O Module (1.812V)	< 1.71V or > 1.89V
8	Voltage Sensor #8	Upper I/O Module (Midplane 5V)	< 4.00V or > 6.00V
9	Voltage Sensor #9	Upper I/O Module (Midplane 12V)	< 11.00V or > 13.00V

TABLE 12-5 Voltage Sensors for Sun StorEdge 3511 SATA Arrays (Continued)

Element ID	Description	Location	Alarm Condition
10	Voltage Sensor #10	Lower I/O Module (1.8V)	< 1.71V or > 1.89V
11	Voltage Sensor #11	Lower I/O Module (2.5V)	< 2.25V or > 2.75V
12	Voltage Sensor #12	Lower I/O Module (3.3V)	< 3.00V or > 3.60V
13	Voltage Sensor #13	Lower I/O Module (1.812V)	< 1.71V or > 1.89V
14	Voltage Sensor #14	Lower I/O Module (Midplane 5V)	< 4.00V or > 6.00V
15	Voltage Sensor #15	Lower I/O Module (Midplane 12V)	< 11.00V or > 13.00V

SES Power Supply Sensors (FC and SATA Only)

Each Sun StorEdge 3510 FC array and Sun StorEdge 3511 SATA array has two fully redundant power supplies, with load-sharing capabilities. The sensors monitor the voltage, temperature and fan units in each power supply.

TABLE 12-6 Power Supply Sensors (FC and SATA)

Element ID	Description	Location	Alarm Condition
0	Left Power Supply 0	Left viewed from the rear	Voltage, temperature, or fan fault
1	Right Power Supply 1	Right viewed from the rear	Voltage, temperature, or fan fault

Viewing Peripheral Device SAF-TE Status (SCSI Only)

A SCSI array's SAF-TE processor is located on the SCSI I/O module. It controls environmental monitoring of SAF-TE devices contained in the chassis such as temperature sensors, cooling fans, the beeper speaker, power supplies, and slot status. These chassis sensors are separate from the controller sensors described in "Viewing Controller Voltage and Temperature Status" on page 270.

▼ To Check the Status of SAF-TE Components (SCSI Only)

1. From the Main Menu, choose “view and edit Peripheral devices →View Peripheral Device Status →SAF-TE Device.”

```

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 Peripheral Device Status

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

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

When a drive slot is filled, the drive slot row displays a SCSI ID number.

In a single-bus configuration, ID numbers 0 through 5 and 8 through 13 are shown if all 12 drives are filled (SCSI IDs 6 and 7 are reserved for host communication). Wherever a slot is empty, the message “No Device Inserted” is displayed. See FIGURE 12-3.

Product ID	StorEdge 3310 A	Drive Slot 1	No Device Inserted
Revision Level	1167	Drive Slot 2	No Device Inserted
Unique ID	3030324134430 0	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	73	Drive Slot 8	SCSI ID 2
Temp Sensor 1	77	Drive Slot 9	SCSI ID 3
Temp Sensor 2	75	Drive Slot 10	SCSI ID 4
Temp Sensor 3	80	Drive Slot 11	SCSI ID 5
Temp Sensor 4	80		
Temp Sensor 5	77		
Temp Sensor 6	75		
Temp Alert	Normal		
Speaker Status	Off or No Speaker		
Drive Slot 0	No Device Inserted		

FIGURE 12-3 Example of the SAF-TE Device Status Window for a Single-Bus Configuration

The SAF-TE protocol does not support a split-bus configuration and recognizes only one bus (half the drives) if you have a split-bus configuration. As a result, in a 12-drive split-bus configuration you see the message “Unknown” for six drives on one channel, but you see the ID numbers for the six drives on the other channel, as shown in FIGURE 12-4.

Note – See “Viewing the Status of a Physical Drive” on page 172 for instructions on how to determine whether you have all slots filled in a split-bus configuration.

Product ID	StorEdge 3310 A	Drive Slot 1	Unknown
Revision Level	1168	Drive Slot 2	Unknown
Unique ID	3030324134430 0	Drive Slot 3	Unknown
Cooling Fan 0	Operational	Drive Slot 4	Unknown
Cooling Fan 1	Operational	Drive Slot 5	Unknown
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	68	Drive Slot 8	SCSI ID 2
Temp Sensor 1	69	Drive Slot 9	SCSI ID 3
Temp Sensor 2	69	Drive Slot 10	SCSI ID 4
Temp Sensor 3	75	Drive Slot 11	SCSI ID 5
Temp Sensor 4	71		
Temp Sensor 5	69		
Temp Sensor 6	68		
Temp Alert	Normal		
Speaker Status	Off or No Speaker		
Drive Slot 0	Unknown		

FIGURE 12-4 Example of SAF-TE Device Status Window for a Split-Bus Configuration

Identifying Fans (SCSI Only)

You can view the status of SAF-TE components, including the pair of fans located in each power supply module. A pair of fans is identified in the SAF-TE Device Status window as Cooling Fan 0 or Cooling Fan 1.

If a fan fails and the Status field does not display the Operational value, you must replace the power supply module and fan.

Cooling elements in the status table can be identified for replacement as shown in TABLE 12-2. Cooling fan locations are identified in FIGURE 12-5.

TABLE 12-7 Location of Cooling Fans

Cooling Element #	Fan # and Power Supply Module #
Cooling Fan 0	FANS 0 AND 1, PS 0
Cooling Fan 1	FAN 2 AND FAN3, PS 1

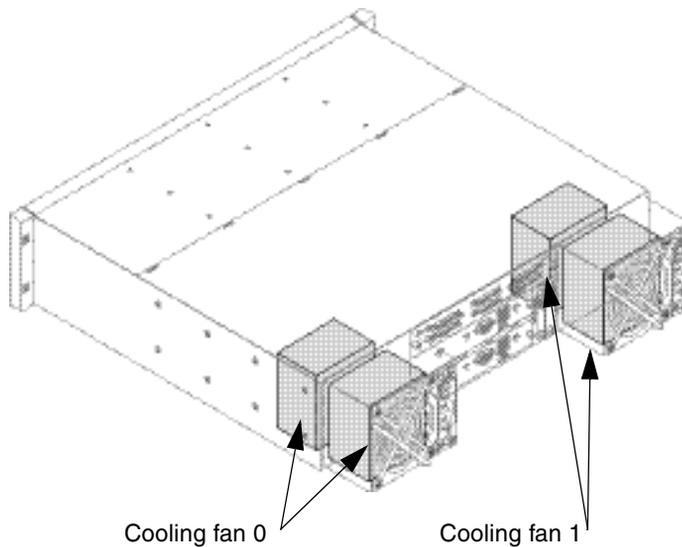


FIGURE 12-5 Cooling Fan Locations

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 in 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 12-8 Temperature Sensor Locations (SCSI)

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)

TABLE 12-8 Temperature Sensor Locations (SCSI) (Continued)

Temp Sensor ID	Description
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

SAF-TE Power Supply Sensors (SCSI Only)

Each Sun StorEdge 3310 SCSI array and Sun StorEdge 3320 SCSI array has two fully redundant power supplies, with load sharing capabilities. The sensors monitor the voltage, temperature and fan units in each power supply.

TABLE 12-9 Power Supply Sensors (SCSI)

Element ID	Description	Location	Alarm Condition
0	Left Power Supply 0	Left viewed from the rear	Voltage, temperature, or fan fault
1	Right Power Supply 1	Right viewed from the rear	Voltage, temperature, or fan fault

Setting Peripheral Device Entry

The Set Peripheral Device Entry menu options include the following:

- Redundant Controller - Primary
- Event Trigger Operations

Redundant Controller - Primary

This menu option enables you to force the failure of either the primary or secondary controller.

Forcing Primary Controller Failure (Reserved)

You can force a primary controller failure to test an array's failover functionality. This feature is normally used only for testing and troubleshooting.

▼ To Force a Primary Controller Failure (Reserved)

- **From the Main Menu, choose “view and edit Peripheral devices →Set Peripheral Device Entry →Redundant Controller →force Primary controller failure,” and then choose Yes to confirm your choice.**

The primary controller is disabled and the array fails over to the secondary controller. A period of time is necessary before the controller's functionality is restored.

Forcing Secondary Controller Failure (Reserved)

You can force a secondary controller failure to test an array's failover functionality. This feature is normally used only for testing and troubleshooting.

▼ To Force a Secondary Controller Failure (Reserved)

1. **From the Main Menu, choose “view and edit Peripheral devices →Set Peripheral Device Entry →Redundant Controller →force Secondary controller failure,” and then choose Yes to confirm your choice.**

A message confirms that the controller has been failed.

```
Controller ALERT: Redundant Controller Failure Detected.
```

2. **Press Escape to clear the message.**

▼ To Restore a Force-Failed Primary or Secondary Controller

1. **From the Main Menu, choose “view and edit Peripheral devices →Set Peripheral Device Entry →Redundant Controller - Primary” to display the following message.**

```
Deassert Reset on Failed Controller ?
```

2. **Choose Yes to restore the controller that you previously force-failed.**

3. Allow several minutes for the failed controller to come back online.

The following message notifies you when the controller is back online:

```
Controller Default Write Policy Restored
```

Event Trigger Operations

Event trigger operations configure an array so that it dynamically switches from write-back-enabled to write-back-disabled (write-through) if a specified failure occurs or threshold is exceeded. Once the problem is corrected, the original write policy is restored.

This change affects the write policy of all logical drives except those whose individual policy has been changed to override the global default write policy for the array.

Except for the “**Temperature exceeds threshold** -” menu option, these trigger operations toggle between being enabled and being disabled each time you change the setting.

Configuring the Controller Failure Event Trigger

If the array has been configured with the write-back cache mode enabled, enable this menu option if you want the array to automatically revert to write-through cache mode (write-back disabled) if one controller in a dual controller array fails.

See “Enabling and Disabling Write-Back Cache” on page 231 for more information about write-back and write-through cache policies.

▼ To Enable or Disable the Controller Failure Event Trigger

- From the Main Menu, choose “view and edit Peripheral devices →Set Peripheral Device Entry →Event Trigger Operations →Controller Failure,” and choose Yes to confirm the change.

Configuring the Battery Backup (BBU) Low Event or BBU Failed Event Trigger

If the array has been configured with the write-back cache mode enabled, enable this menu option if you want the array to automatically revert to write-through cache mode (write-back disabled) if an array’s battery backup fails or falls below its lower threshold.

▼ To Enable or Disable the BBU Low Event or BBU Failed Event Trigger

- From the Main Menu, choose “view and edit Peripheral devices →Set Peripheral Device Entry →Event Trigger Operations →BBU Low or Failed,” and choose Yes to confirm the change.

Configuring the Power Supply Failed Event Trigger

If the array has been configured with the write-back cache mode enabled, enable this menu option if you want the array to automatically revert to write-through cache mode (write-back disabled) if one of the array’s power supplies fails.

▼ To Enable or Disable the Power Supply Failed Event Trigger

- From the Main Menu, choose “view and edit Peripheral devices →Set Peripheral Device Entry →Event Trigger Operations →Power Supply Failed,” and choose Yes to confirm the change.

Configuring the Fan Failure Event Trigger

If the array has been configured with the write-back cache mode enabled, enable this menu option if you want the array to automatically revert to write-through cache mode (write-back disabled) if one of the array’s cooling fans fails.

▼ To Enable or Disable the Fan Failure Event Trigger

- From the Main Menu, choose “view and edit Peripheral devices →Set Peripheral Device Entry →Event Trigger Operations →Fan Failure,” and choose Yes to confirm the change.

Configuring the Temperature Exceeds Threshold Event Trigger

The “**Temperature exceeds threshold -**” menu option differs from other event triggers. It forces a controller shutdown—rather than merely a change in cache policy—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 delay ranging from two minutes to an hour, or disable the controller shutdown entirely. Choose Enable for an immediate shutdown after the

upper threshold limit is exceeded, or choose Disable if you want no trigger for this event. Otherwise, select the time intervals you want to elapse after the threshold is exceeded before the controller shutdown takes place.

▼ To Configure Over-Temperature Controller Shutdown

1. From the Main Menu, choose “view and edit Peripheral devices →Set Peripheral Device Entry →Event Trigger Operations →Temperature exceeds threshold” to display a menu of options and shutdown delay intervals.
2. Select the option or interval you want, and then choose Yes to confirm your choice.

Operating in a NEBS Environment

Sun StorEdge 3000 family products are NEBS Class 3-certified. When this equipment is installed and operated in a NEBS-III or other environment that potentially requires the equipment to be operated outside of the normal temperature range, the Over-Temp Controller Shutdown function must be disabled (this includes testing for such operation).



Caution – Operating this equipment outside of the normal temperature range can adversely affect the operational lifetime of the equipment. The severity of the effect depends on the severity of the overtemp condition and length of time it persists.

If event triggers are enabled for overtemp conditions, the sensors responsible for controller shutdown are:

- Board #1 (85° C)
- Board #2 (85° C). This is the sensor most likely to cause the shutdown, especially the sensor on the IOM in the upper slot.
- CPU (95° C).

Note – There is a thermal switch in each power supply that shuts down the 5VDC & 12VDC current when the switch reaches 95° C. This thermal switch cannot be directly monitored, bypassed, or defeated.



Caution – If two devices make physical contact with each other, thermal conduction between the units can result in higher than expected operating temperatures. Ensure that an air gap exists between RAID controllers and any other equipment in the rack.

Viewing Controller Voltage and Temperature Status

This section describes how to see whether the voltage and temperature of the RAID controller are within normal ranges. These controller sensors differ from the chassis sensors whose status is reported by the SES (FC and SATA) or SAF-TE (SCSI) processor. See “Viewing SES Status (FC and SATA Only)” on page 254 and “Viewing Peripheral Device SAF-TE Status (SCSI Only)” on page 261 for information about chassis sensors.



Caution – The controller sensor settings for your array have been optimized for safe and reliable operations. You might see some menu options that include the word “Default,” which refers only to default firmware settings for a variety of hardware products that use this firmware. However, these menu options are not necessarily the default settings for *your* array. Do not change any voltage or temperature threshold parameters unless specifically advised to do so by service personnel.

▼ To Display Controller Voltage and Temperature Status

- From the Main Menu, choose “view and edit Peripheral devices →Controller Peripheral Device Configuration →View Peripheral Device Status.”

The components checked for voltage and temperature are displayed and defined as normal or out of order.

```

view and edit Logical drives
view and edit logical Volumes
view an
view an
view an
view an
view an
view an
s
v
v
View
Set
Adju
Cont

```

ITEM	VALUE	STATUS
+3.30	3.3840	Operation Normally
+50	5.1530	Operation Normally
+120	12.4420	Operation Normally
CPU Temperature	40.0 (C)	Temperature within Safe Range
Board1 Temperature	42.5 (C)	Temperature within Safe Range
Board2 Temperature	53.0 (C)	Temperature within Safe Range

```

View Peripheral Device Status
Voltage and Temperature Parameters

```

▼ To View or Configure Thresholds



Caution – Equipment damage can result from running equipment outside of normal operating conditions. Do not change any voltage or temperature threshold parameters unless specifically advised to do so by service personnel.

1. From the Main Menu, choose “view and edit Peripheral devices →Controller Peripheral Device Configuration →Voltage and Temperature Parameters” to display a list of trigger threshold categories.
2. Select a parameter whose upper and lower thresholds you want to view or edit.

```

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
View
Set
Adju
Cont
U

```

Trigger Thresholds for +3.30 Events
Trigger Thresholds for +50 Events
Trigger Thresholds for +120 Events
Trigger Thresholds for CPU Temperature Events
Trigger Thresholds for Board Temperature Events

```

Voltage and Temperature Parameters

```

Upper and lower thresholds for the selected parameter are displayed.

3. If you want to change a threshold, select that threshold.

```

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
View Trigger Thresholds for +3.30 Events
Set T
Adj T Upper Threshold for +3.30 Event - Default<3.60>
Cont T Lower Threshold for +3.30 Event - Default<2.90>
U
Voltage and Temperature Parameters

```

Editable threshold values are displayed.

```

view and edit Logical drives
view and edit
view and edit
view and edit Voltage Range from 3.40 to 3.90
view and edit Default Trigger Event: "default"
view and edit Disable Trigger Event: "disable"
view and edit
view and edit Input Voltage Trigger Threshold : default
s
v
v
View Trig
Set T
Adj T Upper Threshold for +3.30 Event - Default<3.60>
Cont T Lower Threshold for +3.30 Event - Default<2.90>
U
Voltage and Temperature Parameters

```

4. If you want to change a threshold, delete the old value, type a new value, and press Return to change the setting.

```

view and edit Logical drives
view and edit
view and edit
view and edit Voltage Range from 3.40 to 3.90
view and edit Default Trigger Event: "default"
view and edit Disable Trigger Event: "disable"
view and edit Input Voltage Trigger Threshold : 3.5
s
v
v
View Trig
Set T
Adj T Upper Threshold for +3.30 Event - Default<3.60>
Cont T Lower Threshold for +3.30 Event - Default<2.90>
U
Voltage and Temperature Parameters

```

Fibre Channel Error Statistics (FC and SATA Only)

You can view FC error statistics that indicate the status of loopback operations 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. The channels and IDs are displayed in hexadecimal format.
- **TYPE.** The device type, such as RAID array, disk, or SES.
- **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 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. From the Main Menu, choose “view and edit Peripheral devices →Fibre Channel Error Statistics →Local Channel Statistics” to display local channel statistics.

CH/ID	TYPE	LIP	LinkFail	LossOfSy	LossOfSi	PrimErr	InvalTxW	InvalCRC
0/28	RAID	1	0	1	1	0	0	0
1/00	RAID	0	0	0	0	0	0	0
2/0F	RAID	1	0	0	0	0	0	0
3/0F	RAID	1	0	0	0	0	0	0
4/2C	RAID	1	0	0	0	0	0	0
5/00	RAID	0	0	0	0	0	0	0

2. Choose “Drive Side Device Statistics” to display drive-side device statistics.

CH/ID	TYPE	LIP	LinkFail	LossOfSy	LossOfSi	PrinErr	InvalTxW	InvalCRC
2/0E	RAID	1	0	0	0	0	0	0
2/04		1	1	35	0	0	1BA3	0
2/03	DISK	1	1	3	0	0	8C5	0
2/08	DISK	1	1	7	0	0	D49	0
2/09	DISK	1	1	2	0	0	C2E	0
2/0F	RAID	1	0	0	0	0	0	0
2/06	DISK	1	1	39	0	0	1B81	0
2/0A	DISK	1	1	20	0	0	BC7	0
2/07	DISK	1	1	1	0	0	C9D	0
2/0C	SES	1	0	0	0	0	0	0
3/0E	RAID	1	0	0	0	0	0	0
3/06	DISK	1	0	30	0	0	AE4	0
3/0A	DISK	1	0	3	0	0	56	0
3/07	DISK	1	0	1	0	0	10B	0
3/0C	SES	1	0	0	0	0	0	0
3/0F	RAID	1	0	0	0	0	0	0

System Functions and Event Logs

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

- “Muting the Beeper” on page 275
- “Setting and Changing the Controller Password” on page 276
- “Resetting the Controller” on page 278
- “Shutting Down the Controller” on page 279
- “The Download Firmware Option (Reserved)” on page 280
- “The Advanced Maintenance Functions Option (Reserved)” on page 280
- “Saving Your Configuration (NVRAM) to Disk” on page 281
- “Restoring Your Configuration (NVRAM) From Disk” on page 284
- “Clearing a Core Dump” on page 286
- “Viewing System Information” on page 286
- “Viewing Event Logs on the Screen” on page 287

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 with event messages and entries in the event log. 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 Appendix E.

The cause of the error condition determines how the alarm is silenced:

- If an alarm is caused by a controller event, use the “**Mute beeper**” menu option to disable the system alarm until another error event occurs.

- If an 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, Mute beeper has no effect.

▼ To Change the Beeper Setting

- From the Main Menu, choose “system Functions →Mute beeper,” and then choose Yes to turn the beeper off until another event occurs.

Setting and Changing the Controller Password

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

If a password is deleted or no password has been set, you may still be prompted for a password. In this case, press Return to continue.

The controller password is also used whenever a user accesses the array using telnet or FTP, if those protocols have been enabled. If a password is deleted or no password has been set, you may still be prompted for a password. In this case, press Return to continue.

See “Network Protocol Support” on page 215 for more information about enabling and disabling network protocols.

Note – The controller verifies the password when you enter the Main Menu from the initial screen, or make configuration changes. If the controller is 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 32-character space. Because the minimum length of the controller name is 1 (when the controller name is empty), the maximum length for the controller password is 31. When the controller password occupies 31 characters, there is only one character left for the controller name, and vice versa.

The procedure for specifying a password differs if no password is currently in effect.

▼ To Create a New Password

1. **From the Main Menu, choose “system Functions →change Password” so you can type in a new password.**
2. **Type the password you want to use, and then press Return.**

Note – The controller password is case-sensitive.

3. **Type the password again, and then press Return to confirm your choice.**
The new password takes effect immediately.

▼ To Change an Existing Password

1. **From the Main Menu, choose “system Functions →change Password” so you can type in a different password.**

You are prompted to type the old password. The password cannot be changed until the old password is first typed correctly.

Note – The controller password is case-sensitive.

2. **Type the current password, and then press Return.**
If the existing password is not typed correctly, an error message is displayed and you cannot change the password.
If the password is correct, you are prompted for a new password.
3. **Type a new password, and then press Return.**
A dialog prompts you to re-enter the password.

4. **Type the password again, and then press Return to confirm your choice.**
The new password takes effect immediately.

▼ To Disable an Existing Password

1. **From the Main Menu, choose “system Functions →change Password.”**

You are prompted to type the old password. The password cannot be changed until the old password is first typed correctly.

Note – The controller password is case-sensitive.

2. **Type the old password in the text area and press Return.**

If the existing password is not typed correctly, an error message is displayed and you cannot change the password.

If the password is correct, you are prompted for a new password.

3. **Press Return without typing anything.**

A dialog prompts you to re-enter the password.

4. **Press Return again to confirm your choice.**

The controller password is deleted and password protection is disabled.

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 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 Reset Controller. Instead, use the “**Shutdown Controller**” menu option and choose Yes when the Reset Controller? prompt is displayed. See “Shutting Down the Controller” on page 279 for more information.

▼ To Reset the Controller Without Saving Cache Contents

1. From the Main Menu, choose “system Functions →Reset controller,” and then choose Yes to reset the controller.

The controller is reset.



Caution – Resetting the controller does not preserve the contents of the cache or write the contents of cache to disk. When the controller is reset, all cache contents are lost. See “Shutting Down the Controller” on page 279 for information about writing the contents of cache to disk before resetting the controller.

Shutting Down the Controller

Always shut down the controller before removing power to the array. After you have chosen this menu option, you can optionally reset the controller so that it is restarted after the shutdown.

The “**Shutdown Controller**” menu option first halts all I/O activity, and so this option should be used only when all I/O activity from hosts has already been halted. The “**Shutdown Controller**” menu option then writes the contents of cache to the drives.

▼ To Shut Down a Controller

1. From the Main Menu, choose “system Functions →Shutdown Controller,” and then choose Yes to confirm that you want to shut down the controller.

A status and confirmation message tells you that the controller shutdown is complete and asks if you want to reset the controller.

```
< Main Menu >
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view
view
view
view
syst
v
v
M
c
R
Shutdown controller
Controller maintenance

**** Shutdown Controller Completed ****
Power off Controller or Reset Controller

Reset Controller ?
Yes No
```

2. Choose Yes if you want to reset the controller.

Note – If you choose No, you must power the controller off and on manually, or use the CLI to restart it. Refer to the *Sun StorEdge 3000 Family CLI User's Guide* for more information.

The Download Firmware Option (Reserved)

This function is no longer available. See the appropriate patch readme file in a firmware patch for firmware download procedures.

The Advanced Maintenance Functions Option (Reserved)

This function is no longer available.

Saving Your Configuration (NVRAM) to Disk

Back up your controller-dependent configuration information. Use the “**save nvrnm to disks**” function to save configuration information whenever a configuration change is made.

When you save your configuration, it is stored in a logical drive.

Note – A logical drive must exist that the controller can write NVRAM content onto.

Note – When you save your configuration, record the configuration information in case you need to refer to it later. Appendix C provides a convenient set of worksheets you can use for this purpose.

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.



Caution – Major upgrades of controller firmware, or replacing a controller with one that has a significantly different version of firmware, might involve differences in non-volatile RAM (NVRAM) that require following special upgrade procedures. For more information, refer to the *Sun StorEdge 3000 Family FRU Installation Guide* and to the release notes for your array.

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 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
- Password Validation Timeout
- 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 or the Sun StorEdge CLI 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 *Sun StorEdge 3000 Family CLI User's Guide*, or the `sccli` man page, for information about the `reset nvram` and `download controller-configuration` commands.

▼ To Save Your Configuration (NVRAM)

1. From the Main Menu, 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.

3. Press Escape to return to the Main Menu.

To restore the configuration, refer to “Restoring Your Configuration (NVRAM) From Disk” on page 284.

Restoring Your Configuration (NVRAM) From Disk

If you have saved your configuration to disk 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 are correct for the array where you are restoring the configuration.

The NVRAM configuration restores all configuration settings (such as channel settings and host IDs) but does not rebuild logical drives. See “Saving Your Configuration (NVRAM) to Disk” on page 281 for information about how to save a configuration file, including advice about saving controller-dependent configuration whenever a configuration change is made.

Note – Using the RAID controller firmware to restore an NVRAM configuration file saved from firmware version 3.2x to a controller that is now running 4.x firmware is not supported. An attempt to do so will fail with a “Restore NVRAM Failed !” error message.

See “Record of Settings” on page 331 for a convenient way to keep a written record of your configuration before saving or restoring configuration files.



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 described in Chapter 4 and Chapter 5 has changed since the configuration file was saved, you might lose access to mismatched channels or drives. You have to change cabling or host or drive channel IDs to correct this mismatch and restore access. On host Solaris workstations, the address of the RAID controller channel must also match what is described in `/etc/vfstab`.

Note – Using Sun StorEdge Configuration Service, you can save a configuration file that can restore all configurations and rebuild all logical drives. However, it also erases all data when it rebuilds the logical drives, so this operation should be performed only when no data has been stored or after all data has been transferred to another array.

▼ To Restore Saved Configuration Settings

1. Choose “system Functions →Controller maintenance →Restore nvram from disks,” and then choose **Yes to restore the NVRAM from disk.**

A confirmation message indicated that the change does not take effect until the controller is reset.

```
It will take effect after resetting controller. Restore NVRAM From  
Disks ?
```

2. Choose **Yes to reset the controller.**

A second confirmation message displays the current NVRAM format and advises you to make sure the NVRAM you are restoring is consistent.

```
The current NVRAM format is 4.x.  
Please make sure that the version of the NVRAM file you want to  
restore is consistent with your firmware version.  
Do you want to restore the NVRAM version previously saved to disk?
```

3. Choose **Yes if you want to restore NVRAM from disk.**

A confirmation message notifies you that the restoration is complete:

```
Controller NOTICE: NVRAM Restore From Disk is Completed
```

4. Chose **No if you do not want to restore NVRAM from disk and leave the current NVRAM version on your active controller unchanged.**

5. Press Escape to continue.

A warning message is displayed:

```
**** WARNING: This is a potentially dangerous operation. ****
The controller will go offline for several minutes.
Data loss may occur if the controller is currently in use.

Do you want to reset the controller now ?
```

6. Choose Yes to continue.

The controller is reset, terminating the telnet connection.

Clearing a Core Dump

Use this function to clear a core dump that has been saved into controller memory in the case of certain specific error messages. Once the message has been recorded and saved in persistent events, clear the core dump so that the message does not recur each time the controller is restarted. See “Controller Alerts” on page 354 for specific details about the error messages that lead to core dumps.

1. Choose “system Functions →Controller Maintenance →Clear core dump.”

A confirmation prompt is displayed:

```
Clear Core Dump ?
```

2. Choose Yes to continue.

The core dump is cleared and the Controller ALERT: Controller Unrecoverable Error 000n [followed by code trap data] event message will not be displayed the next time the controller is reset.

Viewing System Information

You can view a variety of information about your system’s hardware and firmware versions

1. Choose “view system Information.”

Your system information is displayed.

```
< Main Menu >
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
v
v
v
v
s
u
v
```

CPU Type	PPC750
Total Cache Size	1024MB SDRAM(ECC)
Firmware Version	4.23A
Bootrecord Version	1.31H
FW Upgradability	Rev. C
Serial Number	8005640
Battery Backup Unit	Present
Base Board Rev. ID	0
Base Board ID	53
ID of NVRAM Defaults	423A 3510 S491F
Controller Position	Slot A
Current NVRAM Version	4.x

2. Examine the information, paying particular attention to the “Firmware Version,” “ID of NVRAM Defaults,” and “Current NVRAM Version” descriptions.

These settings, which apply to the currently active controller, should be consistent. When you upgrade firmware, the non-volatile RAM version is normally updated as well. If it is not and there is a mismatch, unexpected results can occur. The same is true if an earlier version of NVRAM is subsequently restored.

In the example above, the Firmware Version is 4.23A. This is the version of RAID controller firmware that is running on your active controller. The ID of NVRAM Defaults description shows that the expected version of NVRAM for this firmware version is 423A. The Current NVRAM Version description is 4.x. This shows that there is no serious discrepancy between what is expected and what is present on your active controller.

If, on the other hand, the Current NVRAM Version description was 3.x, it would indicate a serious mismatch between the firmware version and your controller’s NVRAM, which you would need to correct. If the correct version of NVRAM had previously been saved, you would follow the instructions in “Restoring Your Configuration (NVRAM) From Disk” on page 284. For more information about saving NVRAM configuration information, see “Saving Your Configuration (NVRAM) to Disk” on page 281.

Viewing Event Logs on the Screen

When errors occur, you may want to trace the records to see what has happened to your system. The controller event log records up to 100 notification events and another hundred alert and warning events that occur after the system is powered on.

The event log records configuration and operation events and error messages. The event log also shows alarm events reported by the event monitoring unit in each Sun StorEdge 3310 SCSI array or Sun StorEdge 3320 SCSI array and the SES logic in each Sun StorEdge 3510 FC array or Sun StorEdge 3511 SATA array. The status of battery backup units, fans, temperatures, and voltages are sometimes recorded as well.

For each event, a <P> or <S> on the right side of the screen indicates whether the primary or secondary controller in a dual-controller configuration issued the event message.

Powering off or resetting the controller automatically deletes all recorded event log entries.

Note – As you perform the operations described in this guide, you might periodically see event message pop up on the screen. To dismiss an event message after you’ve read it, press Escape. To prevent event messages for displaying so that you can only read them by displaying the event message log, press Ctrl-C. You can press Ctrl-C again at any time to enable pop-up displays of event messages.

▼ To View The Array’s Event Log

1. From the Main Menu, choose “view and edit Event logs” to display a log of recent event messages.

Event Logs	
LG:1 Logical Drive NOTICE: Starting On-Line Initialization	
<Thu Jul 1 13:45:11 2004>	<P>
Creation of Logical Drive 1 Completed	
<Thu Jul 1 13:45:11 2004>	<P>
LG:1 Logical Drive NOTICE: Starting Creation	
<Thu Jul 1 13:45:02 2004>	<P>
LG:0 Logical Drive NOTICE: Starting On-Line Initialization	
<Thu Jul 1 13:44:31 2004>	<P>
Creation of Logical Drive 0 Completed	
<Thu Jul 1 13:44:31 2004>	<P>
LG:0 Logical Drive NOTICE: Starting Creation	
<Thu Jul 1 13:44:30 2004>	<P>
Controller Initialization Completed	
<Thu Jul 1 13:10:13 2004>	<S>
Controller Initialization Completed	
<Thu Jul 1 13:10:07 2004>	<P>

2. Use your arrow keys to move up and down through the list.
3. To clear events from the log after you have read them, use your arrow keys to move down to the first event you want to clear and press Return, and then choose Yes to clear that event log entry and all entries beneath it in the list.

Note – Resetting the controller clears all recorded events. To retain event log entries after controller resets, you can install and use Sun StorEdge Configuration Service.

Maintaining Your Array

Refer to the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array to see hardware-related maintenance and troubleshooting information.

This chapter covers the following firmware-oriented maintenance and troubleshooting topics:

- “Battery Operation” on page 291
 - “Battery Status” on page 292
 - “Battery Support for Cache Operations” on page 293
- “Checking Status Windows” on page 293
 - “Logical Drive Status Table” on page 293
 - “Physical Drive Status Table” on page 296
 - “Channel Status Table” on page 298
- “Upgrading Firmware” on page 301
 - “Patch Downloads” on page 302
 - “Installing Firmware Upgrades” on page 302
 - “Controller Firmware Upgrade Features” on page 303
 - “Upgrading SES and PLD Firmware” on page 304
- “Troubleshooting Your Array” on page 305
 - “Controller Failover” on page 305
 - “RAID LUNs Not Visible to Host” on page 306
 - “Rebuilding Logical Drives” on page 306
 - “Modifying Drive-Side Parameters” on page 310
- “Additional Troubleshooting Information” on page 310

Battery Operation

The battery LED (on the far right side of the I/O controller module) is amber 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

Battery status is displayed at the top of the initial firmware screen. BAT: status displays somewhere in the range from BAD to ---- (charging) to +++++ (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 14-1 Battery Status Indicators

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.
++---	90% charged; adequate to maintain cache memory for 72 hours or more in case of power loss.
+++--	92% charged; adequate to maintain cache memory for 72 hours or more in case of power loss.
++++-	95% charged; adequate to maintain cache memory for 72 hours or more in case of power loss.
+++++	Over 97% 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 77 degrees F (25 degrees C). If the unit is continuously operated at 95 degrees F (35 degrees C) or higher, the battery should be changed every year. The shelf life of your battery is three years.

Note – The RAID controller has a temperature sensor which shuts off battery charging when the temperature reaches 129 degrees F (54 degrees C). When this happens, the battery status might be reported as BAD, but no alarm is written to the event log because 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.

Refer to the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array to see your array's acceptable operating and nonoperating temperature ranges.

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

Battery Support for Cache Operations

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, but you can set an event trigger to make this happen. See "Configuring the Battery Backup (BBU) Low Event or BBU Failed Event Trigger" on page 267 for more information.

Checking Status Windows

The status windows used to monitor and manage the array are described in the following sections:

- "Logical Drive Status Table" on page 293
- "Physical Drive Status Table" on page 296
- "Channel Status Table" on page 298

Logical Drive Status Table

To check and configure logical drives, from the Main Menu choose "**view and edit Logical drives**" and press Return.

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

The status of all logical drives is displayed.

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME	
P0	594EB542	NA	RAID1	34476	GOOD					7	B	2	0	0	
S1	4F342FDA	NA	RAID5	40000	GOOD					7	B	3	0	0	
2			NONE												
3			NONE												
4			NONE												
5			NONE												
6			NONE												
7			NONE												

TABLE 14-2 shows definitions and values for logical drive parameters.

TABLE 14-2 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 indicates no logical volume.
RAID	Assigned RAID level
SIZE (MB)	Capacity of the logical drive
Status1	Logical drive status:
	COPYING The logical drive is in the process of copying from another drive.
	CREATING The logical drive is being initiated.
	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.
	INCOMPLETE Two or more member drives in this logical drive have failed.

TABLE 14-2 Parameters Displayed in the Logical Drive Status Window *(Continued)*

Parameter	Description
	SHUT-DOWN The controller has been shut down using the Shut down command. Restart the controller to restore it to GOOD status.
	REBUILDING The logical drive is being rebuilt.
Status2	Logical Drive status column 2
	I The logical drive is initializing.
	A Adding a physical drive to the logical drive.
	E Expanding a logical drive.
Status3	Logical Drive status column 3
	R The logical drive is rebuilding.
	P Regenerating parity on the logical drive.
O	Stripe size:
	2 4 KB
	3 8 KB
	4 .16 KB
	5 32 KB
	6 64 KB
	7 128 KB
	8 256 KB
C	Write policy setting
	B Write-back
	T Write-through
#LN	Total number of drive members in this logical drive
#SB	Number of standby drives available for the logical drive. This includes local spare and global spare drives available for the logical drive.
#FL	Number of failed 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.

To handle failed, incomplete, or fatal failure status, refer to the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array.

Physical Drive Status Table

To check and configure physical drives, from the Main Menu, choose “**view and edit Drives**” and press Return.

```

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

The Physical Drive Status table is displayed with the status of all physical drives in the array.

Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
2<3>	6	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	7	34732	200MB	0	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	8	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	9	34732	200MB	1	ON-LINE	SEAGATE ST336753FSUN36G
2<3>	10	34732	200MB	GLOBAL	STAND-BY	SEAGATE ST336753FSUN36G
2<3>	12				SES	SUN StorEdge 3510F A

TABLE 14-3 Parameters Displayed in the Physical Drive Status Window

Parameters	Description
Chl	Channel that is assigned to the 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 physical drive member of logical drive x.
Status	COPYING The logical drive is in the process of copying from another drive.
	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. If the drive is a local spare, the LG_DRV column displays the drive number of the logical drive to which the spare is assigned. If the drive is a global spare, the LG_DRV column displays 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 that 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 or the drive is defective and cannot be detected.
	MISSING Drive once existed, but is now missing.
	SB-MISS Spare drive missing.
Vendor and product ID	Vendor and product model information of the drive.

* When a Sun StorEdge 3310 SCSI array or Sun StorEdge 3320 SCSI array is powered up, it can take approximately 30–40 seconds before the drive speed is displayed correctly. Before that happens, the drive speed can display as Async.

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 reinstalled in the array and scanned, the drive status is identified as USED because the drive still has data on it from a logical drive.

When a logical drive is deleted properly, this user 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 Drives**” menu, the drive status changes to NEW.

To replace BAD drives, or if two drives show BAD and MISSING status, refer to the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array.

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

Note – When power is turned on, the controller scans all physical drives that are connected through the drive channels. If a physical drive is connected after a Sun StorEdge 3310 SCSI controller or Sun StorEdge 3320 SCSI controller completes initialization, use the “**Scan scsi drive**” menu option (“**view and edit Drives** →**Scan scsi drive**”) to let the controller recognize the newly added physical drive so you can configure it as a member of a logical drive or as a spare drive.

Channel Status Table

To check and configure channels, from the Main Menu, choose “**view and edit channels**,” and press Return.

The Channel Status table is displayed with the status of all channels on the array.

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
1	Host	40	NA	AUTO	Serial	F	NA	1 GHz	Serial
1	Host	NA	42	AUTO	Serial	F	NA		
2(3;C)	DRV+RCC	14	15	AUTO	Serial	F	NA	2 GHz	Serial
3(2;C)	DRV+RCC	14	15	AUTO	Serial	F	NA	2 GHz	Serial
4	Host	44	NA	AUTO	Serial	F	NA	2 GHz	Serial
5	Host	NA	46	AUTO	Serial	F	NA		

Note – Each controller has an RS232 port and an Ethernet port. This architecture ensures continuous access for communication should either controller fail. Since the connection is established with only one controller at a time, even when the array is in redundant mode, the CurSynClk and CurWid settings are displayed only for the connected controller. Therefore, if a user maps one LUN to the primary controller, and another LUN to a secondary controller, only the LUN mapped with the currently connected controller is displayed through the serial and Ethernet port.



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

TABLE 14-4 Parameters Displayed in the Channel Status Table

Parameters	Description
Ch1	Channel's ID.
Mode	Channel mode.
	RCCOM Redundant controller communication channel. Displays as RCC in the Channel Status table.
	Host The channel is functioning as a host channel.
	Drive The channel is functioning as a drive channel.
	DRV+RCC The channel is functioning as a drive channel with a redundant controller communication channel. (Fibre Channel only).
PID	Primary controller's ID mapping:
	* Multiple IDs were applied (host channel mode only).
	# The ID to which host LUNs are mapped in the host channel mode. The ID for the primary controller in the drive channel mode.
	NA No ID applied.

TABLE 14-4 Parameters Displayed in the Channel Status Table *(Continued)*

Parameters	Description
SID	Secondary controller's ID mapping:
	* Multiple IDs (host channel mode only).
	# The ID to which host LUNs are mapped in the host channel mode. The ID for the secondary controller in drive channel the mode.
	NA No ID applied.
DefSynClk	Default bus synchronous clock:
	xx.x MHz Maximum synchronous transfer rate (SCSI array only)
	x GHz Maximum synchronous transfer rate (FC array only).
	Async Channel is set for asynchronous transfers (SCSI arrays only).
	Auto Channel is configured to communicate at 1 or 2 GHz (FC arrays only).
DefWid	Default bus width:
	Wide Channel is set to allow wide (16-bit) transfers (SCSI arrays only).
	Narrow Channel is set to allow narrow (8-bit) transfers (SCSI arrays only).
	Serial Channel is using serial communication.
S	Signal:
	S Single-ended
	L LVD
	F Fibre
Term	Terminator status:
	On Termination is enabled (SCSI arrays only).
	Off Termination is disabled (SCSI arrays only).
	NA For a redundant controller communications (RCCOM) channel (SCSI arrays) and all FC array channels.
CurSynClk	Current bus synchronous clock. This field only displays values for channels that are assigned to the primary controller.
	xx.x MHz The current speed at which a SCSI array channel is communicating.
	x GHz The current speed at which a FC array channel is communicating.

TABLE 14-4 Parameters Displayed in the Channel Status Table (*Continued*)

Parameters	Description	
	Async	The channel is communicating asynchronously or no device is detected.
	(empty)	The default bus synchronous clock has changed. Reset the controller for changes to take effect.
CurWid	Current bus width. This field only displays values for channels that are assigned to the primary controller.	
	Wide	The channel is currently servicing wide 16-bit transfers (SCSI arrays only).
	Narrow	The channel is currently servicing narrow 8-bit transfers (SCSI arrays only).
	Serial	Channel is using serial communication.
	(empty)	The default bus width has changed. Reset the controller for the changes to take effect.

Upgrading Firmware

From time to time, firmware upgrades are made available as patches. Check the release notes for your array to find out the current patch IDs available for your array.

You can download RAID controller firmware patches from SunSolve Online, located at:

<http://sunsolve.sun.com>

Each patch applies to one or more particular piece of firmware, including:

- Controller firmware
- SES firmware
- PLD firmware
- SATA router firmware (SATA only)
- MUX firmware (SATA only)

Note – Disk drive firmware is provided through Sun disk firmware patches, which include the required download utility. Sun disk firmware patches are separate from Sun StorEdge 3000 family firmware patches. Do not use Sun StorEdge Configuration Service or the Sun StorEdge CLI to download disk drive firmware.

SunSolve has extensive search capabilities that can help you find these patches, as well as regular patch reports and alerts to let you know when firmware upgrades and other patches become available. In addition, SunSolve provides reports about bugs that have been fixed in patch updates.

Each patch includes an associated README text file that provides detailed instructions about how to download and install that patch. But, generally speaking, all firmware downloads follow the same steps:

- Locating the patch on SunSolve that contains the firmware upgrade you want
- Downloading the patch to a location on your network
- Using your array software (Sun StorEdge Configuration Service or the Sun StorEdge CLI) to “flash” the firmware to the device it updates

Note – For instructions on how to download firmware to disk drives in a JBOD directly attached to a host, refer to the README file in the patch that contains the firmware.



Caution – Be particularly careful about downloading and installing PLD firmware. If the wrong firmware is installed, or the firmware is installed on the wrong device, your controller might be rendered inoperable. Always be sure to upgrade your SES firmware first before trying to determine if you need a PLD upgrade.

Patch Downloads

1. Once you have determined that a patch is available to update firmware on your array, make note of the patch number or use SunSolve Online’s search capabilities to locate and navigate to the patch.
2. Read the README text file associated with that patch for detailed instructions on downloading and installing the firmware upgrade.
3. Follow those instructions to download and install the patch.

Installing Firmware Upgrades

It is important that you run a version of firmware that is supported by your array. Before updating your firmware, make sure that the version of firmware you want to use is supported by your array.

Refer to the release notes for your array for Sun Microsystems patches containing firmware upgrades that are available for your array. Refer to SunSolve Online for subsequent patches containing firmware upgrades.

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



Caution – Major upgrades of controller firmware, or replacing a controller with one that has a significantly different version of firmware, might involve differences in non-volatile RAM (NVRAM) that require following special upgrade procedures. For more information, refer to the *Sun StorEdge 3000 Family FRU Installation Guide* and to the release notes for your array.

To download new versions of controller firmware, or SES and PLD firmware, use one of the following tools:

- The Sun StorEdge CLI (with an in-band connection, for Linux and Microsoft Windows hosts, and for servers running the Solaris operating system)
- Sun StorEdge Configuration Service (with an in-band connection, for Solaris and Microsoft Windows hosts)

Note – Do not use both in-band and out-of-band connections at the same time to manage the array. You might cause conflicts between multiple operations.

Note – Disk drive firmware is provided through Sun disk firmware patches which include the required download utility. Sun disk firmware patches are separate from the Sun StorEdge 3000 family firmware patches. Do not use the Sun StorEdge CLI or Sun StorEdge Configuration Service to download disk drive firmware.

Controller Firmware Upgrade Features

The following firmware upgrade features apply to the controller firmware:

- Redundant Controller Rolling Firmware Upgrade

When downloading is performed on a dual-controller system, firmware is flashed onto both controllers without interrupting host I/O. When the download process is complete, the primary controller resets and lets the secondary controller take over the service temporarily. When the primary controller comes back online, the secondary controller hands over the workload and then resets itself for the new firmware to take effect. The rolling upgrade is automatically performed by controller firmware, and the user's intervention is not necessary.

- Automatically Synchronized Controller Firmware Versions

A controller that replaces a failed unit in a dual-controller system often has a newer release of the firmware installed than the firmware in the controller it replaced. To maintain compatibility, the surviving primary controller automatically updates the firmware running on the replacement secondary controller to the firmware version of the primary controller.

Note – When you upgrade your controller firmware in the Solaris operating system, the `format(1M)` command still shows the earlier revision level.

Upgrading SES and PLD Firmware

When you replace an I/O controller, the new controller might have a version of SES or PLD firmware different from the other controller in your array. If this mismatch occurs, when you install a controller you hear an audible alarm and see a blinking amber Event LED.

To synchronize the SES firmware and hardware PLD versions, you must download new SES firmware through Sun StorEdge Configuration Service or the Sun StorEdge CLI.

If you have not installed this software, you must install it from the software CD that shipped with your array. Refer to the *Sun StorEdge 3000 Family Configuration Service User's Guide* for your array to see instructions for downloading firmware for devices. Refer to the *Sun StorEdge 3000 Family CLI User's Guide*, or the `scccli(1M)` man page for similar instructions for using the Sun StorEdge CLI. Refer to the release notes for your array for instructions about where to obtain the firmware that you need to download.

When you open Sun StorEdge Configuration Service or the Sun StorEdge CLI and connect to the array, an error message alerts you to the mismatched version problem.

Troubleshooting Your Array

For hardware troubleshooting information, refer to the Installation, Operation and Service Manual for your array. For additional troubleshooting tips, refer to the release notes for your array.

Controller Failover

Controller failure symptoms include:

- The surviving controller sounds an audible alarm.
- The Controller Status LED turns solid amber 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 takes over for the failed controller.

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. The surviving controller 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 hosts.

Controllers are hot-swappable if they are in a redundant configuration. Replacing a failed controller takes only a few minutes. Since the I/O connections are on the controllers, you might experience some unavailability between the time when cables on the failed controller are disconnected and the time when a new controller is installed and its cables are connected.

To maintain your redundant controller configuration, replace a failed controller as soon as possible. For details, refer to the *Sun StorEdge 3000 Family FRU Installation Guide*.

RAID LUNs Not Visible to Host

By default, all RAID arrays are preconfigured with one or two logical drives. For a logical drive to be visible to the host server, its partitions must be mapped to host LUNs. For mapping details, see “Mapping a Partition to a Host LUN” on page 74 for SCSI arrays or “LUN Mapping” on page 112 for FC and SATA arrays.

To make the mapped LUNs visible to a specific host, perform any steps required for your operating system. Refer to the Installation, Operation and Service Manual for your array to see host-specific information about different operating systems.

Rebuilding Logical Drives

This section describes automatic and manual procedures for rebuilding logical drives. The time required to rebuild a logical drive is determined by the size of the logical drive, the I/O that is being processed by the controller and the array’s Rebuild Priority setting. With no I/O being processed, the time required to build a 2-Tbyte RAID 5 logical drive can be approximately:

- 4.5 hours for a Sun StorEdge 3310 SCSI array or Sun StorEdge 3510 FC array
- 6.5 hours for a Sun StorEdge 3511 SATA array

Note – As disks fail and are replaced, the rebuild process regenerates the data and parity information that was on the failed disk. However, the NVRAM configuration file that was present on the disk is not re-created. After the rebuild process is complete, restore your configuration as described in “Restoring Your Configuration (NVRAM) From Disk” on page 284.

Automatic Logical Drive Rebuild

Rebuild with Spare. When a member drive in a logical drive fails, the controller first determines whether there is a local spare drive assigned to the logical drive. If there is a local spare drive, the controller automatically starts to rebuild the data from the failed drive onto the spare.

If there is no local spare drive available, the controller searches for a global spare drive. If there is a global spare, the controller automatically uses the global spare to rebuild the logical drive.

Failed Drive Swap Detect. If neither a local spare drive nor a global spare drive is available, and 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 Periodic Auto-Detect Failure Drive Swap Check Time, perform the following steps:

1. **From the Main Menu, choose “view and edit Configuration parameters →Drive-side Parameters →Periodic Auto-Detect Failure Drive Swap Check Time.”**

A list of check time intervals is displayed.

2. **Select a Periodic Auto-Detect Failure Drive Swap Check Time interval.**

A confirmation message is displayed.

3. **Choose Yes to confirm.**

When Periodic Auto-Detect Failure Drive Swap Check Time is enabled (that is, when a check time interval has been selected), the controller detects whether the failed drive has been replaced by checking the failed drive’s channel and ID. Once the failed drive has been replaced, the rebuild begins immediately.

Note – This feature requires system resources and can impact performance.

If the failed drive is not replaced but a local spare is added to the logical drive, the rebuild begins with the spare.

FIGURE 14-1 illustrates this automatic rebuild process.

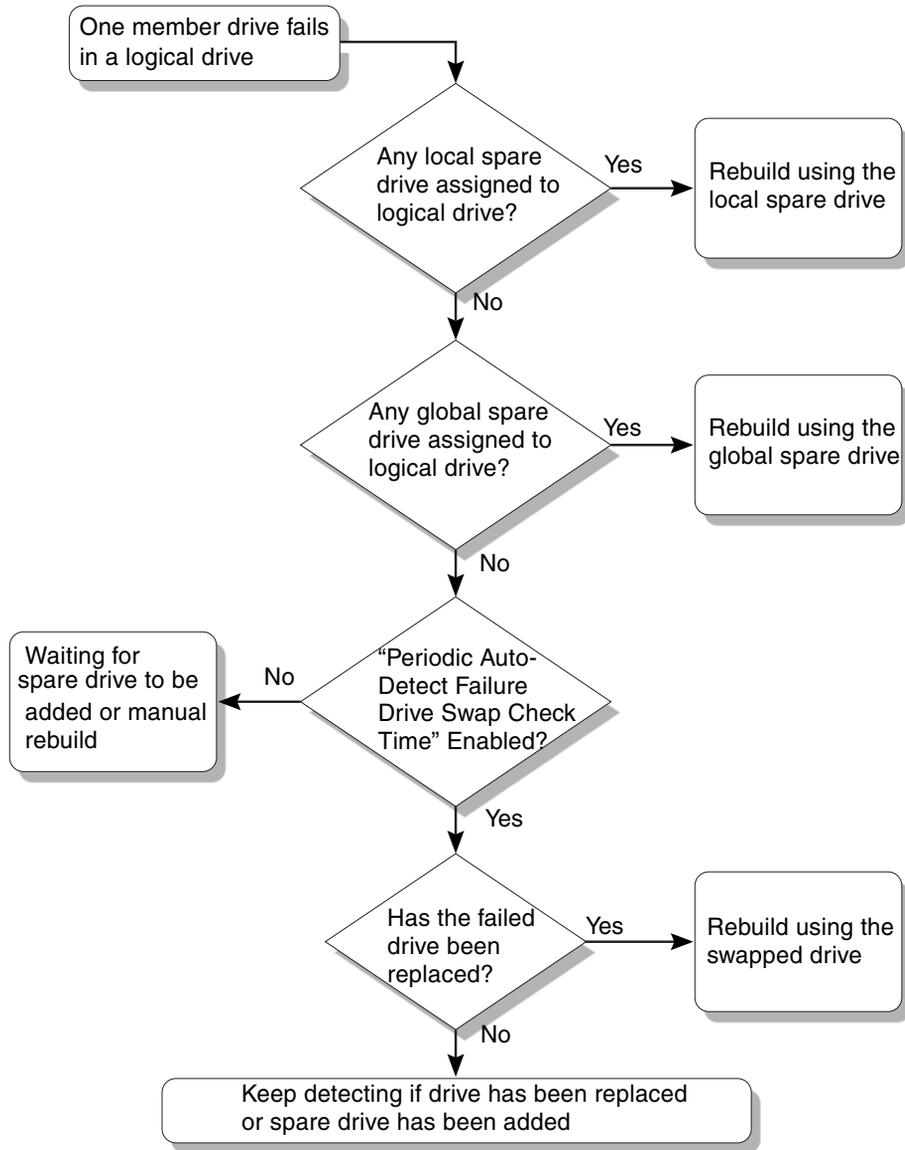


FIGURE 14-1 Automatic Rebuild

Manual Rebuild

When a user applies forced-manual rebuild, the controller first determines whether there is a local spare drive assigned to the logical drive. If a local spare drive is available, the controller automatically starts to rebuild onto the spare drive.

If no local spare drive is available, the controller searches for a global spare drive. If there is a global spare drive, the controller begins to rebuild the logical drive immediately. FIGURE 14-2 illustrates this manual rebuild process.

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

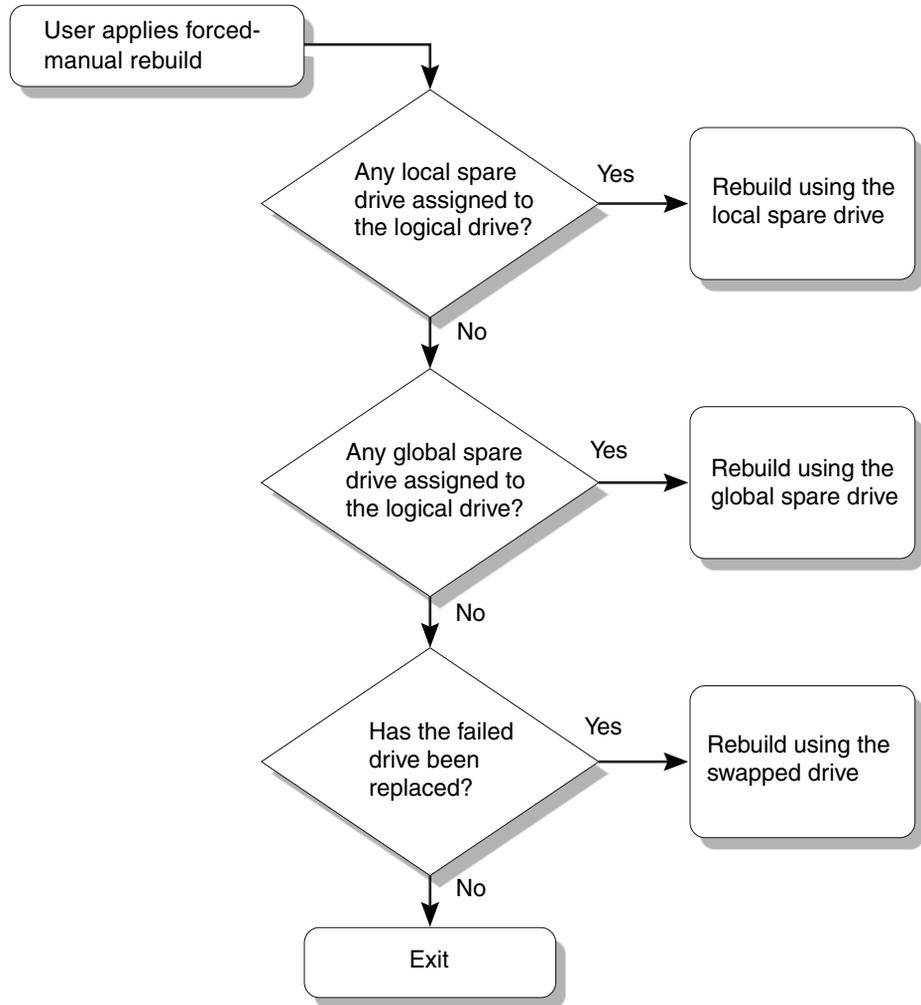


FIGURE 14-2 Manual Rebuild

Concurrent Rebuild in RAID 1+0

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

Modifying Drive-Side Parameters

There are a number of interrelated drive-side configuration parameters you can set using the “**view and edit Configuration parameters**” menu option. It is possible to encounter undesirable results if you experiment with these parameters. Only change parameters when you have good reason to do so.

See “Drive-Side Parameters Menu” on page 240 for cautions about changing sensitive drive-side parameter settings. In particular, do not set Periodic SAF-TE and SES Device Check Time to less than one second, and do not set Drive I/O Timeout to anything less than 30 seconds for FC or SATA arrays.

Additional Troubleshooting Information

For additional troubleshooting tips, refer to the Installation, Operation, and Service manual for your array, and to the release notes for your array.

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:

- “RAID Terminology Overview” on page 312
 - “Logical Drives” on page 312
 - “Logical Volumes” on page 313
 - “Channels, Partitions, and LUN Mapping” on page 313
- “RAID Levels” on page 316
 - “RAID 0” on page 319
 - “RAID 1” on page 319
 - “RAID 1+0” on page 320
 - “RAID 3” on page 321
 - “RAID 5” on page 322
 - “Advanced RAID Levels” on page 323

See also:

- “Local and Global Spare Drives” on page 5
- “Using Both Local and Global Spare Drives” on page 7

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.

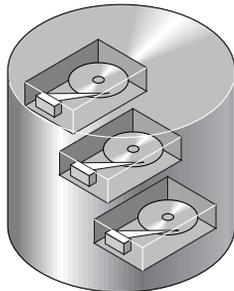


FIGURE A-1 Logical Drive Including Multiple Physical Drives

Logical drives can be configured to provide several distinct RAID levels. For descriptions of each RAID level, see “RAID Levels” on page 316.

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 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 Chapter 6.

A logical volume can be divided into a maximum of 32 partitions for Sun StorEdge 3000 family arrays.

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

Channels, Partitions, and LUN Mapping

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).

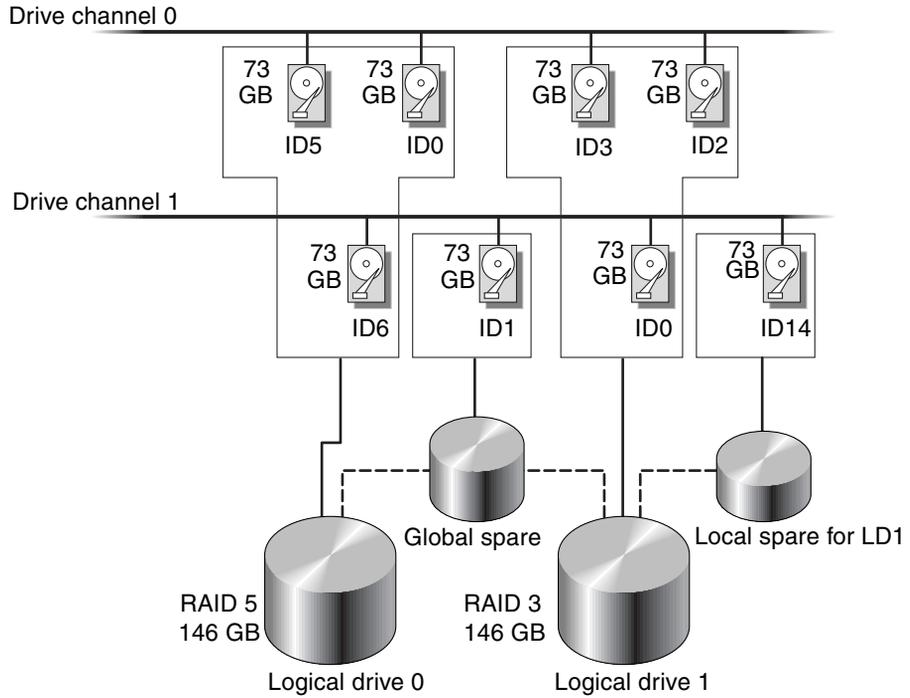


FIGURE A-2 Allocation of Drives in Logical Drive Configurations

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

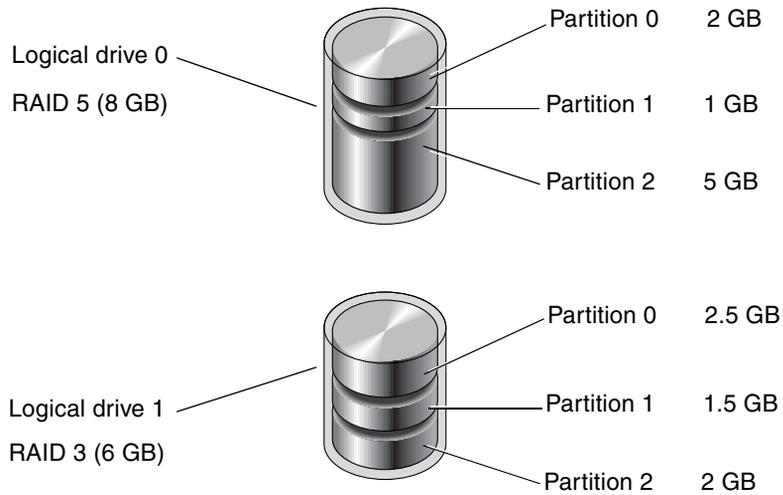


FIGURE A-3 Partitions in Logical Drive Configurations

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.

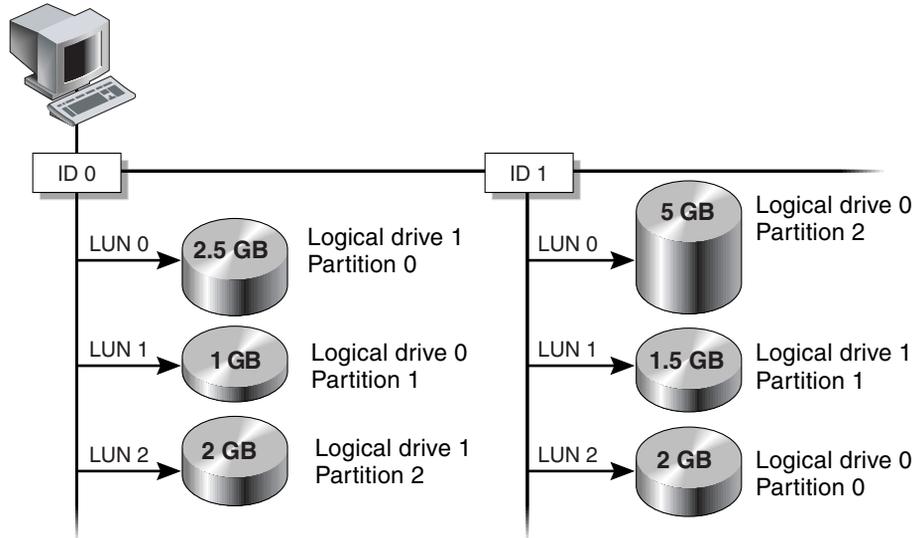


FIGURE A-4 Mapping Partitions to Host ID/LUNs

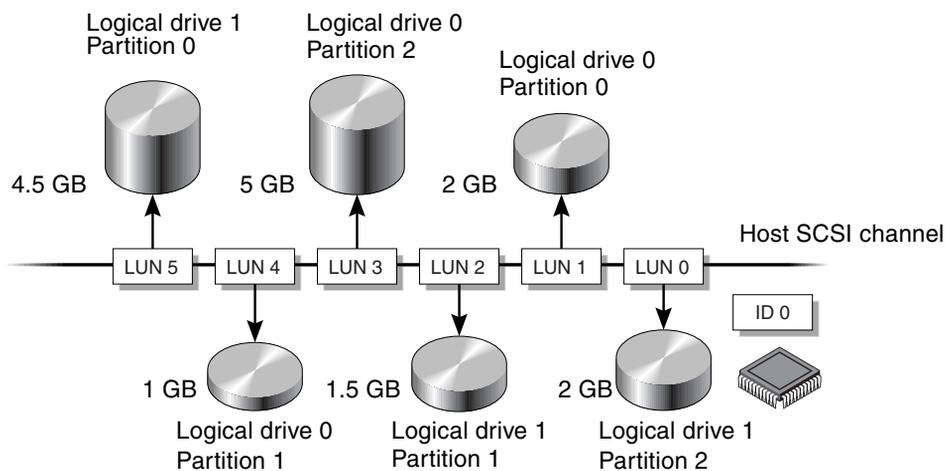


FIGURE A-5 Mapping Partitions to LUNs Under an ID

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, RAID 1, RAID 3, RAID 5, RAID 1+0, RAID 3+0 (30), and RAID 5+0 (50). RAID 1, RAID 3, and RAID 5 are most commonly used.

Note – NRAID does not provide data redundancy. The NRAID option that appears in some firmware menus 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 A-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 A-2 RAID Level Characteristics

RAID Level	Description
NRAID	NRAID is an acronym for Non-RAID. NRAID does not provide data redundancy and is not recommended.
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	<p>Striping with fault tolerance; this is the best-suited RAID level for multitasking or transaction processing.</p> <p>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.</p> <p>The primary advantages of RAID 5 are that:</p> <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. <p>RAID 5 requires at least three drives.</p>
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.

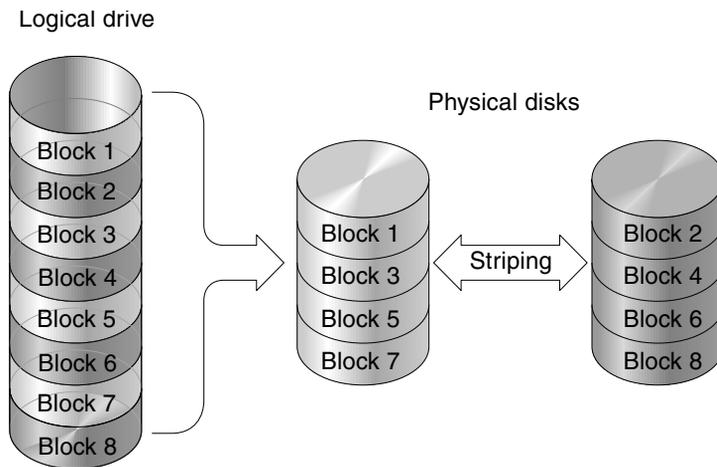


FIGURE A-6 RAID 0 Configuration

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.

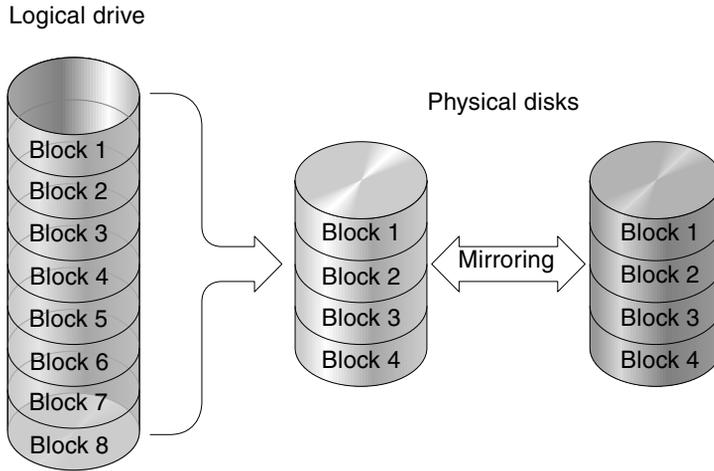


FIGURE A-7 RAID 1 Configuration

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.

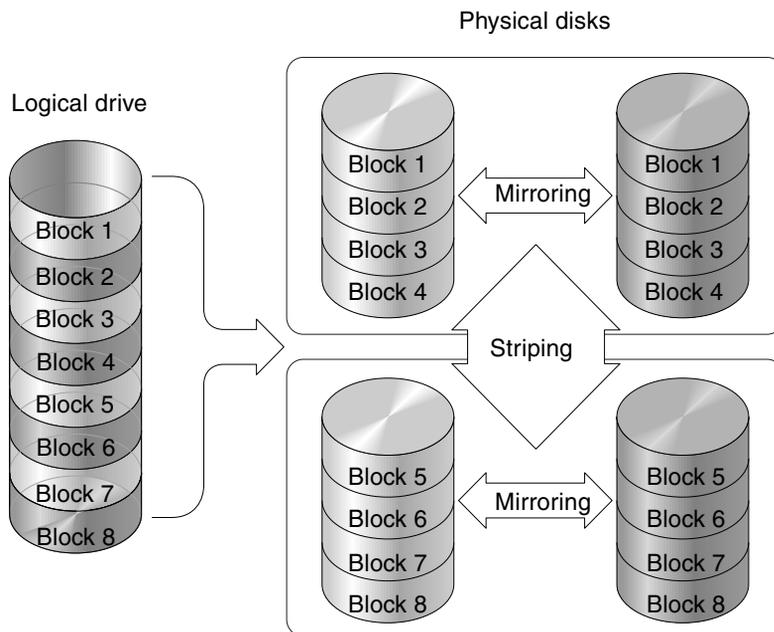


FIGURE A-8 RAID 1+0 Configuration

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.

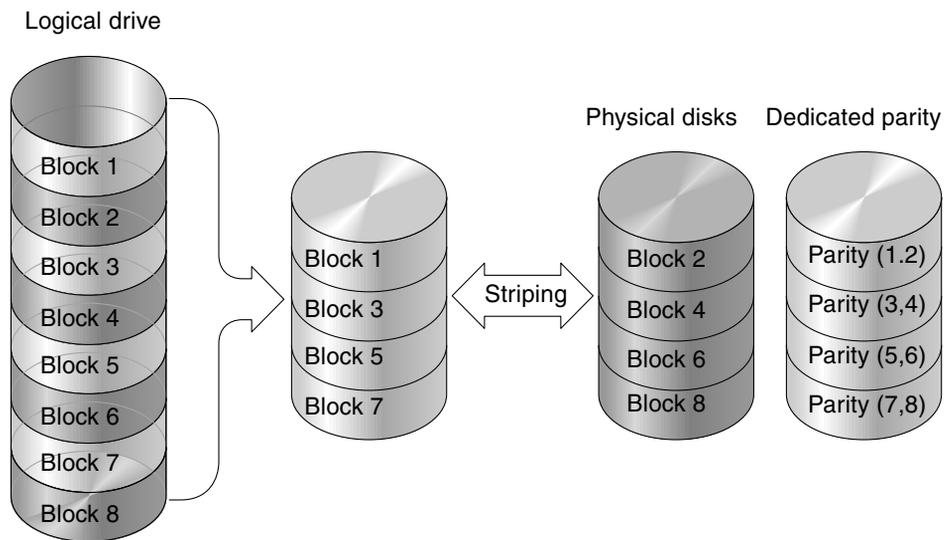


FIGURE A-9 RAID 3 Configuration

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.

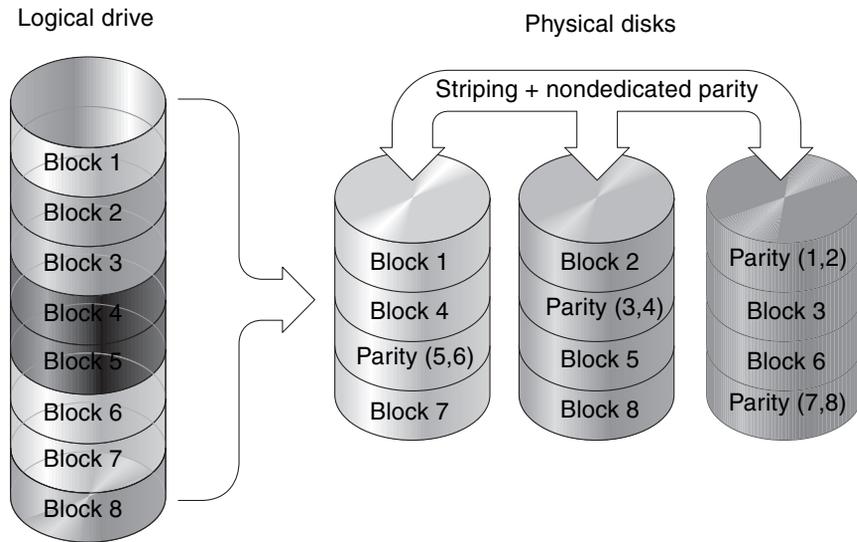


FIGURE A-10 RAID 5 Configuration

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 A-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.

Firmware Specifications

This appendix contains the firmware specifications in the following tables:

- “Basic RAID Management” on page 325
- “Advanced Features” on page 326
- “Caching Operation” on page 327
- “RAID Expansion” on page 327
- “Redundant Controller” on page 328
- “Data Safety” on page 328
- “Security” on page 329
- “Environment Management” on page 329
- “User Interface” on page 330

TABLE B-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	32
Raid Level Dependency to Each Logical Drive	Independent. Logical drive configured in different RAID levels can coexist in an array.
Logical Drive Identification	Unique, controller-generated logical drive ID; logical drive name is user-configurable.
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.

TABLE B-1 Basic RAID Management (*Continued*)

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	Supported.
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 B-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.
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.

TABLE B-2 Advanced Features *(Continued)*

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 B-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 0, RAID 1, 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 B-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 B-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 B-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.

TABLE B-6 Data Safety (Continued)

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: Detect and Clone+Replace
Clone Failing Drive	Users can choose to clone data from a failing drive to a backup drive manually.

TABLE B-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 B-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).
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 B-9 User Interface

Feature	Description
RS-232C Terminal	Supports terminal modes: ANSI, VT-100, ANSI Color. Provides menu-driven, user-friendly, text-based interface.
Ethernet Interface	Supports 10BASE-T communication for out-of-band management.
Beeper Alarm	Warns user when any failure or critical event occurs.

Record of Settings

This appendix provides tables for recording configuration data. Topics covered are:

- “View and Edit Logical Drives” on page 332
 - “Logical Drive Information” on page 332
 - “Logical Drive Partition Information” on page 333
- “LUN Mappings” on page 334
- “View and Edit Drives” on page 335
- “View and Edit Channels” on page 336
- “View and Edit Peripheral Devices” on page 337
 - “View System Information” on page 337
- “Save NVRAM to Disk and Restore From Disk” on page 338

In addition to saving the configuration data in NVRAM to disk, keeping a hard copy of the controller configuration is also recommended. This speeds the re-creation of the RAID in the event of a disaster.

The following tables are provided as a model for recording configuration data.

Note – The configuration data in the NVRAM should be saved to disk or file whenever a configuration change is made.

View and Edit Logical Drives

Logical Drive Information

LG	ID	LV	RAID Level	Size (MB)	Status	O	#LN	#SB	#FL	Name	Disk Reserved Space

View and Edit Drives

Slot	Channel	ID	Size (MB)	Speed	LG DRV? Global Spare? Local Spare?		Vendor and Product ID

View and Edit Peripheral Devices

View System Information

Firmware version
Boot record version
Serial number

Save NVRAM to Disk and Restore From Disk

Update Firmware	Date	Save NVRAM to Disk or File	Date/Location	Restore NVRAM From Disk	Date

Parameter Summary Tables

This appendix lists the firmware parameter settings for Sun StorEdge 3510 FC arrays, Sun StorEdge 3511 SATA arrays, Sun StorEdge 3310 SCSI arrays, and Sun StorEdge 3320 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 339
- “Basic Default Parameters” on page 340
- “Default Configuration Parameters” on page 341
- “Default Peripheral Device Parameters” on page 348
- “Default System Functions” on page 350
- “Parameter Defaults That Must Not Change” on page 351

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.

Basic Default Parameters

These parameters are the primary settings for each array.

TABLE D-1 Logical Drive Parameters (View and Edit Logical Drives)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
Create Logical Drives	1 RAID 0 per array. No spares per array. 1–32 drives per array.	1 RAID 0 per array. No spares per array. 1–16 drives per array.	
Change a Logical Drive Controller Assignment	Primary.		Secondary.

TABLE D-2 Logical Volume Parameters (View and Edit Logical Volumes)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
Create a Logical Volume	Primary controller.	Primary controller.	Secondary.

TABLE D-3 Host LUN Parameters (View and Edit Host LUNs)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting
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 D-4 Drive Parameters (View and Edit Drives)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
FC Drive ID Switch Settings	0	0	0–7

TABLE D-5 Channel Parameters (View and Edit Channels)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	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 parameters in the following tables can be set using the View and Edit Configuration Parameters menu:

- “Introducing Default Parameters” on page 339
- “Caching Parameters” on page 343
- “Peripheral Device Type Parameters” on page 344
- “Host-Side and Drive-Side Parameters” on page 345
- “Other Configuration Parameters” on page 347

The most important parameters 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 next three tables show the default settings and range of available values for various communication parameter settings:

TABLE D-6 Communication Parameters > RS-232 Port Configuration

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
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.

TABLE D-7 Communication Parameter > Internet Protocol (TCP/IP)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
TCP/IP Addresses	DHCP Client	DHCP Client	DHCP client, RARP client, or type IP address, netmask, and gateway

TABLE D-8 Communication Parameters > Network Protocol Support

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
Telnet	Enabled.	Enabled.	Enabled. Disabled.
HTTP	Enabled.	Enabled.	Enabled. Disabled.
HTTPS	Disabled.	Disabled.	Enabled. Disabled.
FTP	Enabled.	Enabled.	Enabled. Disabled.
SSH	Disabled.	Disabled.	Enabled. Disabled.
PriAgent	Enabled	Enabled	Do not change this setting Enabled Disabled

TABLE D-8 Communication Parameters > Network Protocol Support (Continued)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
SNMP	Disabled.	Disabled.	Enabled. Disabled.
DHCP	Enabled.	Enabled.	Enabled. Disabled.
Ping	Enabled.	Enabled.	Enabled. Disabled.

TABLE D-9 Communication Parameter > Telnet Inactivity Timeout Time

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
Telnet Inactivity Timeout Time	Disabled.	Disabled.	Disabled or 60 seconds through 2700 seconds

TABLE D-10 Caching Parameters

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
Write-Back Cache	Enabled.	Enabled.	Disabled. Enabled.
Optimization for Random/Sequential	Sequential.	Sequential.	Sequential or Random
Periodic Cache Flush Time	Disabled.	Disabled.	Disabled. Continuous Sync 1/2 minute 1 minute 2 minutes 5 minutes 10 minutes

TABLE D-11 Peripheral Device Type Parameters

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	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 D-12 Host-Side and Drive-Side Parameters

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
Host-side Parameters			
Maximum Queued I/O Count	1024	1024	Auto, or 1-1024
LUNs per Host SCSI ID	32 (FC)	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)
In-band Management	Enabled	Enabled	Enabled. Disabled.
Host-side Parameters > Host Cylinder/ Head/ Sector Mapping Parameters			
Sector Ranges	Variable	Variable	32, 64, 127, 255, Variable sectors.
Head Ranges	64	64	64, 127, 255, Variable heads.
Cylinder Ranges	< 65536	< 65536	1024, 32768, 65536, Variable cylinders.
Host-side Parameters > Fibre Connection Option			
Fibre Connection Option	Loop only	N/A	FC Range of Values: Point to point only Loop only

TABLE D-12 Host-Side and Drive-Side Parameters (*Continued*)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
Drive-side Parameters			
Drive Motor Spin-Up	Disabled.	Disabled.	Do not change this parameter. Enabled. Disabled.
Disk Access Delay Time	15 seconds.	15 seconds.	Do not change this parameter. None to 75 seconds.
Drive I/O Timeout	30 seconds.	30 seconds.	500 milliseconds to 30 seconds.
Queue Depth	32	32	32
Maximum Tag Count	32 (FC) 4 (SATA)	32	1–256 or Disabled.
Periodic Drive Check Time	30 seconds.	30 seconds.	1/2 to 30 seconds.
Periodic SAF-TE and SES Check Time	30 seconds.	30 seconds.	Disabled to 60 seconds.
Periodic Auto-Detect Failure Drive Swap Check Time	Disabled.	Disabled.	5 to 60 seconds. Disabled.
Drive Predictable Failure Mode (SMART)	Detect and Clone + Replace.	Detect and Clone + Replace.	Detect Only Detect and Perpetual Clone. Detect and Clone + Replace.
Auto-Assign Global Spare Drive	Disabled.	Disabled.	Enabled. Disabled.
Media Scan at Power-Up	Disabled.	Disabled.	Enabled. Disabled.

TABLE D-13 Other Configuration Parameters

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	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.
Mixed Drive Support	Enabled.	N/A	Enabled (FC and SATA Only) Disabled (FC and SATA Only)
Redundant Controller Parameters			
Cache Synchronization on Write-Through	Enabled.	Enabled.	Enabled. Disabled.
Controller Parameters			
Controller Name	Not Set.	Not Set.	Type a name.
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.

Default Peripheral Device Parameters

The following peripheral device parameters are available.

TABLE D-14 Peripheral Device Type Parameters (View and Edit Peripheral Devices)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	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			
Controller Failure	Enabled.	Enabled.	Enabled. Disabled.
BBU Low or Failed	Enabled.	Enabled.	Enabled. Disabled.
Power Supply Failed	Enabled.	Enabled.	Enabled. Disabled.
Fan Failure	Enabled.	Enabled.	Enabled. Disabled.
Temperature exceeds threshold	Shutdown Period: 30 minutes.	Shutdown Period: 30 minutes.	Disabled. Enabled. Shutdown periods from two minutes to one hour.
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.

TABLE D-14 Peripheral Device Type Parameters (View and Edit Peripheral Devices) (*Continued*)

User-Defined Parameter	FC and SATA Default Setting	SCSI Default Setting	Range of Values
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	Default 95° C.	Default 95° C.	Disable, 50–100° C.
Lower Trigger Threshold for CPU Temperature Events	Default 0° C.	Default 0° C.	Disable, 0–20° C.
Upper Trigger Threshold for Board Temperature Events	Default 85° C.	Default 85° C.	Disable, 50–100° C.
Lower Trigger Threshold for Board Temperature Events	Default 0° C.	Default 0° C.	Disable, 0–20° C.

Default System Functions

The following system function parameters are available.

TABLE D-15 System Function Parameters

User-Defined Parameter	FC and SATA 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.

Parameter Defaults That Must Not Change

Do not change the following default parameters unless directed to do so by Technical Support.

TABLE D-16 Default Parameters That Must Not Change (unless directed to do so by Technical Support)

Parameter	Default Setting
Fault Management	
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).
Network Protocol Support	
PriAgent	Enabled.

Event Messages

This appendix lists the following event messages:

- “Controller Events” on page 354
 - “Controller Alerts” on page 354
 - “Controller Warnings” on page 357
 - “Controller Notifications” on page 358
- “Drive Events” on page 359
 - “Drive Alerts” on page 359
 - “Drive Warnings” on page 362
 - “Drive Notifications” on page 363
- “Channel Events” on page 363
 - “Channel Alerts” on page 363
 - “Channel Notifications” on page 367
- “Logical Drive Events” on page 367
 - “Logical Drive Alerts” on page 368
 - “Logical Drive Notifications” on page 371
- “General Target Events” on page 375
 - “SAF-TE Device Events” on page 375
 - “Controller Self-Diagnostic Events” on page 377
 - “I²C Device Events” on page 378
 - “SES Device Events” on page 378
 - “General Peripheral Device Events” on page 380

There are three categories of events as shown in TABLE E-1:

TABLE E-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 Escape to clear the message.
Notification	Informational message sent from controller firmware; press Escape to clear the message.

Controller Events

The controller records all array events during power on; it records up to one thousand events.

Note – Powering off or resetting the controller automatically deletes all recorded event log entries.

Controller event messages include the following:

Controller Alerts

Controller ALERT: Redundant Controller Failure Detected

A failure in a dual-redundant configuration has occurred and the other controller is managing all controller functions. When you reset the active controller or power-cycle the array, the failed controller is restarted and the following alert is displayed.

Controller ALERT: Controller Unrecoverable Error 000n [followed by code trap data]

This message is displayed *each time* a failed controller is restarted after a redundant controller failure and *does not* indicate a new controller failure. The date and time of the event message indicates the time the controller was restarted, not the time of the failure.

Note – If the unrecoverable error recurs, clear the core only on the advice of your support representative.

Controller SDRAM ECC Multi-bits Error Detected

SDRAM ECC multi-bit error.

Controller SDRAM ECC Single-bit Error Detected

SDRAM ECC single-bit error.

Controller SDRAM Parity Error Detected

SDRAM parity error.

Controller PCI Bus Parity Error Detected

PCI bus parity error.

Controller ALERT: Power Supply Unstable or NVRAM Failed

Power supply unstable, NVRAM has failed, firmware update failure, or incorrect configuration (for example, a controller combined with the wrong backplane type).

```
Memory Not Sufficient to Fully Support Current Config.
```

Memory capacity not sufficient to support current configuration.

```
CHL:_ FATAL ERROR (_)
```

CHL:_ channel fatal error.

```
Controller ALERT: +12V Low Voltage Detected (_)
```

This event message indicates that voltage dropped below the low voltage threshold (shown in parentheses).

Controller Warnings

```
WARNING: BBU Absent or Failed! Correct It and Reset Ctrl to Take Effect
```

As above.

```
WARNING: Controller BBU Absent or Failed!
```

As above.

```
WARNING: Controller BBU Failure Detected!
```

As above.

```
WARNING: Controller BBU Not Fully Charged !
```

Battery under charge and unable to support a configuration change. It is normal to see this message periodically as the battery regularly discharges and recharges.

```
WARNING: Controller BBU Thermal Shutdown/Enter Sleep-Mode !
```

There is a temperature sensor on the charger board. The upper threshold is 65° C. The controller will resume charging when normal temperature is restored.

```
Force Controller Write-Through on Trigger Cause !
```

The controller has been forced to adopt a safe caching mode on event-trigger conditions. The safety trigger can shutdown the controller or change the caching mode. The trigger causes for safety mechanisms are user-configurable, including battery condition, overheated board temperature, or peripheral device failure.

Controller Notifications

```
Controller NOTICE: NVRAM Factory Defaults Restored
```

Firmware settings have been restored to factory defaults. Options for restoring defaults are not available to users and are only reserved for qualified engineers.

```
Controller BBU Present !
```

BBU has once been removed and is now installed.

```
Controller BBU Back On-Line !
```

Battery once was absent or failed, and is now restored to normal functionality; charging has resumed.

```
Controller BBU Fully Charged !
```

The battery has been recharged. It is normal to see this message periodically as the battery regularly discharges and recharges.

```
Memory is Now Sufficient to Fully Support Current Config.
```

Memory is now sufficient to support current configuration.

```
Force Controller Write-Through on Trigger Cause
```

An environmental trigger event occurred that caused the controller to switch the cache policy to write-through (see the following message).

```
Controller Default Write Policy Restore
```

The environmental event that caused the cache policy switch (see above message) was corrected and the previous write policy was reestablished.

Drive Events

Physical drive event messages include the following:

Drive Alerts

```
CHL:_ ID:_ SCSI Target ALERT: Unexpected Select Timeout
```

Drive SCSI target select timeout. The specified hard drive cannot be selected by the controller. This can occur if a disk drive hangs a bus during the selection phase, resulting in a selection timeout. Since the bus is hung, the controller resets the bus which results in a Gross Phase error being reported by the controller. Drives are failed because the controller can no longer communicate properly. This is not necessarily due to a faulty disk drive but can be the result of any connection to the specific SCSI bus including the disk drive, controller, cable or I/O module.

```
CHL:_ ID:_ SCSI Target ALERT: Gross Phase/Signal Error Detected
```

Drive-side SCSI phase/signal abnormality detected. This can occur if a disk drive hangs a bus during the selection phase, resulting in a selection timeout. Since the bus is hung, the controller resets the bus which results in a Gross Phase error being reported by the controller. Drives are failed because the controller can no longer communicate properly. This is not necessarily due to a faulty disk drive but can be the result of any connection to the specific SCSI bus including the disk drive, controller, cable or I/O module.

```
CHL:_ ID:_ SCSI Target ALERT: Unexpected Disconnect Encountered
```

Drive-side SCSI target unexpected disconnect detected.

```
CHL:_ ID:_ SCSI Target ALERT: Timeout Waiting for I/O to Complete
```

Drive-side SCSI target I/O timeout. Possible drive-side cabling/termination and canister connection abnormal or drive malfunctioning.

```
CHL:_ ID:_ SCSI Target ALERT: SCSI Parity/CRC Error Detected
```

SCSI parity/CRC error detected while communicating with the specified hard drive.

```
CHL:_ ID:_ SCSI Drive ALERT: Unexpected Drive Not Ready (__B)
```

Drive installed does not respond with "Ready."

```
CHL:_ ID:_ SCSI Drive ALERT: Bad Block Encountered -  
____Block_number (Sense_key Sense_code)
```

Hard drive media error reported. A bad block is encountered in the specified hard drive. The RAID controller will ask the hard drive to retry. If the host attempts a read to this location, a "media error" status will be returned. If it attempts a write, the block will be recovered and the "recovered" message will be displayed.

```
CHL:_ ID:_ SCSI Drive ALERT: Drive HW Error (Sense_key Sense_code)
```

Drive-Side SCSI drive unrecoverable hardware error reported.

```
CHL:_ ID:_ SCSI Target ALERT: Unit Attention Received (Sense_key  
Sense_code)
```

Unit attention received on the SCSI drive target.

```
CHL:_ ID:_ SCSI Drive ALERT: Aborted Command (Sense_key  
Sense_code)
```

SCSI drive aborted command reported.

```
CHL:_ ID:_ SCSI Drive ALERT: Unexpected Sense Received (Sense_key  
Sense_code)
```

Drive-side SCSI drive unexpected sense data received.

Note – The three-digit code in parentheses 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

```
CHL:_ ID:_ SCSI Drive ALERT: Block Successfully Reassigned -  
Block_number (Sense_key Sense_code)
```

Rewrites attempted and bad blocks have been successfully reassigned.

```
CHL:_ ID:_ SCSI Drive ALERT: Block Reassignment Failed -  
Block_number (Sense_key Sense_code)
```

Drive-side block reassignment failed. Drive will be considered as having media errors or failed.

```
CHL:_ ID:_ SCSI Target ALERT: Data Overrun/Underrun Detected
```

Drive-side SCSI target data overrun or underrun detected.

```
CHL:_ ID:_ SCSI Target ALERT: Negotiation Error Detected
```

Drive-side SCSI target sync/wide negotiation abnormality detected.

```
CHL:_ ID:_ SCSI Target ALERT: Invalid Status/Sense Data Received  
(Sense_key Sense_code)
```

Drive-side SCSI invalid status/sense data received from target.

```
CHL:_ ID:_ ALERT: Redundant Path for CHL:_ ID:_ Failure Detected
```

Disconnection with the pair loop of the loop connection where CHL:_ ID:_ resides may have occurred.

Drive Warnings

```
SMART-CH:_ ID:_ Predictable Failure Detected
```

The SMART detect function has detected a Recovered Error (0x01) check condition.

```
SMART-CH:_ ID:_ Predictable Failure Detected (TEST)
```

(Test Mode) This message appears when simulating the SMART detect function. This message shows that your drives support SMART functions.

```
SMART-CH:_ ID:_ Predictable Failure Detected-Starting Clone
```

SMART errors detected; a spare is conducted to rebuild and/or replace the faulty drive. This is done according to the preset scheme.

```
SMART-CH:_ ID:_ Predictable Failure Detected-Clone Failed
```

SMART errors detected and a spare is conducted to rebuild. The cloning process is halted due to power interruption or yet another member drive has failed. Interruption to array integration will halt the cloning process; for example, drive failure.

Drive Notifications

```
CHL:_ ID:_ SCSI Drive NOTICE: Scan SCSI Drive Successful
```

Scanning new/missing drives from a SCSI channel successful.

```
CHL:_ NOTICE: Fibre Channel Loop Connection Restored
```

CHL:_ loop connection restored.

```
CHL:_ ID:_ NOTICE: Redundant Path for CHL:_ ID:_ Restored
```

Alternate connection to the dual-ported device, CHL:_ ID:_ is restored.

Channel Events

Channel event messages include the following:

Channel Alerts

```
CHL:_ Drive SCSI Channel ALERT: Unexpected Select Timeout
```

Drive channel CHL:_ select timeout. The specified drive channel cannot be selected by the controller. The channel has been disconnected; or the mode, cabling, termination, or canister for the channel is out of order.

```
CHL:_ RCC Channel ALERT: Gross Phase/Signal Error Detected
```

Gross phase/signal error found on the channel path used for redundant controller communications. This can occur if a disk drive hangs a bus during the selection phase, resulting in a selection timeout. Since the bus is hung, the controller resets the bus which results in a Gross Phase error being reported by the controller.

Drives are failed because the controller can no longer communicate properly. This is not necessarily due to a faulty disk drive but can be the result of any connection to the specific SCSI bus including the disk drive, controller, cable or I/O module.

```
CHL:_ Drive SCSI Channel ALERT: Gross Phase/Signal Error Detected
```

Unexpected disconnect detected on the channel path used for redundant controller communications. This can occur if a disk drive hangs a bus during the selection phase, resulting in a selection timeout. Since the bus is hung, the controller resets the bus which results in a Gross Phase error being reported by the controller. Drives are failed because the controller can no longer communicate properly. This is not necessarily due to a faulty disk drive but can be the result of any connection to the specific SCSI bus including the disk drive, controller, cable or I/O module.

```
CHL:_ Drive SCSI Channel ALERT: Unexpected Disconnect Encountered
```

Unexpected disconnect detected on the drive channel CHL:._

```
CHL:_ RCC Channel ALERT: Timeout Waiting for I/O to Complete
```

I/O timeout on the channel path used for redundant controller communications. Possible channel path cabling/termination and canister connection abnormal or malfunctioning.

```
CHL:_ Drive SCSI Channel ALERT: Timeout Waiting for I/O to Complete
```

I/O timeout on the drive channel path CHL:._

```
CHL:_ RCC Channel ALERT: SCSI Parity/CRC Error Detected
```

SCSI parity/CRC error detected on the channel path used for redundant controller communications.

```
CHL:_ SCSI Drive Channel ALERT: SCSI Parity/CRC Error Detected
```

SCSI parity/CRC error detected on the drive channel path CHL:._

CHL:_ RCC Channel ALERT: Unit Attention Received

Unit attention received on the channel path used for redundant controller communications.

CHL:_ SCSI Drive Channel ALERT: Unit Attention Received

Unit attention received on the drive channel CHL:_.

CHL:_ RCC Channel ALERT: Data Overrun/Underrun Detected

Data overrun or underrun detected on the channel path used for redundant controller communications.

CHL:_ Drive SCSI Channel ALERT: Data Overrun/Underrun Detected

Data overrun or underrun detected on the drive channel CHL:_.

CHL:_ RCC Channel ALERT: Negotiation Error Detected

SCSI target sync/wide negotiation abnormality detected on the channel path used for redundant controller communications.

CHL:_ Drive SCSI Channel ALERT: Negotiation Error Detected

SCSI target sync/wide negotiation abnormality detected on the drive channel CHL:_.

CHL:_ RCC Channel ALERT: Invalid Status/Sense Data Received

Invalid status/sense data received on the channel path used for redundant controller communications.

```
CHL:_ Drive SCSI Channel ALERT: Invalid Status/Sense Data Received
```

Invalid status/sense data received on the drive channel CHL:_.

```
CHL:_ SCSI Host Channel Alert: SCSI Bus Reset Issued
```

Host SCSI bus CHL:_ reset issued.

```
CHL:_ ALERT: Redundant Loop Connection Error Detected on ID:_
```

One of the dual loop members may have failed or been disconnected. Make sure all channels are properly connected and topological configuration properly set.

```
CHL:_ SCSI Host Channel ALERT: SCSI Channel Failure
```

Specific drive channel CHL:_ may have failed or disconnected.

```
CHL:_ ALERT: Fibre Channel Loop Failure Detected
```

Fibre channel loop failure is detected.

```
CHL:_ ALERT: Redundant loop for CHL:_ Failure Detected
```

The pair loop of CHL:_ has failed.

```
CHL:_ ALERT: Redundant Path for CHL:_ ID:_ Expected but Not Found
```

Disconnection with the pair loop of the loop connection where CHL:_ ID:_ resides may have occurred.

Channel Notifications

```
CHL:_ LIP( ) Detected
```

Fibre Loop LIP issued on CHL:_.

```
CHL:_ SCSI Host Channel Notification: SCSI Bus Reset Issued
```

SCSI bus reset issued on CHL:_

```
CHL:_ NOTICE: Fibre Channel Loop Connection Restored
```

CHL:_ loop connection restored.

Logical Drive Events

Logical drive event messages often begin with the letters LG, an abbreviation for Logical Group that identifies the logical drive number to which the message applies.

Logical drive event messages include the following:

Logical Drive Alerts

```
LG:_ Logical Drive ALERT: CHL:_ ID:_ SCSI Drive Missing
```

A member hard drive in the specified logical drive is missing.

```
LG:_ Logical Drive ALERT: CHL:_ ID:_ Drive Failure
```

A member hard drive in the specified logical drive has failed.

```
LG:_ Logical Drive ALERT: Creation Aborted
```

The creation process of logical drive LG_ is aborted.

```
LG:_ Logical Drive ALERT: Creation Failed
```

The creation process of logical drive LG_ has failed.

```
LG:_ Logical Drive ALERT: Initialization Failed
```

The initialization process of logical drive LG_ has failed.

```
LG:_ Logical Drive ALERT: Expansion Failed
```

A member drive or other hardware failed, bad blocks were encountered, or the user cancelled the operation.

```
LG:_ Logical Drive ALERT: Rebuild Aborted
```

The rebuilding operation on logical drive LG_ is aborted.

```
LG:_ Logical Drive ALERT: Rebuild Failed
```

The rebuilding operation on logical drive LG_ has failed. It can be the result of the following conditions:

- The rebuild has been canceled by a user.
- The drive used for rebuild failed during the rebuild process.
- Bad blocks have been encountered on yet another member drive during the rebuild process.

```
LG:_ Logical Drive ALERT: Parity Regeneration Failed
```

During the parity-regeneration process, one member drive failed.

```
LG:_ ALERT: CHL:_ ID:_ Media Scan Failed
```

Media scan failed on the member of logical drive LG_ (CHL_, ID_)

```
LG:_ ALERT: CHL:_ ID:_ Media Scan Aborted
```

Media scan canceled by user or aborted on the member of logical drive LG_ (CHL_, ID_) for array integrity concerns.

```
LG:_ Logical Drive ALERT: CHL:_ ID:_ Clone Failed
```

Cloning process failed when proceeding with the member of logical drive LG_, CHL_, ID_.

```
LG:_ Logical Drive ALERT: Logical Drive Bad Block Table FULL
```

Bad block table full with entries found in logical drive LG_.

```
LG:_ Logical Drive ALERT: Logical Drive Bad Block Table BAD
```

Logical drive LG_ bad block table has failed.

```
LG:_ Logical Drive ALERT: Logical Drive On-Line Init Table Bad
```

The table storing information about online initialization progress of logical drive LG_ has failed.

```
LG:_ Logical Drive ALERT: Logical Drive Block Marked BAD
```

One or more bad blocks found during media scan, parity regeneration, or normal write check operations on logical drive LG_. The block was marked BAD so that the host can deal with it appropriately without risking data..

```
CHL:_ ID:_ ALERT: Media Scan Bad Block Unrecoverable - 0x_____
```

Bad blocks found irrecoverable even after the controller attempts to rewrite data onto it. Block address is 0x_____.

```
CHL:_ ID:_ SCSI Drive ALERT: Bad Block Encountered - _____ (___)
```

Bad blocks found on drive CHL_ ID_. Block address is _____ (___).

```
CHL:_ ID:_ ALERT: Bad Block Encountered - 0x_____
```

Bad blocks encountered on CHL_ ID_. Block address is 0x_____.

```
LG:_ Logical Drive ALERT: Fatal Fail
```

A Fatal Fail condition occurred on Logical Drive LG:_.

```
LG:_ Logical Drive ALERT: Cache Data Purged
```

A Fatal Fail condition occurred on LG:_ while under load. Data in cache was discarded.

Logical Drive Notifications

```
LG:_ Logical Drive NOTICE: Starting Creation
```

A message related to “Immediate Array Availability.” The controller/subsystem starts assembling member hard drives into a logical drive, LG_. The logical drive will be ready for I/O when creation is done, and the controller/subsystem will find appropriate time to conduct parity initialization.

```
LG:_ Logical Drive NOTICE: Starting On-Line Initialization
```

A message related to “Immediate Array Availability.” The controller/subsystem starts initializing the logical drive. “On-Line” means the array is immediately accessible, even before the initialization process is completed.

```
LG:_ Logical Drive NOTICE: Starting Off-Line Initialization
```

“Off-Line” means the array is accessible only after the initialization process is completed. The controller/subsystem starts initializing the logical drive once the array is configured.

```
On-Line Initialization of Logical Drive_ Completed
```

A message related to “Immediate Array Availability.” Initialization of logical drive, LG_, is completed.

```
Off-Line Initialization of Logical Drive_ Completed
```

Initialization of logical drive LG_ is completed.

```
Creation of Logical Drive_ Completed
```

A message related to “Immediate Array Availability.” Member hard drives have been successfully grouped into a logical drive, LG_. The logical drive is now ready for I/O, and the controller/subsystem will find appropriate time to complete parity initialization.

```
LG:_ Logical Drive NOTICE: Starting Rebuild
```

The rebuild process on logical drive LG_ has started.

```
Rebuild of Logical Drive_ Completed
```

Logical drive LG_ has been successfully rebuilt.

```
LG:_ Logical Drive NOTICE: Starting Parity Regeneration
```

Start regenerating parity data of logical drive LG_.

```
Parity Regeneration of Logical Drive_ Completed
```

Parity regeneration on logical drive_ completed.

```
LG:_ Logical Drive NOTICE: Starting On-Line Expansion
```

Start expanding the logical drive. Data re-striping is carried out later in the background.

```
LG:_ Logical Drive NOTICE: Starting Off-Line Expansion
```

Start expanding the logical drive. Data re-striping is carried out immediately.

```
On-Line Expansion of Logical Drive_ Completed
```

Logical drive expansion completed.

```
Off-Line Expansion of Logical Drive_ Completed
```

Logical drive expansion completed.

```
LG:_ Logical Drive NOTICE: Starting Add Drive Operation
```

Expansion “by adding new drive” has started.

```
Add Drive to Logical Drive_ Completed
```

The expansion “by adding new drive” is completed.

```
LG:_ Logical Drive NOTICE: Add SCSI Drive Operation Paused
```

The expansion process is halted because of one of the following events:

- Logical drive expansion canceled by user.
- One of the member drives failed during the “Add Drive” operation.
- Bad blocks encountered on one of the member drives.
- Hardware failure.

```
LG:_ Logical Drive NOTICE: Continue Add Drive Operation
```

The “Add Drive” process had once been paused and is now resumed. The target logical drive has been restored to its previous status, and the system can continue with the “Add Drive” operation.

```
LG:_ Logical Drive NOTICE: CHL:_ ID:_ Starting Clone
```

This message is displayed when a member drive is manually cloned to a spare, or a spare is automatically applied to clone a faulty member on SMART-detected errors.

```
LG:_ Logical Drive NOTICE: CHL:_ ID:_ Copy and Replace Completed
```

This message is displayed when a spare is used to replace a member drive suspected of imminent faults. This message indicates completion of cloning.

```
LG:_ Logical Drive NOTICE: CHL:_ ID:_ Clone Completed
```

Cloning process on the member of LG_, CHL_, ID_, has been completed.

```
LG:_ NOTICE: CHL:_ ID:_ Starting Media Scan
```

Starting media scan on the members of logical drive LG_. Each member being scanned is recognized by its channel and channel ID. This message is shown when member drives are being scanned.

```
LG:_ NOTICE: Media Scan of CHL:_ ID:_ Completed
```

Media scan is completed on a member drive (CHL:_ and ID:_).

```
LG:_ Logical Drive NOTICE: Logical Drive Block Recovered
```

Bad block recovered by rewriting data onto it.

```
CHL:_ ID:_ NOTICE: Media Scan Bad Block Recovered-ox_____
```

Bad block recovered by rewriting data onto it. Block address is 0x_____ .

```
LG:_ NOTICE: Inconsistent Parity Encountered Block _____
```

Inconsistent parity of logical drive LG:_, found on block address _____.

```
CHL:_ ID:_ SCSI Drive NOTICE: Scan SCSI Drive Successful
```

Scanning new/missing drives on a SCSI channel successful.

```
CHL:_ NOTICE: Fibre Channel Loop Connection Restored
```

CHL:_ loop connection restored.

```
CHL:_ ID:_ NOTICE: Redundant Path for CHL:_ ID:_ Restored
```

Alternate connection to the dual-ported device, CHL:_ ID:_ is restored.

General Target Events

General target event messages include SAF-TE device messages, controller self-diagnostic messages, I²C messages, SES device messages, and general peripheral device messages.

SAF-TE Device Events

SAF-TE device event messages include the following:

SAF-TE Device Alerts

```
SAF-TE Device (_) ALERT: Power Supply Failure Detected (Idx: __)
```

Power supply (device __; device ID__) failure detected by enclosure management.

```
SAF-TE Device (_) ALERT: Cooling Fan Not Installed (Idx: _)
```

Fan (_) is missing from device slot.

```
SAF-TE Device (_) ALERT: Cooling Fan Failure Detected (Idx: _)
```

Cooling fan_ has failed.

```
SAF-TE Device (_) ALERT: Elevated Temperature Alert
```

Temperature exceeding threshold on SAF-TE device_.

```
SAF-TE Device (_) ALERT: UPS Power Failure Detected
```

UPS power failure detected through SAF-TE device_.

SAF-TE Device Notifications

```
SAF-TE Device ( ) NOTICE: Fan Back On-Line (Idx: _)
```

Device _ failed fan back on-line (device ID:_).

```
SAF-TE Device ( ) NOTICE: Temperature Back to Non-Critical Levels
```

Temperature restored to within safety range.

```
SAF-TE Device ( ) NOTICE: Power Supply Back On-Line (Idx:_)
```

Power supply module_ back on-line (device ID:_), reported through SAF-TE device (_).

```
SAF-TE Device ( ) NOTICE: UPS Power Back On-Line
```

UPS power restored, reported through SAT-TE device (_).

Controller Self-Diagnostic Events

Controller self-diagnostic event messages include the following:

Controller Self-Diagnostic Alerts

```
Peripheral Device ALERT: Controller FAN_ Not Present or Failure  
Detected
```

This event refers to the cooling fan in front bezel. Check cable connection and see if a fan has failed.

```
ALERT: Controller FAN_ Low Speed Detected ( _ RPM)
```

This message refers to the cooling fan in controller's front bezel. Low rotation speed detected.

```
ALERT: +3.3V Low Voltage Detected ( __. _V)
```

The detected +3.3V voltage source is now lower than the preset threshold.

Controller Self-Diagnostic Notifications

```
Board1 Cold Temperature Back to Non-Critical Levels
```

Main board temperature restored to within safety range.

```
Controller NOTICE: Redundant Controller Firmware Updated
```

Firmware updated for both controllers in the dual-controller configuration.

```
+12V Upper Voltage Back within Acceptable Limits ( _._V)
```

+12V restored to within upper safety threshold.

```
+12V Lower Voltage Back within Acceptable Limits (._.V)
```

+12V restored to within lower safety threshold.

I²C Device Events

I²C event messages include the following:

I2C Device Notifications

```
NOTICE: Fan Module _ Back On-Line (Fan_ _RPM)
```

Fan module _ back online (Fan_, _RPM).

```
NOTICE: Controller Fan_ Back On-Line ( _RPM)
```

Controller fan_ (fan on the front bezel) back online (_RPM).

SES Device Events

SES event messages include the following:

SES Device Alerts

```
SES (C_ I_) Power Supply_: Device Not Supported !
```

Unrecognizable device type on C_ I_. (SES.)

```
SES (C_ I_) Cooling Fan_: Device Not Supported !
```

Unrecognizable device type on C_ I_. (SES.)

```
SES (C_ I_) Temp Sensor _: Device Not Supported !
```

As above.

```
SES (C_ I_) device not supported
```

As above.

```
SES (C_ I_) UPS_: Device Not Supported !
```

As above.

```
SES (C_ I_) UPS device not supported!
```

As above.

```
SES(C2 I12)Voltage sensor 3:Under Voltage Critical
```

A voltage sensor has detected a critical under-voltage condition.

```
SES(C2 I12)Voltage sensor 2:Power Supply Failure Detected
```

A voltage sensor has detected a power supply failure. The value can be either 1 or 2, depending on which power supply was detected.

SES Device Notifications

```
SES (C_ I_) Cooling Fan _: Fan Back On-Line
```

Cooling fan_ back online, reported through SES (C_I_)

```
SES (C_ I_) Temp Sensor _: Temperature Back to Non-Critical Levels
```

Temperature restored to within safety range; detected by SES (C_ I_) sensor.

```
SES (C_ I_) Power Supply _: Power Supply Back On-Line
```

Power supply_ back online, reported through SES (C_I_).

```
SES (C_ I_) UPS _: UPS Power Back On-Line
```

UPS_ power back online, reported through SES (C_I_).

General Peripheral Device Events

General peripheral device event messages include the following:

General Peripheral Device Alerts

```
Peripheral Device ALERT: Power Supply_ Failure Detected
```

Power supply failure detected.

```
Peripheral Device ALERT: Power Supply_ Not Present
```

Power supply module installed but not present now.

```
ALERT: Low Power Supply_ Voltage Detected (._.V)
```

Low voltage detected from power supply module __.

```
Peripheral Device ALERT: FAN_ Not Present
```

Fan module_ installed but not present now.

Peripheral Device ALERT: FAN_ Failure Detected

Fan_ failure detected.

Peripheral Device ALERT: Low FAN_ Speed Detected (__ RPM)

Fan module_ low rotation speed detected (__RPM).

Peripheral Device ALERT: CPU Cold Temperature Detected (._ _C)

CPU temperature dropped below preset threshold.

Peripheral Device ALERT: Elevated Temperature Alert

Elevated ambient temperature within chassis.

Peripheral Device ALERT: Temperature Sensor _ Failure Detected

Peripheral device temperature sensor_ failure detected.

Peripheral Device ALERT: Temperature Sensor _ Not Present

Peripheral device temp sensor_ installed but not present now.

Peripheral Device ALERT: Cold Temperature _ Detected (_C)

Cold temperature detected by device_ (_C).

Peripheral Device ALERT: UPS_ AC Power Failure Detected

UPS AC power failure.

Peripheral Device ALERT: UPS_ Battery Failure Detected

UPS battery failure.

```
Peripheral Device ALERT: Boardx Hot Temperature Detected(56.0(C)
(Secondary)
```

A board temperature event has been detected for Board 1, Board 2, or a CPU.

General Peripheral Device Notifications

```
Peripheral Device NOTICE: Fan Back On-Line
```

Fan module back online.

```
NOTICE: FAN_ Back On-Line
```

Fan module_ back online.

```
NOTICE: Fan_ is present
```

Fan module_ is present.

```
NOTICE: Fan_ Back On-Line (_ RPM)
```

Fan module_ back online (_ RPM).

```
NOTICE: Temperature _ Back to Non-Critical Levels
```

Temperature detected through sensor_ restored to within safety range.

```
NOTICE: Temperature _ Back to Non-Critical Levels (_ C)
```

Temperature detected through sensor_ restored to within safety range (_ C).

```
NOTICE: Temperature _ is Present
```

Temperature sensor_ is present

```
Power Supply _ Back-Online
```

Power supply module_ back online.

Power Supply _ Back-Online (._V)

Power supply module_ back online (._V).

UPS Connection Detected

UPS connection detected.

UPS AC Power Restored

UPS AC Power Restored.

UPS Battery Restored to Safe Level

UPS battery charge restored to within safe levels.

Peripheral Device Notice: UPS _ AC Power Back On-Line

UPS _ AC power back online.

Peripheral Device Notice: UPS _ Battery Back On-Line

UPS _ battery back online.

UPS Connection is Absent

Controller/subsystem lost connection with UPS device.

Warning: UPS AC Power-Loss Detected

UPS AC power loss detected.

UPS Battery Low - _%

UPS battery found under-charge, charge percentage _%.

SNMP Management Using HP OpenView

HP OpenView is a complete enterprise network management, maintenance, and monitoring solution. Included with the HP OpenView software are tools for sending SNMP GET and SET messages and tools for listening for SNMP traps.

This appendix discusses using HP OpenView for SNMP management with Sun StorEdge 3000 family arrays. Topics covered include:

- “Loading MIBs” on page 386
- “Configuring Events” on page 389
- “Viewing and Setting System Group Objects” on page 391
- “Viewing Enterprise Objects” on page 394

Loading MIBs

By itself, OpenView can listen for and dispatch SNMP traps. However, MIBs are supplied to make the best use of the management feature.

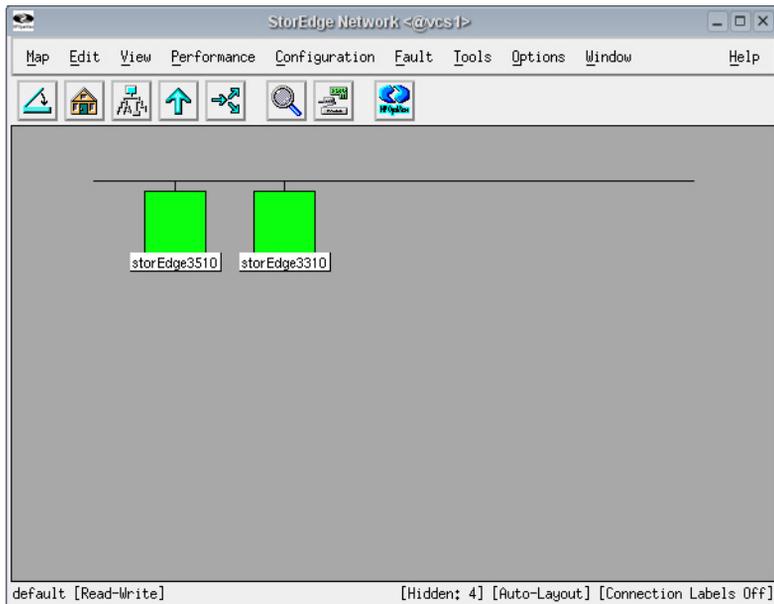
Note – The agent .ini file must be configured with trap destinations in order to receive traps. See “To Create and Save the NPC Configuration File (agent.ini)” on page 227.

It is assumed that HP OpenView has discovered the Sun StorEdge 3000 family nodes. Refer to your HP OpenView documentation for details on node discovery.

On Solaris, type the following command to launch OpenView.

```
# /opt/OV/bin/ovw &
```

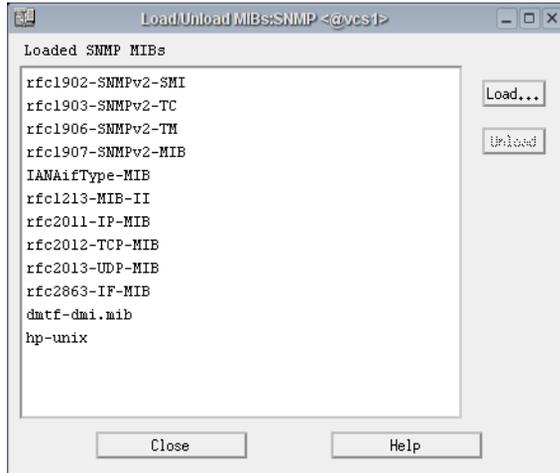
From the Root dialog, navigate to the network segment on which the Sun StorEdge nodes reside. The following screen shows a typical network segment displayed by OpenView with two Sun StorEdge nodes.



To load MIBs, perform the following steps.

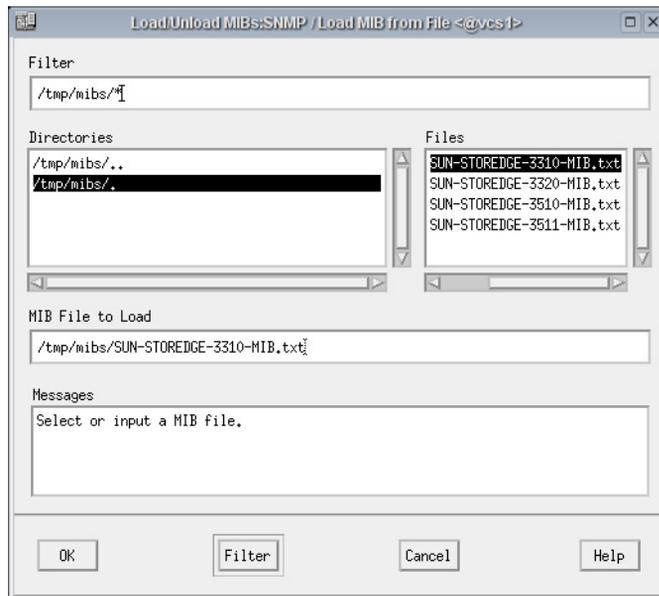
1. Choose Options →Load/Unload MIBs: SNMP.

The following dialog is displayed.



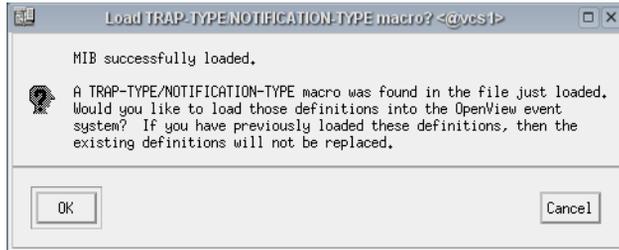
2. Click Load...

The following dialog is displayed.



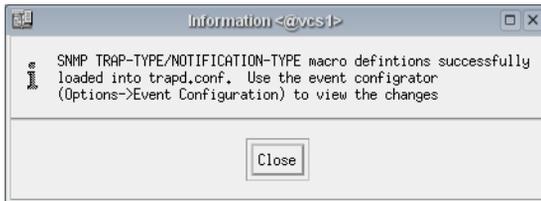
3. Select the MIBs to be loaded and click OK.

Since the MIB contains TRAP/NOTIFICATION information, OpenView will detect this. A dialog is displayed requesting confirmation to load the definitions into the OpenView event system.



4. Click OK to continue.

If the definitions load successfully, an information dialog is displayed.



5. Click Close.

If more than one array product type is being monitored (for example, a Sun StorEdge 3320 SCSI array and a Sun StorEdge 3510 FC array), you can load multiple MIBs.

6. Repeat the preceding steps for each MIB to be loaded.

Note – You can also unload MIBs by selecting them from the list and clicking “Unload”.

Note – MIBs are distributed with firmware patches. You might have to reload MIBs if firmware changes impact the SNMP agent on the array controller.

Configuring Events

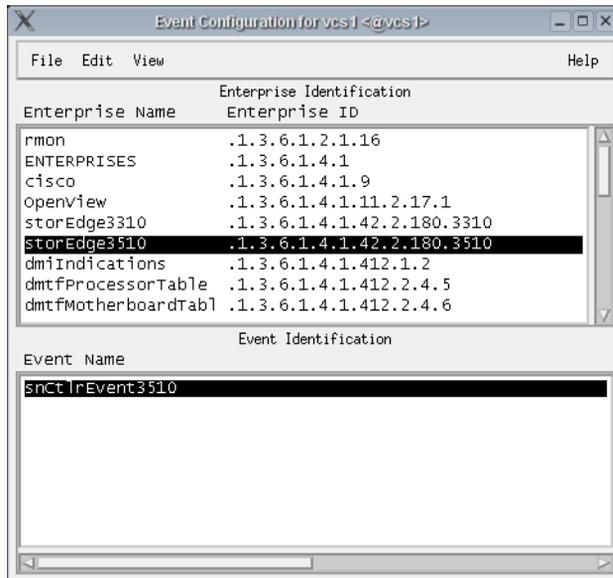
Since the MIBs contain information on traps, it is possible to configure these events. Events can be displayed in the Alarm browser, in a pop-up window, forwarded to other hosts, and logged to files. Refer to your HP OpenView documentation for details.

If the MIBs have been loaded successfully, the Enterprise Name and Enterprise ID will be displayed in the Event Configuration dialog.

To view and modify events, perform the following steps.

1. Choose Options →Event Configuration.

The Event Configuration dialog is displayed.



2. To modify an event, select the Enterprise Name in the upper pane and double-click on the Event Name in the lower pane.

The Modify Event dialog is displayed.

Event Configurator / Modify Event for vcs1 <@vcs1>

Event Name: snCtrlrEvent3510 Event Type: Enterprise Specific Event Object Identifier: .1.3.6.1.4.1.42.2.180.3510.0.1

Event Description

Long Descr.:
"There is a controller event occurred"
Variables:

Event Sources (all sources if list is empty)

Source: Add From Map Delete Delete All Add

Category: Status Alarms Forward Event Severity: Normal

Event Log Message: StorEdge 3510 Event: \$1

Pop-up Notification (Optional):

Command for Automatic Action (Optional):

OK Reset Cancel Help

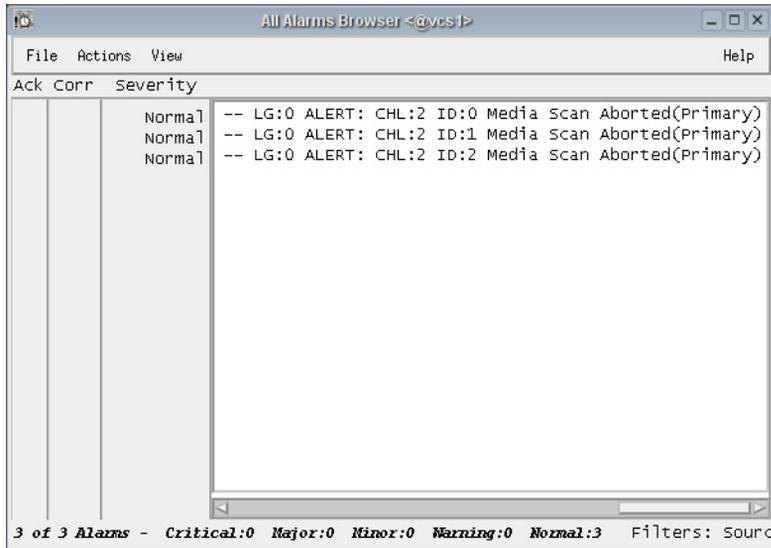
3. To display the event in the Alarm browser, select an appropriate category.

Since only one trap type is supported, all events will go into this category. In this example, a category of "Status Alarms" is used.

4. Enter a message for the event in the Event Log Message text box.

In this example, "\$1" represents the variable string of the SNMP trap. See the HP OpenView documentation for details on other "\$" variables available as part of the event message.

- To view alarms, from the main menu bar, choose **Fault** → **Alarms**.
The following dialog is displayed.



Viewing and Setting System Group Objects

SNMP must be enabled on the array to view and set system group objects. From the RAID firmware Main Menu, choose “view and edit Configuration parameters → Communication Parameters → Network Protocol Support → SNMP” and verify that SNMP is enabled.

- From **HP OpenView**, browse the system group objects for a node by selecting the node on the segment map.
- Choose **Tools** → **SNMP MIB Browser**.
Confirm that the correct Name or IP Address is displayed.
- Navigate to the following **MIB Object ID**:
`iso.org.dod.internet.mgmt.mib-2`

4. Select system from the list and click Start Query.

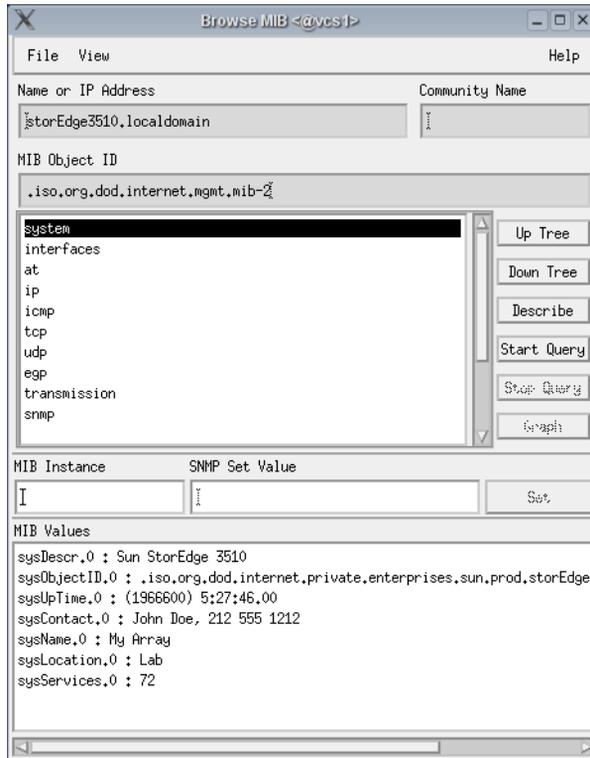


Read/Write values can be set from this dialog.

5. To set the system name, select `sysName.0` from the list and enter a new value in "SNMP Set Value".

6. Click Set.

A new query on the system group shows the new value.



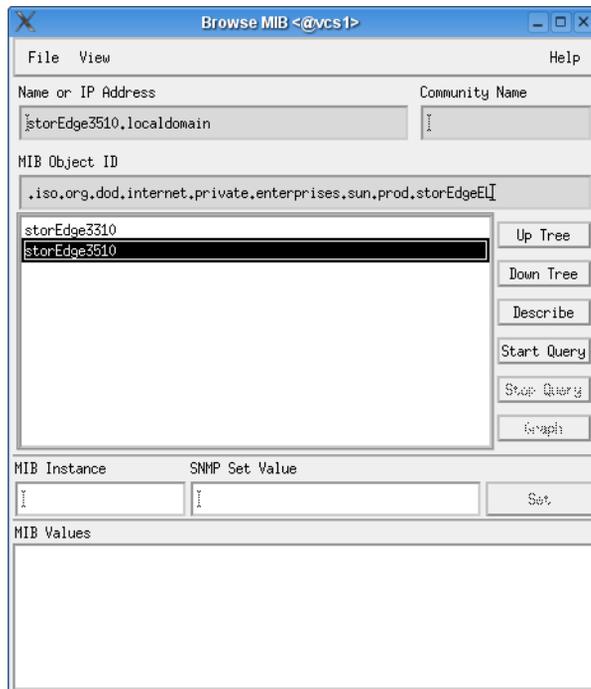
Viewing Enterprise Objects

SNMP must be enabled on the array to view enterprise objects. From the RAID firmware Main Menu, choose “view and edit Configuration parameters → Communication Parameters → Network Protocol Support → SNMP” and verify that SNMP is enabled.

Certain array information can be obtained by browsing the MIB. These objects are located under the following subtree:

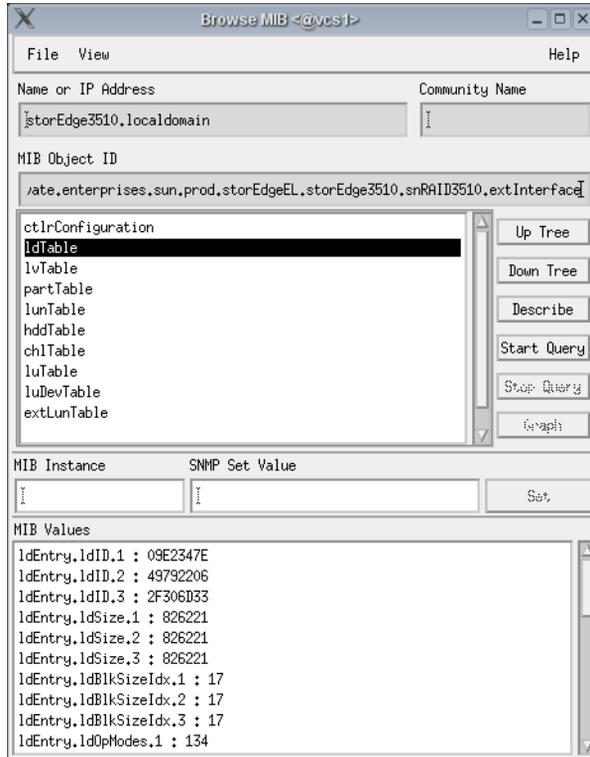
```
.iso.org.dod.internet.private.enterprises.sun.product.storEdgeEL
```

1. **Navigate to this object ID with the MIB browser and select the appropriate branch for the product.**



2. **Continue to browse to the `extInterface` node and select the desired item from the list.**

3. Click Start Query.



For details on OIDs and descriptions available for enterprise objects, see the MIB for your array.

Glossary

The glossary lists acronyms and defines RAID terms found throughout 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 firmware application to restart the rebuilding process.
- block striping** *See striping.*
- block striping with dedicated parity** (RAID 3) A technique that 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 Mbyte).
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.
CISPR	International Special Committee on Radio Interference.
DHCP	Dynamic Host Configuration Protocol.
disk mirroring	<i>See</i> mirroring (RAID1).
EMC	Electromagnetic compatibility.
EMU	Event monitoring unit.
Fabric	Fibre Channel network built around one or more switches.
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. 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.
Gbyte	(Gigabyte) 1024 Mbyte or 1,073,741,824 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	A group is a data object that enables multiple servers to be contained under a single category. Groups are similar in concept to domains, and enable you to organize servers.
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 RAID array remains powered on and operational.
ID	Identifier number.
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 storage device that consist of drives with no controllers.
LAN	Local area network.
LD	Logical drive.
logical drive	A section of disk storage space that is presented to the host operating system as a single physical drive. A logical drive might be located on one or more physical 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 a local disk drive.
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	(Low-Voltage Differential) 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 10/100BASE-T Ethernet port that is used to configure a RAID array.
Mbyte	(Megabyte) 1024 Kbyte or 1,048,576 bytes
media scan	A background process that continuously checks physical drives for bad blocks or other media errors.
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.
NDMP	Network Data Management Protocol.
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.
OBP	OpenBoot™ PROM (OBP). When you first start Solaris, it shows an OK prompt, which is the OBP. It is a command-line interface.

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. For RAID 1 configurations, data is compared with mirrored data, but since RAID 1 does not store parity, no automatic correction is possible.
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.
PID	Primary controller identifier number
PLA	(Programmable Logic Array) Offers flexible features for more complex designs.
PLD	(Programmable logic device) A generic term for an integrated circuit that can be programmed in a laboratory to perform complex functions.
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. 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.
RAS	(Reliability, Availability, and Serviceability) A variety of features and initiatives all designed to maximize equipment uptime and mean time between failures, minimize downtime and the length of time necessary to repair failures, and eliminate or decrease single points of failure in favor of redundancy.
read policy	A storage device parameter that determines whether the storage device holds 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	Enables the RAID controller to serve other I/O requests while rebuilding the logical drives. Priority ranges from low, which uses the controller's minimum resources to rebuild, to high, which uses the controller's maximum resources to complete the rebuilding process.

- SAN** (Storage Area Network) 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 physical conditions within an enclosure, and enable access to the status reporting and configuration features of the enclosure (such as indicator LEDs on the enclosure).
- 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 system 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 RAID array stores the states of drives, logical drives, and the controller in its nonvolatile memory. This information is retained across power interruptions.
- stripe size** The amount of data in kilobytes that is striped across each physical drive in a logical drive. Generally, large stripe sizes are more effective for arrays with sequential reads.
- striping** The storing of sequential blocks of incoming data on all the different physical 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	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 system 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 system 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

Numerics

- 1024 LUNs in loop mode with added host IDs (FC and SATA) 93
- 1024 LUNs, loop mode (FC and SATA) 163
- 128 LUNs in point-to-point mode (FC and SATA) 93
- 128 LUNs with added host IDs (SCSI) 58
- 64 LUNs default (SCSI) 58
- 64 LUNs, redundant point-to-point 164

A

- agent.ini file
 - creating for SNMP 226
 - SNMP parameters 229
- audible alarms, silencing 275
- auto-assigning global spare drives (FC and SATA) 245

B

- BAD drive status 297
- battery
 - charging status 12
 - operation 291
- baud rate, setting 212
- beeper, muting 275

C

- cache
 - battery support 293
 - default parameters 343
 - enabling write-back or write-through 231
 - status 12

- cache optimization (FC and SATA) 84
- cache optimization (SCSI) 52
- caching parameters 231
- capacities, maximum 25
- channel event messages 363
- channel status table 298
- channels
 - default parameters 22, 341
 - default settings (SCSI) 56
 - defining 313
 - deleting host IDs 201
 - drive IDs 202
 - ID ranges (FC and SATA) 94
 - setting host and drive (FC and SATA) 88
 - setting host and drive (SCSI) 56
 - setting termination (SCSI) 202
 - status table 200
 - transfer clock speed (SCSI) 203
 - transfer width (SCSI) 204
 - viewing and editing 199
- chip information 205
- clock speed (SCSI) 203
- cloning a failing drive 184
- cloning, viewing status 189
- COM port baud rate 212
- communication parameters 212
- concurrent host LUN connections, maximum 234
- concurrent rebuild 310
- configuration
 - saving NVRAM to disk (FC and SATA) 125
 - saving NVRAM to disk (SCSI) 80

- configuration parameters 211
 - configurations
 - adding physical drives to logical drives 137
 - creating 1024 LUNs in loop mode 163
 - DAS loop example (FC/SATA) 41
 - default parameters 341
 - default parameters that should not change 351
 - drive-side parameters 240
 - expanding the capacity of logical drives 133
 - identifying a failed drive 180
 - loop 97, 238
 - maximum number of drives 24
 - maximum number of logical drives and LUNs 24
 - other default parameters 347
 - point-to-point 238
 - point-to-point (FC and SATA) 97
 - redundant point-to-point configurations 164
 - restoring from disk 284
 - SAN example (FC/SATA) 35
 - saving to disk 281
 - single-bus configuration (SCSI) 173
 - split-bus configuration (SCSI) 173
 - summary of key steps 45
 - viewing logical drive (FC and SATA) 82
 - viewing logical drive (SCSI) 50
 - connections
 - loop 97
 - point-to-point 97
 - controller
 - assignment of logical drive to (FC and SATA) 108
 - assignment of logical drive to (SCSI) 70
 - date and time 251
 - event messages 354
 - failover 305
 - IDs 202
 - muting beeper 275
 - name 12
 - naming 248
 - over-temp shutdown 268
 - parameters 248
 - password validation timeout 249
 - rebuild priority 246
 - redundant controller parameters 247
 - resetting 278
 - restoring a configuration 284
 - saving configuration to disk 281
 - shutting down 279
 - unique identifier 250
 - upgrading firmware 302
 - voltage status 270
 - conventions and terminology 17
 - cylinder ranges, changing (FC and SATA) 98
 - cylinder ranges, changing (SCSI) 61
 - cylinder/head/sector mapping 236
- ## D
- DAS loop configuration 41
 - date and time
 - setting controller 251
 - default parameters that should not change, list of 351
 - default parameters, summary 339
 - device capacities, units of 19
 - device parameters, default 348
 - DHCP 213
 - default IP address 8
 - enabling and disabling 215
 - setting array IP address using 214
 - disk access delay time 241
 - disk array parameters 246
 - disk drive firmware 301
 - downloading firmware 301
 - drive channel defaults 22
 - drive channel IDs 202
 - drive check time, periodic 243
 - drive event messages 359
 - drive firmware 301
 - drive IDs (FC and SATA) 175
 - drive IDs (SCSI) 173
 - drive swap check time, periodic 244
 - drives
 - adding and deleting entries (SCSI) 179
 - adding to a logical drive 137
 - cloning a failing 184
 - cloning status 189
 - combining FC and SATA 87
 - default parameters 340, 345
 - flashing to identify good or bad 181
 - global or local spare identified 297
 - I/O timeout 241
 - identifying failed 180
 - local and global spare 5
 - maximum number per array 24
 - media scan for bad blocks 145

- motor spin-up 240
- number supported by RAID level 317
- periodic check time 243
- perpetual clone 186
- read/write testing 196
- reformatting 195
- replacing after cloning 184
- reserved space 197
- revision number, serial number, and disk
 - capacity 176
- scanning SCSI 179
- status table 172
- USED DRV drive status 297
- vendor ID 297
- viewing available (FC and SATA) 86
- viewing available (SCSI) 54

drive-side parameters 240

DRV FAILED status 294

DRV MISS status 294

Dynamic Host Configuration Protocol. See DHCP

E

email

- using SNMP to send event messages 226

error statistics, Fibre Channel 273

error statistics, loopback 273

event logs

- viewing 288

event messages 353

- channel 363
- controller 354
- drive 359
- logical drive 367
- sending using SNMP 226
- types of 353

event trigger operations 267

expansion of logical drives 133

external interface, enabling and disabling 238

F

failed drives, identifying 180

failover, controller 305

fan status, identifying 256

fans, identifying 263

fault management, default parameters 351

fault protection measures

cloning a failing drive 183

- perpetual clone 186
- SMART 183
- verifying written data 247

FC arrays, configuration and reconfiguration of 81

Fibre Channel error statistics 273

Fibre Connection options

- loop 238
- point-to-point 238

Fibre protocol 97

filter mode (FC and SATA) 121

filter type (FC and SATA) 120

firmware

- advanced features 326
- basic components 11
- enabling and disabling external interface of 238
- initial window 12
- main menu 15
- navigation keys 15
- upgrading 301

firmware downloads 301

firmware specifications 325

firmware window components 12

first-time configuration

- for FC and SATA arrays 81
- for SCSI arrays 49

first-time configuration summary 45

FTP, enabling and disabling 215

G

global spare drives 5

- adding 178
- auto-assigning (FC/SATA) 245
- deleting 178

H

head ranges, changing (FC and SATA) 98

head ranges, changing (SCSI) 61

host applications

- planning 4

host channel defaults 22

host channel IDs

- deleting 201

host configuration

- setting communication parameters 9

- host cylinder/head/sector mapping, changing (FC and SATA) 98
- host cylinder/head/sector mapping, changing (SCSI) 61
- host filter mode (FC and SATA) 121
- host filter type (FC and SATA) 120
- host filters
 - viewing and modifying information 168
- host filters (FC and SATA) 118
- host IDs
 - number supported (FC and SATA) 92
- host IDs, creating and adding (FC and SATA) 92
- host LUNs
 - adding WWN entries 168
 - deleting a mapping 165
 - mapping to partitions 160
 - maximum concurrent connections 234
 - number of tags reserved for 235
 - number supported per host ID 234
- host-side parameters 233
 - defaults 345
- host-side parameters (FC and SATA) 97
- HTTP, enabling and disabling 215
- HTTPS, enabling and disabling 215

I

- I/O count, maximum queued 233
- I/O timeout, configuring drive 241
- ID ranges
 - changing channel (FC and SATA) 94
- inactivity, telnet timeout 216
- in-band EI management
 - enabling and disabling 238
- INCOMPLETE status 294
- initialization mode, logical drive (FC and SATA) 104
- initialization mode, logical drive (SCSI) 66
- INITING status 294
- INVALID status 294
- IP address
 - setting 8, 213
 - using RARP or DHCP to set 9

K

- keyboard shortcuts 15

L

- local spare assignment 177
- local spare assignment (FC and SATA) 102
- local spare assignment (SCSI) 64
- local spare drives 5
 - deleting 178
- logical drive, initialization mode when creating (FC and SATA) 104
- logical drive, initialization mode when creating (SCSI) 66
- logical drives
 - adding physical drives to 137
 - assigning a local spare (FC and SATA) 102
 - assigning a local spare (SCSI) 64
 - assigning RAID level (FC and SATA) 100
 - assigning RAID level (SCSI) 61
 - changing controller assignment (FC and SATA) 108
 - changing controller assignment (SCSI) 70
 - checking parity 139
 - copying 143
 - creating (FC and SATA) 97, 99
 - creating (SCSI) 59, 61
 - default parameters 340
 - defining 312
 - deleting 51
 - deleting (FC and SATA) 83
 - deleting a partition 130
 - enabling and disabling parity reporting 141
 - expanding the capacity of existing 133
 - ID 294
 - larger than 253 Gbyte (FC and SATA) 98
 - larger than 253 Gbyte (SCSI) 60
 - LG number 294
 - mapping a partition to a LUN (SCSI) 116
 - mapping partitions (FC and SATA) 112
 - mapping partitions (SCSI) 74
 - mapping partitions to LUNs 160
 - mapping partitions to LUNs (SCSI) 75
 - maximum number of disks per 25
 - maximum number per configuration 24
 - maximum physical capacity (FC and SATA) 102
 - maximum physical capacity (SCSI) 64
 - maximum usable capacity 25
 - media scan for bad blocks 145
 - naming (FC and SATA) 109
 - naming (SCSI) 71
 - partitioning (FC and SATA) 109

- partitioning (SCSI) 72
- preparing for larger than 253 Gbyte 237
- RAID level 294
- rebuild priority 246
- rebuilding 132, 306
- regenerating the parity of 140
- replacing larger-capacity drives in 143
- saving NVRAM configuration 281
- size 294
- status table 293
- viewing configuration (FC and SATA) 82
- viewing configuration (SCSI) 50
- logical volumes
 - avoiding failure 153
 - creating 155
 - default parameters 340
 - defining 313
 - deleting 158
 - deleting partitions of 153
 - expanding 153, 158
 - limitations 152
 - partitions 153
 - partitions, deleting 153
 - RAID levels 154
 - spare drives 154
 - status table 155
 - understanding 152
- loop connection 97
- loop DAS configuration 41
- loop options 238
- loopback error statistics 273
- low-level formatting of a physical drive 195
- LUNs
 - adding WWN entries 168
 - creating 1024 in loop mode 163
 - creating host filters for (FC and SATA) 118
 - default parameters 340
 - deleting a mapping 165
 - labelling (FC and SATA) 123
 - labelling (SCSI) 78
 - mapping (FC and SATA) 112
 - mapping (SCSI) 74
 - mapping logical drive partitions to (SCSI) 75
 - mapping partitions to 160
 - maximum concurrent host connections 234
 - maximum number 24
 - number of tags reserved for 235
 - number supported per host ID 234
 - overview (FC and SATA) 113
 - overview (SCSI) 74
 - Solaris device files for (FC and SATA) 124
 - Solaris device files for (SCSI) 79
- M**
 - main menu, firmware 15
 - manual rebuild 308
 - mapping
 - cylinders, heads, and sectors 236
 - mapping partitions (FC and SATA) 112
 - mapping partitions (SCSI) 74
 - maximum drive capacity (FC and SATA) 102
 - maximum drive capacity (SCSI) 64
 - maximum number of drives 24
 - maximum queued I/O count 233
 - maximum tag count 242
 - maximum usable capacities 25
 - maximum usable capacity per logical drive 25
 - media scan
 - aborting 146
 - continuous 146
 - for bad blocks in drives 145
 - one-time 146
 - priority 146
 - terminating 146
 - using on individual drives 193
 - messages
 - controller events 354
 - drive 359
 - event 353
 - logical drive events 367
 - messages, event types 353
 - messages. See also event messages
 - MISSING drive status 297
 - mixed drive support, enabling and disabling (FC and SATA) 87
- N**
 - naming a controller 248
 - naming a logical drive (FC and SATA) 109
 - naming a logical drive (SCSI) 71
 - navigation keys 15
 - network protocol support 215
 - NEW DRV drive status 297

- node name 176
- NVRAM
 - restoring configuration from disk 284
 - saving configuration to disk 281
 - saving configuration to disk (FC and SATA) 125
 - saving configuration to disk (SCSI) 80
- O**
- offline initialization (FC and SATA) 104
- offline initialization (SCSI) 66
- online initialization (FC and SATA) 104
- online initialization (SCSI) 66
- optimization modes
 - (FC and SATA) 84
 - (SCSI) 52
 - random and sequential (FC and SATA) 85
 - random and sequential (SCSI) 53
 - verifying and changing (FC and SATA) 85
 - verifying and changing (SCSI) 53
- over-temp controller shutdown 268
- P**
- parameters
 - communication 212
 - configuration 211
 - controller 248
 - disk array 246
 - drive-side 240
 - host-side 233
 - host-side (FC and SATA) 97
 - peripheral device type 236
 - redundant controller 247
 - summary of default 339
- parity
 - checking 139
 - enabling and disabling reporting inconsistent 141
 - enabling SCSI bus byte 204
 - overwriting inconsistent 140
 - regenerating 140
- partition
 - mapping of logical drive to a LUN (SCSI) 75
- partitioning logical volumes 153
- partitions
 - deleting logical drives 130
 - deleting logical volume 153
 - logical drive (FC and SATA) 109
 - logical drive (SCSI) 72
 - mapping of logical drives to LUN (SCSI) 116
 - mapping to LUNs 160
 - maximum number per logical drive 24
- password
 - changing 277
 - disabling 278
 - setting new 276
- password validation timeout 249
- patch downloads 302
- periodic drive check time 243
- periodic drive swap check time 244
- periodic SAF-TE and SES check time 243
- peripheral device type parameters 236
- peripheral devices 253
 - default parameters 348
 - event trigger operations 267
 - redundant controller mode 265
 - SAF-TE status 261
 - viewing status 253
- perpetual clone of a failing drive 186
- physical drives
 - cloning a failing 184
 - cloning status 189
 - flashing to identify good or bad 181
 - periodic check time 243
 - reformatting 195
 - replacing after cloning 184
 - reserved space 197
 - revision number, serial number, and disk capacity 176
 - scanning SCSI 179
 - size and speed 296
 - status table 172, 296
 - viewing information 176
- ping, enabling and disabling 215
- planning
 - host applications 4
 - RAID configurations 4
- PLD, firmware upgrading 304
- point-to-point connection 97
- point-to-point options 238
- point-to-point, sample SAN configuration 35
- primary controller SCSI IDs 202
- progress indicators 18
- protocol support, setting network 215

R

RAID

- advantages over non-RAID 316
- multilevel (logical volumes) 152
- planning considerations 4
- terminology 312

RAID 0 319

RAID 1 319

RAID 1+0 320

- concurrent rebuild in 310

RAID 3 321

RAID 3+0 323

RAID 5 322

RAID 5+0 323

RAID levels

- 1+0, 3+0, 5+0, 5+1, 5+5 154
- 10, 30, 50 154
- assigning (FC and SATA) 100
- assigning (SCSI) 61
- defining 316
- descriptions 318
- mirroring and striping 317
- number of drives per 317
- redundancy in 317

RARP 213

- setting array IP address using 214
- setting automatic IP address for array 9

RCCOM (FC and SATA) 247

read/write testing 196

rebuild priority 246

rebuild, concurrent in RAID 1+0 310

rebuilding logical drives 306

rebuilding logical drives automatically 306

rebuilding logical drives manually 308

records, keeping 331

redundant controller parameters 247

reserved space on drives 197

reset button, use in silencing alarms 276

resetting controller 278

restoring NVRAM configuration from disk 284

Reverse Address Resolution Protocol. See RARP

RS-232 port configuration 212

S

SAF-TE

check time 243

temperature sensor locations 264

viewing status 261

SAN configuration example (FC/SATA) 35

SATA arrays, configuration and reconfiguration of 81

saving configuration to disk 281

SB-MISS drive status 297

scanning

drive media 145

for bad blocks with media scan 193

scanning drives (SCSI) 179

SCSI arrays

configuration and reconfiguration of 49

SCSI bus byte

enabling parity 204

SCSI drive utilities

low-level format 195

SCSI Enclosure Services. See SES

SCSI transfer width 204

secondary controller SCSI IDs 202

sector ranges, changing (FC and SATA) 98

sector ranges, changing (SCSI) 61

Self-Monitoring Analysis and Reporting Technology.

See SMART

sensor locations

SAF-TE temperature 264

SES temperature 259

SES

enclosure information 176

firmware upgrading 304

temperature sensor locations 259

SES check time 243

SES status, viewing 254

settings, how to record 331

shelf life 292

shutting down the controller 279

single-bus configuration (SCSI) 173

size of drive 296

SMART 183

disabling detection 193

how cloning options work 191

testing a drive 192

SMART detection

enabling 191

- SMART technology defined 190
- SNMP
 - configuring 226
 - enabling and disabling 215
- SNMP agent.ini file example 228
- SNMP traps, enabling 227
- Solaris device files for LUNs (FC and SATA) 124
- Solaris device files for LUNs (SCSI) 79
- spare drives
 - assigning a global 178
 - assigning a local 177
 - deleting 178
 - global 5
 - local 5
 - logical volumes 154
- specifications, firmware 325
- speed of drive 296
- split-bus configuration (SCSI) 173
- SSH, enabling and disabling 215
- STAND-BY drive status 297
- status table
 - channel 298
 - host and drive channels 200
- status windows
 - checking 293
 - logical drive table 293
 - physical drive status table 296
- stripe size
 - configuring (FC and SATA) 105
 - configuring (SCSI) 67
- Sun StorEdge 3000 family, product descriptions 2
- sync transfer clock speed (SCSI) 203
- system functions 275
 - default parameters 350

T

- tag command queuing 242
- tag count, maximum 242
- tags, number reserved per host LUN 235
- telnet
 - enabling and disabling 215
 - inactivity timeout 216
- temperature
 - exceeds threshold trigger 268
- temperature limit 268

- temperature sensor locations
 - SAF-TE 264
 - SES 259
- temperature status, controller 270
- terminology, basic 312
- terminology, navigation 17
- threshold trigger
 - temperature exceeds 268
- timeout
 - password validation 249
 - telnet inactivity 216
- transfer clock speed, setting (SCSI) 203
- transfer rate indicator 12
- transfer width (SCSI) 204
- troubleshooting 305

U

- unique identifier, controller 250
- units of device capacity 19
- upgrading firmware 301

V

- verification on writes 247
- voltage status, controller 270

W

- worldwide name. See also WWNs
- worldwide names, finding 166
- write errors, avoiding 247
- write policy 231
 - guidelines 32
 - selecting (FC and SATA) 104
 - selecting (SCSI) 66
- write-back cache 293
- write-back cache, enabling and disabling 231
- write-through cache, enabling and disabling 231
- WWN. See also worldwide name
- WWNs 176
 - manually adding entries 168
 - selecting while creating a host filter 118